Quantum mechanics

Predict measurements from quantum states, approximate confined energies, and determine how driving, exchange and loss of phase change observable spectra.

Read the subject in depth

Each volume develops its arguments in sequence, with worked problems and links to experiments.

- ### Quantum states, dynamics and measurement

  Build the connection between preparation, time evolution and measurement. Develop the algebra, approximations and scattering theory needed to predict a spectrum or a set of count rates.

- ### Computational quantum experiments

  Follow a quantum state from its preparation to observable frequencies, correlations and counting statistics. Test how basis size, finite sampling and detector settings change each inference.

Advanced study across the subject

Each sequence connects the course to detailed chapters in neighbouring subjects. The complete course PDF includes their derivations, examples and exercises.

Use the angular momentum and central-potential material with the radial quantisation chapter. The approximation methods provide a second route into the coupled-oscillator problem. Identical particles connect those methods to the equilibrium strand. Torsten Fließbach · Quantenmechanik, 6th edition

1. ### Translational symmetry, bands and edges

   Apply a translation operator, reduce a many-site Hamiltonian to a small matrix and use the resulting eigenvectors to predict an edge state. Compare a band touching with a gapped system.

2. ### Matter and fields in quantum equilibrium

   Derive quantum occupations from particle counting. Use them to predict responses of electrons, acoustic modes and photons. Look for the boundary between a one-particle and a collective description.

3. ### From a field to an observable spectrum

   Separate orbital and spin responses, count magnetic states and compare a quantum prediction with the corresponding orbit. Revisit the assumptions behind a nondissipative screening current.

4. ### Approximate an eigenstate and check what it predicts

   Use projection, orthogonalisation and energy stationarity to construct approximations. Compare them with phase-integral and explicitly propagating pictures of the same mathematical operators.

How the course fits together

1. **Chapters 1–4**

   Normalisation turns a complex amplitude into probabilities. Match amplitudes and currents across interfaces. Operator eigenstates then connect a measurement to oscillator and atomic spectra.

2. **Chapters 5–8**

   Mix states through perturbations or by combining subsystems. Variational trial states supply energy bounds when exact eigenfunctions are unavailable. A phase integral gives a different approximation in the semiclassical limit.

3. **Chapters 9–12**

   Drive transitions and impose exchange symmetry before tracing out unobserved degrees of freedom. The final billiard problem asks which features of a confining boundary its measured spectrum can reveal.

Contents

1. Chapter 1 · Wave packets and probability current
2. Chapter 2 · Scattering and tunnelling in one dimension
3. Problem sheet 1 · Flux, packets and barriers
4. Chapter 3 · Operators and the quantum harmonic oscillator
5. Chapter 4 · Angular momentum and hydrogenic atoms
6. Problem sheet 2 · Quantised vibrations and atomic states
7. Chapter 5 · Perturbations, degeneracy and level repulsion
8. Chapter 6 · Composite states and spin correlations
9. Problem sheet 3 · Small perturbations and paired measurements
10. Chapter 7 · Estimating a ground-state energy
11. Chapter 8 · Semiclassical quantization
12. Problem sheet 4 · Energy bounds and semiclassical actions
13. Chapter 9 · Transitions under a timed drive
14. Chapter 10 · Exchange and two-particle energies
15. Problem sheet 5 · Driven transitions and exchange symmetry
16. Chapter 11 · Density matrices and loss of phase
17. Chapter 12 · Confinement and spectral geometry
18. Problem sheet 6 · Coherence and the confinement spectrum
19. Chapter 13 · When a free quantum packet initially gets narrower
20. Chapter 14 · Boundary conditions that carry a current around a ring
21. Chapter 15 · Why the Coulomb radial series terminates
22. Chapter 16 · Measure an interaction by its low-energy scattering length
23. Chapter 17 · A phase accumulated by slowly turning a magnetic field
24. Chapter 18 · A sudden expansion and the distribution of measured energy
25. Chapter 19 · Adiabatic following and a finite sweep through an avoided crossing
26. Chapter 20 · What a spin echo can recover
27. Problem sheet 7 · Measuring two phases with one interferometer
28. Problem sheet 8 · How a pair of neutral oscillators acquires an attractive interaction
29. Quantum states, dynamics and measurement
30. Computational quantum experiments
31. Translational symmetry, bands and edges
32. Matter and fields in quantum equilibrium
33. From a field to an observable spectrum
34. Approximate an eigenstate and check what it predicts

35. Applied study
36. Notation and assumptions
37. Further reading

Chapter 1 · Wave packets and probability current

Normalise a quantum wave packet and derive a local balance law that predicts the passage of probability through a detector plane.

FROM FIRST PRINCIPLES

Derive local probability balance from Schrödinger evolution

Take a spinless particle on a line with constant mass and a real local potential $V(x,t)$. The inputs are the Schrödinger equation as a dynamical postulate and the Born rule $\rho = |\psi|^2$ as the probability density. To see how probability moves, differentiate $\psi^*\psi$ and substitute the time derivative of the state and its complex conjugate:

$$i\hbar\partial_t\psi = -\frac{\hbar^2}{2m}\partial_x^2\psi + V\psi,$$

$$\partial_t\rho = \frac{i\hbar}{2m}(\psi^*\partial_x^2\psi - \psi\partial_x^2\psi^*).$$

The potential terms cancel because V is real. The remaining difference is a spatial derivative: applying the product rule makes the two first-derivative cross terms cancel. This identifies the current and its sign:

$$\partial_t\rho + \partial_x j = 0,$$

$$j = \frac{\hbar}{2mi}(\psi^*\partial_x\psi - \psi\partial_x\psi^*).$$

Integrate over a detector interval $[a,b]$. The probability inside changes at rate $j(a) - j(b)$, giving direct physical meaning to the local derivative. Vanishing exterior current conserves the total normalization. A complex absorbing potential would add a local sink and requires a different balance.

Background · State normalisation

$$\int |\psi(x)|^2\,dx = 1$$

For a normalisable one-dimensional state, squared modulus is probability per length. Stationary incoming scattering states require a flux or continuum normalisation.

DERIVATION AND APPLICATIONS

Predict the minimum width of a freely propagating packet

Two spatial snapshots of a beam can differ even when no forces act. To calculate this spreading without evaluating a Fourier integral, work with the position and momentum operators for a free, nonrelativistic particle of mass m. Assume motion along a line, a normalized state with finite second moments, and

negligible interaction during flight. In the Heisenberg picture $H = p^2/(2m)$ is constant, momentum is conserved and velocity is p/m. Integrating the operator equations gives

$$p(t) = p(0),$$
$$x(t) = x(0) + \frac{t}{m}p(0).$$

Subtract the mean trajectory. The square of $\delta x(t) = \delta x(0) + (t/m)\delta p(0)$ contains both product orders because position and momentum do not commute. Defining $C_{xp} = \langle \delta x\, \delta p + \delta p\, \delta x \rangle/2$, the exact variance is

$$\sigma_x^2(t) = \sigma_x^2(0) + \frac{2t}{m}C_{xp} + \frac{t^2}{m^2}\sigma_p^2(0).$$

Now specialize to an unchirped Gaussian at release, with standard deviation s in position, constant mean phase gradient and no quadratic phase. Then $C_{xp} = 0$ and the momentum standard deviation is $\hbar/(2s)$. Substitution gives

$$\sigma_x(t) = s\sqrt{1 + (t/\tau)^2},$$
$$\tau = 2ms^2/\hbar.$$

The characteristic time marks the point at which propagation and initial preparation contribute equally to the variance. A narrower launch has a larger momentum spread and reaches this crossover sooner. The general variance equation also shows why free evolution can initially narrow a correlated preparation with $C_{xp} < 0$.

WORKED EXAMPLE

Predict an electron-beam width at the detector

Change the inputs and investigate this derivation.

Take $m = 9.109 \times 10^{-31}$ kg, $s = 20.0$ nm and $\hbar = 1.055 \times 10^{-34}$ J s. The spreading time is 6.91 ps. After 10.0 ps the width is

$$20.0\sqrt{1 + (10.0/6.91)^2} = 35.2\,\mathrm{nm}.$$

An overall boost translates the centroid while leaving this predicted width unchanged.

SECOND APPLICATION

Choose the launch width for the smallest final spot

Change the inputs and investigate this derivation.

At fixed flight time t, minimize $s^2 + (\hbar t/2m)^2/s^2$ over unchirped Gaussian preparations. Differentiating with respect to s^2 gives $s^2 = \hbar t/(2m)$ at the minimum. For the same electron and 10.0 ps, the best launch width is 24.1 nm and the final width is $\sqrt{2}$ times larger, 34.0 nm. Tightening the initial spot below that value makes this detector spot wider.

Start from this chapter. A hard wall sets the wavefunction to zero. Ask which stationary spatial modes satisfy this condition on a polygon, and how changing the length scale changes their quantum energies.

OneSimulation laboratory · Turn spatial eigenvalues into confinement energies

Specify a preparation

Prepare many independent particles of the same mass in the same one-dimensional state and record one position from each preparation. The frequencies in small position bins reveal a probability density. Quantum theory assigns a complex wavefunction $\psi(x,t)$ to a pure state. The probability to record a particle in an interval and the normalisation of the preparation are

$$P(a<x<b;t)=\int_a^b |\psi(x,t)|^2\,dx,$$
$$1=\int_{-\infty}^{\infty} |\psi(x,t)|^2\,dx.$$

The density $|\psi|^2$ has units of inverse length and can numerically exceed one. Integrating it over a position interval gives a dimensionless probability between zero and one.

A single position histogram fixes the modulus of the wavefunction but contains no direct information about its phase. A global phase is unobservable, whereas a phase that varies across space changes the momentum distribution and subsequent evolution. Two preparations with the same initial histogram can reach a detector at different times. To predict how probability moves we need the phase information together with an equation of motion.

Obtain a continuity equation

For a particle in a real scalar potential $V(x,t)$, the Schrödinger equation is

$$i\hbar\partial_t\psi=-(\hbar^2/2m)\partial_{xx}\psi+V\psi.$$

Multiply by ψ^*, subtract the conjugate equation multiplied by ψ and use the product rule. The two potential terms cancel because V is real. The result takes the conservative form

$$\partial_t|\psi|^2+\partial_x j=0,$$
$$j=\frac{\hbar}{2mi}(\psi^*\partial_x\psi-\psi\partial_x\psi^*).$$

Integrated between two fixed points, the rate of change of probability inside equals the inward current at the left minus the outward current at the right. In this one-dimensional normalisation j has units of inverse seconds. Vanishing current at both infinities ensures that a normalised state stays normalised even if the shape of its density changes appreciably.

Writing $\psi=\sqrt{\rho}\,e^{i\varphi}$ gives $j=(\hbar/m)\rho\,\partial_x\varphi$ where the phase is regular. For an initially real packet the current vanishes at that instant. A spatial phase ramp e^{ikx} leaves ρ unchanged yet gives current $\hbar k\rho/m$. Current is a

signed net flux. Interpreting every crossing of a point as a first arrival requires additional information about the detector and possible returns.

Background · Does a common phase leave every measurement unchanged?

Multiplying a wavefunction by a spatially constant phase preserves the density and the current. Now prepare a second wavepacket with the spatial phase ramp $\exp(ik_0x)$ in a region free of vector potentials. Both packets give the same position histogram at the preparation time. The derivative in the probability current also differentiates the ramp.

$$j[\mathrm{e}^{ik_0x}\psi] = j[\psi] + \frac{\hbar k_0}{m}|\psi|^2.$$

After free expansion the second packet's centroid has the added speed $\hbar k_0/m$. A spatial phase gradient is measurable through a later position histogram even when the initial density is identical. For a charged particle in a magnetic vector potential the mechanical current includes that potential too. A joint gauge change of wavefunction and vector potential leaves the physical current unchanged.

Count momenta as well as positions

Use the Fourier convention $\psi(x) = (2\pi)^{-1/2}\int \tilde{\psi}(k)e^{ikx}dk$. Parseval's identity preserves the normalisation. The probability in a wave-number interval is $|\tilde{\psi}(k)|^2dk$, with momentum $p = \hbar k$. A Gaussian packet having density variance σ_x^2 at preparation can be written as

$$\psi = (2\pi\sigma_x^2)^{-1/4}\exp[-(x-x_0)^2/(4\sigma_x^2) + ik_0x].$$

Fourier transforming this Gaussian gives mean momentum $\hbar k_0$ and standard deviation $\hbar/(2\sigma_x)$. Tighter initial localisation broadens the momentum distribution.

For a free particle the plane-wave components advance with $\omega(k) = \hbar k^2/(2m)$. Different wave numbers give different group velocities. An initially minimal, unchirped Gaussian spreads to $\sigma_x(t) = \sigma_x\sqrt{1 + [\hbar t/(2m\sigma_x^2)]^2}$. Transferring this expression to an initially focused or phase-curved packet requires its position–momentum correlation as well as its initial width.

When does a wave packet follow Newton's equation?

A position measurement on a single particle produces one draw from its distribution. Suppose we repeat the preparation at successive flight times. What determines the motion of the measured centroid? Write the ensemble mean of a Hermitian observable as $\langle A\rangle = \langle\psi|A|\psi\rangle$. Differentiate both states and the operator. The Schrödinger equation gives $|\dot{\psi}\rangle = -iH|\psi\rangle/\hbar$, while its adjoint gives the opposite sign on the bra. Thus

$$\frac{d\langle A\rangle}{dt} = \frac{i}{\hbar}\langle[H, A]\rangle + \left\langle\frac{\partial A}{\partial t}\right\rangle.$$

For a particle on a line let $H = p^2/(2m) + V(x)$, with a real differentiable potential and states decaying sufficiently rapidly to drop boundary terms. Since $[p^2, x] = -2i\hbar p$ and $[V, p] = i\hbar V'(x)$, the first moment satisfies

$$m\frac{d^2\langle x\rangle}{dt^2} = -\langle V'(x)\rangle.$$

A uniform force accelerates every part of the density equally. In a harmonic trap the force is linear, so averaging it gives the force at the centroid for a packet of any width. An anharmonic force samples higher moments. The centroid and width can then influence each other even when no external noise is present.

Consider a quartic trap $V = \lambda x^4/4$ with $\lambda = 1.00\,\mathrm{meV\,nm^{-4}}$. At one instant the position density is Gaussian, centred at $a = 2.00\,\mathrm{nm}$ with standard deviation $s = 1.00\,\mathrm{nm}$. Writing $x = a + \xi$ eliminates odd central moments, giving $\langle x^3\rangle = a^3 + 3as^2 = 14.0\,\mathrm{nm^3}$. The mean force is therefore

$$-14.0\,\mathrm{meV/nm} = -2.24\,\mathrm{pN}.$$

Evaluating the force only at the centroid gives $-8.00\,\mathrm{meV/nm}$. The high-force side of the distribution contributes disproportionately. As $s/|a| \to 0$ the fractional correction $3s^2/a^2$ vanishes. This is an instantaneous prediction. Evolution in a quartic trap need not preserve a Gaussian shape.

Practise on sheet 1 · Infer a force from the motion of a centroid. Further reading · Rutgers University · Intermediate Quantum Mechanics · Lecture 9, The Ehrenfest Theorem.

Predict a free-flight measurement

An electron has initial width 100 nm and mean wave number 1.0×10^8 m⁻¹. Its mean speed is $\hbar k_0/m_e = 1.16 \times 10^4$ m/s and its velocity spread is about 579 m/s. After 1.0 ns the centre has moved 11.6 µm. For the unchirped preparation the width grows to $\sqrt{(0.100\,\mu\mathrm{m})^2 + (0.579\,\mu\mathrm{m})^2} = 0.588$ µm. The packet is well inside the nonrelativistic regime.

The solution predicts two independent observables: the motion of the centre and the broadening about that centre. Reversing the phase ramp reverses the former while leaving the latter unchanged. The initial position histograms for the two preparations are identical. A downstream record is needed to discriminate their motion.

One can measure spreading by repeating the preparation and recording the position distribution after several flight times. For an initially uncorrelated position and momentum, the free-particle variance is

$$\sigma_x^2(t) = \sigma_x^2(0) + t^2\sigma_p^2/m^2.$$

This follows from the operator equation $x(t) = x(0) + p(0)t/m$. A known initial position–momentum covariance adds a term linear in time and can even make the packet contract initially. Measuring a minimum downstream width alone therefore does not establish the width or the momentum uncertainty at release.

Infer current from two interfering momentum components

An interference pattern can make density strongly position dependent even when probability current is uniform. Put a free particle on a ring of circumference L and use periodic boundary conditions. Two states with wavenumbers $+k$ and $-k$, where $k = 2\pi n/L$ and n is a nonzero integer, have the same energy $E = \hbar^2 k^2/(2m)$. A normalized superposition has wavefunction $\psi = [a e^{ikx} + b e^{-ikx}] e^{-iEt/\hbar}/\sqrt{L}$ with $|a|^2 + |b|^2 = 1$. Integration around the ring makes the two opposite travelling waves orthogonal.

Squaring the amplitude gives a density containing a constant term and a spatial fringe,

$$\rho = [1 + 2\mathrm{Re}(a b^* e^{2ikx})]/L.$$

To find current, differentiate ψ before taking the imaginary part of $\psi^* \partial_x \psi$. The two diagonal products contribute with opposite signs. The cross terms in that product are real, so their contribution to current vanishes. Thus

$$j = \frac{\hbar k}{mL}(|a|^2 - |b|^2),$$

$$\partial_t \rho = -\partial_x j = 0.$$

The relative weights determine the net rate of passage past a fixed point; the relative phase shifts the fringe positions without changing that rate. Equal weights give a standing wave with stationary nodes and zero current. An unequal pair can have a stationary density pattern with nonzero circulating current, because every segment loses probability at one end at exactly the rate it gains it at the other.

Distinguish two preparations with one density pattern

Take an electron on an ideal ring of circumference 1.00 µm with $n = 2$, so $k = 1.257 \times 10^7$ m⁻¹. Choose $a = \sqrt{0.80}$ and $b = \sqrt{0.20}$ real. Fringe spacing is $\pi/k = 0.250$ µm and normalized density ranges from $0.20/L$ to $1.80/L$. The velocity scale is

$$\hbar k/m_e = 1.46 \times 10^3\,\mathrm{m/s}.$$

The probability current is this speed times $0.60/L$, or 8.73×10^8 s⁻¹. These units describe crossings per unit time for a normalized one-dimensional state rather than probability density per metre.

Interchange the weights assigned to clockwise and anticlockwise components while keeping their relative phase fixed and real. Density everywhere on the ring is identical but current reverses sign. A position histogram alone therefore cannot resolve which preparation was made. A momentum measurement over repeated preparations would instead return the two momenta with swapped probabilities. For a physical charged particle, multiplication by its charge gives the circulating electric current, whose sign is also reversed relative to probability current for an electron. A weak magnetic probe couples to that circulation rather than to the fringe visibility.

A relative phase change offers another comparison. Multiplying the anticlockwise amplitude by $e^{i\alpha}$ moves the density pattern around the ring by $\alpha/(2k)$ with the sign fixed by the chosen wave convention. At $\alpha = \pi$ a previous maximum becomes a minimum, yet neither the probabilities of the two momentum outcomes nor the current changes. Conversely replacing the coherent preparation by a random mixture with the same two

momentum probabilities leaves the circulating current intact while removing the interference fringes. Current, momentum populations and position fringes are therefore independent pieces of information about even this simple state.

Follow a freely spreading wave packet

A freely evolving minimum-uncertainty Gaussian starts at rest with position width σ_0. Its momentum spread is $\hbar/(2\sigma_0)$. Free motion gives $\sigma(t)^2 = \sigma_0^2 + [\hbar t/(2m\sigma_0)]^2$ in the absence of an initial position–momentum correlation.

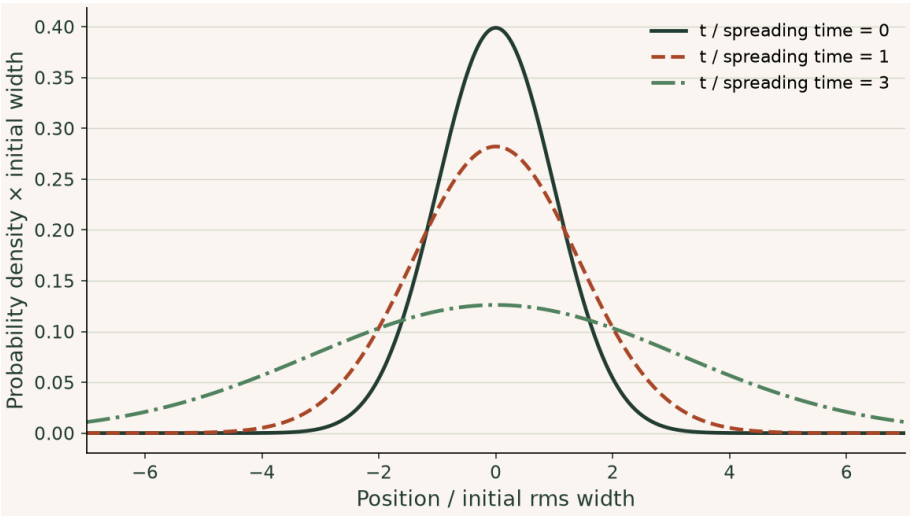

After one spreading time the rms width grows by √2 and the central probability density falls by the same factor. The centre has not moved. A probability density becoming smaller at the origin does not imply loss of total probability.

Try it · Gaussian packet

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Elapsed time / spreading time 1

Initial rms width (nm) 1

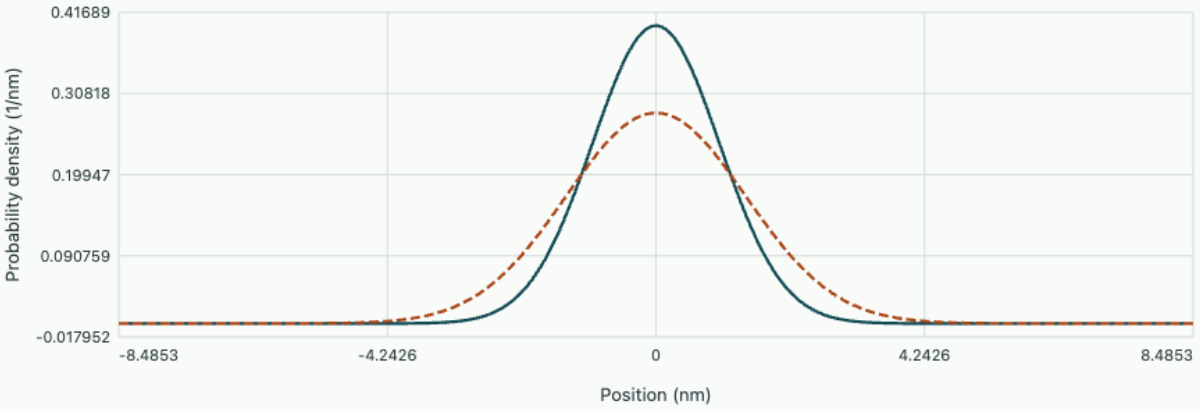

Rms width 1.4142 nm. The full spatial integral stays one; spreading does not reduce total probability.

▶ **Read values from the plot**

Study the full derivation

From two-path interference to the geometry of a quantum state · Quantum states, dynamics and measurement

Two preparations give the same counts at one analyser. Place a second analyser in a rotated basis and the counts separate. From this experiment develop amplitudes, phase, projective measurement, mixed states and the geometry of distinguishability.

Chapter 2 · Scattering and tunnelling in one dimension

Match stationary waves across a potential discontinuity and use probability current to distinguish reflection, propagation and an evanescent field.

FROM FIRST PRINCIPLES

Obtain reflected and transmitted probabilities at a potential step

Suppose a constant-mass particle reaches a flat potential step from the left with energy above both sides. In a region of constant potential the stationary Schrödinger equation has waves $e^{\pm ikx}$ with $\hbar^2k^2/2m = E - V$. Set the incoming amplitude to one and use asymptotic waves $e^{ik_1x} + re^{-ik_1x}$ at left and te^{ik_2x} at right. Integration of the differential equation across the finite step makes the derivative continuous; the wave function itself is continuous as well:

$$1 + r = t,$$
$$k_1(1 - r) = k_2 t.$$

Solving the two linear conditions gives

$$r = (k_1 - k_2)/(k_1 + k_2),$$
$$t = 2k_1/(k_1 + k_2).$$

The current from the previous chapter for one travelling component is $\hbar k|A|^2/m$ in its propagation direction. Ratios to incoming current therefore yield

$$R = \frac{(k_1 - k_2)^2}{(k_1 + k_2)^2},$$
$$T = \frac{4k_1k_2}{(k_1 + k_2)^2},$$
$$R + T = 1.$$

Equal potentials give unity transmission. Near a right-hand threshold $k_2 \to 0$, a finite boundary amplitude can coexist with vanishing transmitted flux because the outgoing velocity also tends to zero.

Background · Probability current

$$j = (\hbar/m)\,\mathrm{Im}(\psi^* \partial_x \psi)$$

This expression assumes a scalar particle with no vector potential. Transmission probabilities require outgoing relative to incoming current, even when their wave numbers differ.

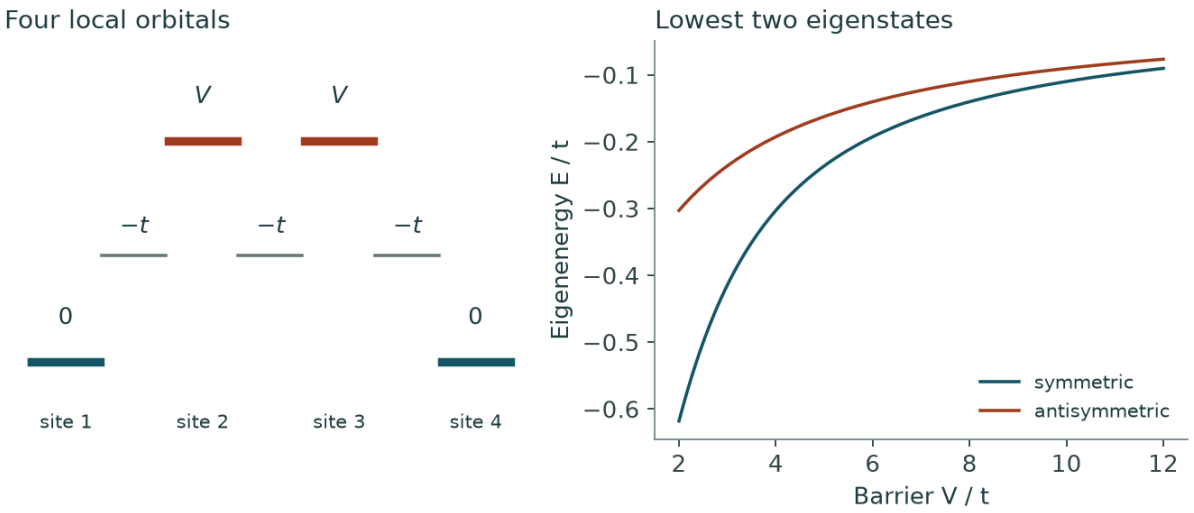

From a barrier drawing to a measurable splitting

Put the four on-site energies on the diagonal of a matrix and put –t between neighbours. An end-to-end transition visits both central sites. At large V it requires three hops and two off-resonant denominators, so its amplitude scales as t^3/V^2. The graph shows the two exact low energies whose difference tests that prediction.

Try it · Physics lab double well splitting

Explore the same physical relation with adjustable inputs. This independent parameter study has its own displayed reference settings; compare its assumptions and units before matching the static example.

Open this live calculation

Try the calculation

Middle site energy V (meV) **4**

Hopping t (meV) **1**

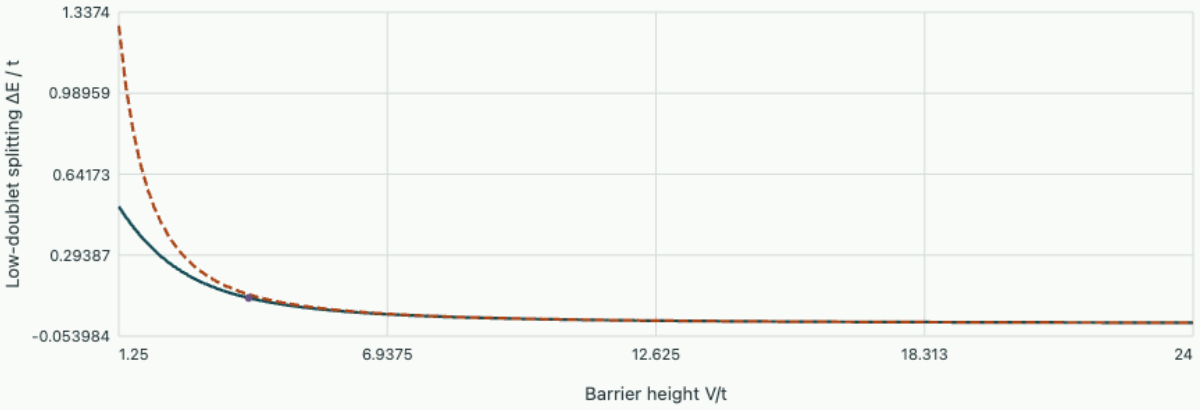

Splitting 0.11019 meV. First transfer between dressed end states at 18.766 ps.

Download the OneSimulation model

In your OneSimulation environment, run the model with these inputs.

```
python spectra-and-orbits.py wells --barrier 4 --hopping 1
```

DERIVATION AND APPLICATIONS

Match both faces of a finite barrier

A decaying wave inside a classically inaccessible interval does not alone determine its transmission. Boundary matching at the exit also contributes. Let $V(x) = V_0$ for $0 < x < L$ and zero outside, with a particle of constant mass and $0 < E < V_0$ incident from the left. Define $k = \sqrt{2mE}/\hbar$ and $\kappa = \sqrt{2m(V_0 - E)}/\hbar$. Set the phase reference at the exit so the outgoing wave is $te^{ik(x-L)}$. At $x = L$, continuity gives both the exit value t and its derivative ikt.

Inside, $\psi'' = \kappa^2\psi$. Propagate these exit data back to the entrance using the even and odd solutions:

$$\psi(0) = t[\cosh(\kappa L) - i(k/\kappa)\sinh(\kappa L)],$$
$$\psi'(0) = t[-\kappa\sinh(\kappa L) + ik\cosh(\kappa L)].$$

On the left write $e^{ikx} + re^{-ikx}$. Adding its value to its derivative divided by ik removes the reflected amplitude. Inserting the propagated data yields

$$2 = t\,[2\cosh(\kappa L) + i(\kappa/k - k/\kappa)\sinh(\kappa L)].$$

Both asymptotic waves propagate in the same material, so their velocity factors cancel in the current ratio. Take the squared magnitude and use $\cosh^2 u = 1 + \sinh^2 u$:

$$T = |t|^2 = \left[1 + \frac{(k^2+\kappa^2)^2}{4k^2\kappa^2}\sinh^2(\kappa L)\right]^{-1}.$$

A barrier of vanishing width transmits fully. An exponential estimate at large κL neglects the interface-dependent prefactor. Close to the top, the $\sinh$ factor and its apparently singular coefficient must instead be taken to the limit together.

WORKED EXAMPLE

Infer a width from a tunnelling measurement

Change the inputs and investigate this derivation.

At $E = V_0/2$, $k = \kappa$ and the result simplifies to $T = \operatorname{sech}^2(\kappa L)$. A device with $\kappa = 2.00$ nm⁻¹ has $T = 0.250$. Thus $\cosh(\kappa L) = 2$ and $L = \operatorname{arcosh}(2)/2.00 = 0.658$ nm. Doubling that width gives $\cosh(2\kappa L) = 2(2)^2 - 1 = 7$, hence

$$T = 1/49 = 0.0204.$$

SECOND APPLICATION

Approach the top from below

Change the inputs and investigate this derivation.

Let $E \to V_0$ at fixed L. With $\sinh(\kappa L) \simeq \kappa L$ the transmission becomes

$$T = [1 + (kL)^2/4]^{-1}.$$

For $k = 3.00$ nm⁻¹ and $L = 0.500$ nm this is 0.640. A vanishing inverse decay length therefore does not guarantee complete transmission. Exactly at the barrier top the internal solution becomes linear and still has to match the slopes on both sides.

Try it · Barrier transmission

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Barrier height / incident energy 2

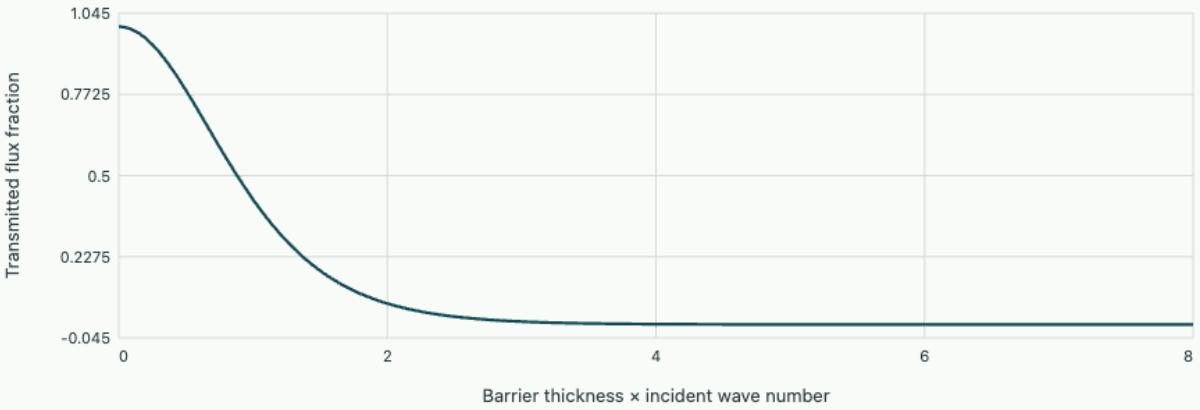

Exact rectangular barrier

At threshold V0=E the limiting formula is 1/[1+(kL)²/4]. Below threshold repeated internal reflections give resonances; above it an evanescent region suppresses transmission.

Build a scattering state

Send a narrow band of particles from the left towards a region where the potential changes. Far from the obstacle a stationary energy component has incoming and outgoing travelling waves. Such plane waves are not normalisable on the whole line, but their flux ratios describe what happens to a sufficiently long incident packet whose momentum distribution is narrowly centred. Assume a particle of constant mass and a real, finite, time-independent potential.

The stationary Schrödinger equation is

$$-\hbar^2\phi''/(2m) + V\phi = E\phi.$$

Integrating it over a shrinking interval around a finite step gives continuity of ϕ'; a finite jump in ϕ itself would introduce a more singular derivative, so ϕ is continuous as well. These matching laws change for an ideal delta potential, an infinite wall or an effective mass discontinuity. Stating the model at an interface matters as much as solving the constant-potential regions.

Partly reflect above a potential step

Take $V = 0$ for $x < 0$ and $V = V_0$ on the right. For $E > V_0$, write $\phi_L = e^{ikx} + re^{-ikx}$ and $\phi_R = te^{iqx}$, with $k = \sqrt{2mE}/\hbar$ and $q = \sqrt{2m(E - V_0)}/\hbar$. Matching gives $1 + r = t$ and $k(1 - r) = qt$. Thus $r = (k - q)/(k + q)$ and $t = 2k/(k + q)$. A wave can reflect from a change of wavelength even if a classical particle has enough energy to proceed.

A right-moving component Ae^{ikx} carries current $\hbar k|A|^2/m$. Consequently $R = |r|^2$ while $T = (q/k)|t|^2$, and direct substitution gives $R + T = 1$. For $E = 4V_0/3$, $q = k/2$, so $r = 1/3$, $t = 4/3$, $R = 1/9$ and $T = 8/9$. Although the transmitted wavefunction amplitude is larger than the incident one, its reduced group velocity keeps the transmitted probability flux below the incident flux.

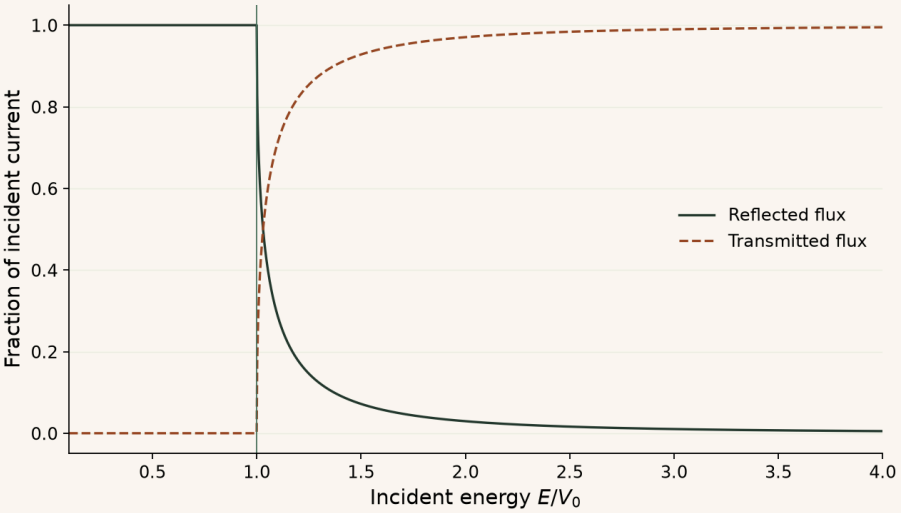

For a semi-infinite potential step the flux fractions sum to unity at each energy. Partial reflection persists above the step; below threshold the decaying wave carries no net transmitted current.

Try it · Step scattering

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Selected energy / step height 2

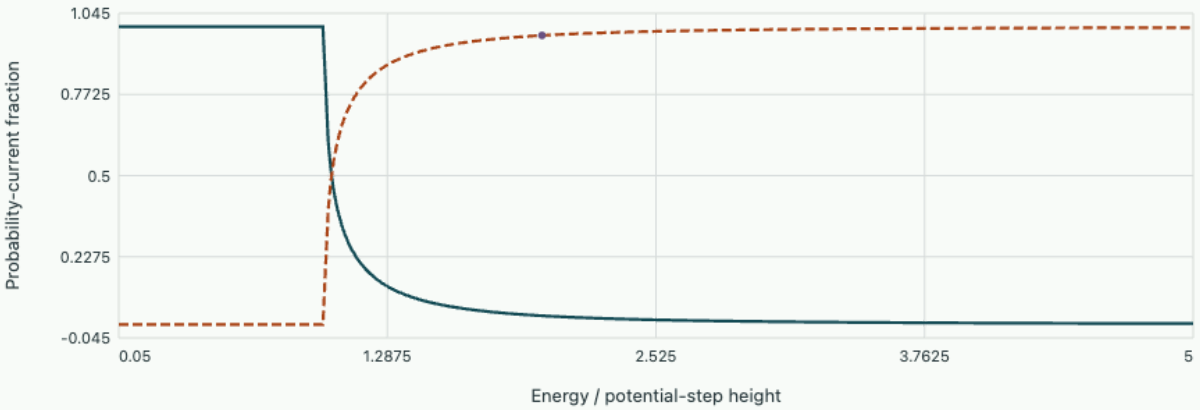

R + T = 1. At the selected energy T = 0.97056. The step extends to infinity; evanescent density below threshold carries no transmitted current.

Background · Why can a transmitted amplitude exceed one?

For incidence from the left, match the wavefunction and its derivative at a potential step with equal particle masses and positive real wavenumbers k_L, k_R. With unit incident amplitude the transmitted amplitude is $t = 2k_L/(k_L + k_R)$. A detector downstream counts probability crossing its position per unit time. The current in a right-travelling plane wave includes its group velocity.

$$T = \frac{k_R}{k_L}|t|^2,$$

$$R = \left|\frac{k_L - k_R}{k_L + k_R}\right|^2.$$

For $k_R = k_L/3$, the transmitted amplitude is 1.5, its intensity factor is 2.25, and its transmitted current fraction is 0.75. The reflected fraction supplies the other 0.25. The density increases because the transmitted particles propagate more slowly. Using flux-normalised channel states absorbs the velocity factor into the scattering amplitude.

Follow an evanescent field

If $0 < E < V_0$, the right-hand wave number becomes $i\kappa$ with $\kappa = \sqrt{2m(V_0 - E)}/\hbar$. The physical solution for a semi-infinite step decays as $e^{-\kappa x}$. This single exponential times a constant has zero current, since its spatial derivative has the same complex phase. It nevertheless has finite density a short distance into the forbidden region. Matching at the interface now gives a reflected wave of unit magnitude with an energy-dependent phase.

A finite barrier differs because an allowed region begins again at $x = L$. Both decaying and growing exponentials must be retained within the barrier, and their relative complex coefficient allows nonzero current. Match the wavefunction and derivative at both ends. For equal zero potentials on the outside the transmission probability reduces to

$$T = \left[1 + \frac{V_0^2 \sinh^2(\kappa L)}{4E(V_0 - E)}\right]^{-1}.$$

The large-width dependence comes from $\sinh(\kappa L) \sim e^{\kappa L}/2$. At fixed barrier height it is the exponential sensitivity to width that makes tunnelling useful as a distance probe. A single semi-infinite exponential fitted to density cannot by itself supply the finite-barrier prefactor.

Estimate the effect of a subnanometre change

For an electron with $E = 1.0$ eV and $V_0 = 2.0$ eV, $\kappa \simeq 5.12$ nm⁻¹ and the prefactor of $\sinh^2$ is exactly one. At width 0.50 nm,

$$T = 1/\cosh^2(2.56) \simeq 0.0236.$$

At 1.00 nm, $T \simeq 1.43 \times 10^{-4}$. Doubling this subnanometre gap suppresses transmitted current by roughly a factor of 165 for equal incident current.

The calculation assumes a sharp rectangular barrier, one occupied incident energy and elastic coherent scattering. A measured electric tunnel current generally combines many available occupied and empty states across an energy window. A current ratio is therefore most directly compared with the single-energy exponential when barrier shape and occupation stay sufficiently similar across the distance change.

A more detailed test uses scattering from opposite sides of a real, static barrier at the same energy. Time-reversal symmetry implies equal transmission probabilities for a single channel, even when the spatial barrier profile is asymmetric. The two reflection phases can differ. If the exterior potentials differ, the transmitted amplitude ratios from the two sides need not coincide because their travelling states carry different current per unit amplitude. Normalizing the prepared and detected states to unit flux removes that ambiguity. Checking incident current against the sum of outgoing currents separately tests whether the model conserves probability.

Solve for scattering and binding at one narrow defect

A point interaction illustrates how the same Hamiltonian predicts both scattering and bound states. On an infinite line take $V(x) = g\delta(x)$, with g having units of energy times length. This limiting model replaces a narrow feature when the wavelength of interest greatly exceeds its microscopic width. The wavefunction remains continuous, but integrating the stationary Schrödinger equation from $-\eta$ to $+\eta$ before taking $\eta \to 0$ gives the derivative jump

$$\psi'(0^+) - \psi'(0^-) = \frac{2mg}{\hbar^2}\psi(0).$$

For a particle incident from the left, write $\psi = e^{ikx} + re^{-ikx}$ on the left and te^{ikx} on the right. Continuity gives $t = 1 + r$. Inserting both derivatives into the jump condition and defining $\xi = mg/(\hbar^2 k)$ gives

$$t = 1/(1 + i\xi),$$
$$r = -i\xi/(1 + i\xi).$$

The external speeds are equal on both sides, so $T = 1/(1 + \xi^2)$ and $R = \xi^2/(1 + \xi^2)$. The sum is one for real g. Transmission probability depends only on the magnitude of the interaction, but the complex phase of the transmitted wave distinguishes its sign.

For negative energies the exterior solution must decay as $e^{-\kappa|x|}$. Its derivatives jump by $-2\kappa\psi(0)$, giving $\kappa = -mg/\hbar^2$. A normalizable state requires $\kappa > 0$, hence only an attractive defect binds. The bound energy is $-\hbar^2\kappa^2/(2m)$. A one-dimensional narrow attractive defect can therefore support binding even when a scattering-probability measurement alone cannot distinguish its sign.

Compare two defects with the same transmission

Use electron mass and take $|g| = 0.0762$ eV nm. Since $\hbar^2/(2m_e) = 0.0381$ eV nm², $m_e|g|/\hbar^2 = 1.00$ nm⁻¹. An incoming wave with $k = 2.00$ nm⁻¹ has energy 0.152 eV and $|\xi| = 0.500$. Both signs of g transmit 0.800 and reflect 0.200 of the incident probability current. The transmission phases, however, are $\mp\arctan(0.500) = \mp 26.6^\circ$, a difference observable against a reference path.

For the attractive sign the decay length of the bound state is $1/\kappa = 1.00$ nm and its energy is -0.0381 eV. The repulsive sign has no such state. At high incident k both become nearly transparent; at small k both produce nearly complete reflection for fixed nonzero strength. The order of these limits matters: taking g to zero first gives perfect transmission at any finite incident wavenumber.

An attractive defect with the inferred decay constant has normalized bound state $\psi = \sqrt{\kappa}\, e^{-\kappa|x|}$. The probability outside distance b from the defect is $e^{-2\kappa b}$ after integrating both tails. With $b = 1.00$ nm in the example this gives 0.135, despite each wave amplitude at that distance being $1/e$ of its central value. A numerical calculation in a box can use this estimate to choose outer boundaries. Placing artificial zero-value boundaries only one decay length away would remove substantial tail probability and shift the apparent bound energy upward.

Increasing the box symmetrically and tracking the bound energy is more informative than merely refining the spatial mesh in the original box. Mesh convergence at fixed outer boundaries can certify a well-resolved solution of the wrong finite-domain physical problem.

Study the full derivation

Inferring a short-range interaction from elastic scattering · Quantum states, dynamics and measurement

A beam of neutral particles encounters a weak, localised static potential. Derive differential cross section from current, extract phase shifts by radial matching and compare exact low-energy scattering with the Born approximation and its failure near a bound state.

Problem sheet 1 · Flux, packets and barriers

Use the probability balance of Lecture 1 and the interface matching of Lecture 2 to interpret measured particle counts.

Give two packets the same density

Two packets have wavefunctions $f(x)e^{ikx}$ and $f(x)e^{-ikx}$, where f is real and normalised. Compare their densities and currents. For electrons with $k = 2.0 \times 10^{8}$ m⁻¹ and local density 4.0×10^{5} m⁻¹, find both signed currents. Does either current by itself give the chance of finding a particle inside a chosen interval?

▼ **Differentiate the phase of each packet**

Both densities are f^2, while their phase gradients are $+k$ and $-k$. The speeds entering the flux are $\pm\hbar k/m_e = \pm 2.32 \times 10^{4}$ m/s, giving currents $\pm 9.27 \times 10^{9}$ s⁻¹ at the stated point. The probability within an interval instead comes from integrating the density. Its rate of change requires the difference of the currents at both ends. A large numerical rate per second can describe a normalised packet passing a point in a very short time.

Determine an unknown step

A beam with kinetic energy 9.0 eV approaches an upward step. Its reflected probability fraction is 0.040. Infer the ratio of wave numbers, the step energy and the transmitted amplitude for unit incident amplitude. Then recover the transmitted probability from current.

▼ **Invert the matching conditions**

For an upward step below the incident energy r is positive, so $r = 0.20$. The ratio $q/k = (1-r)/(1+r) = 2/3$ gives $(E - V_0)/E = 4/9$ and $V_0 = 5.0$ eV. Displacement continuity gives $t = 1.20$. The current fraction is $(2/3)(1.20)^2 = 0.96$, closing the balance with $R = 0.04$. The step energy in this inverse calculation depends on having specified the sign of the step.

Infer a tunnelling decay length

At fixed incident energy a wide rectangular barrier transmits 2.0×10^{-4} of incident particles. Increasing its width by 0.20 nm reduces that fraction to 2.0×10^{-5}. Estimate κ, the field-amplitude decay length and $V_0 - E$ for an electron. State which part of the inference would fail for a thin barrier.

▼ **Use the ratio to eliminate the prefactor**

For $\kappa L \gg 1$ the ratio is $e^{-2\kappa \Delta L} = 0.10$, so $\kappa = \ln 10/(0.40\,\mathrm{nm}) = 5.76$ nm⁻¹ and $1/\kappa = 0.174$ nm. With $\hbar^2/(2m_e) = 0.0381$ eV·nm², $V_0 - E = 0.0381(5.76)^2 = 1.26$ eV. In a thin barrier the complete hyperbolic-sine matching expression is needed; taking a log ratio then need not yield a width-independent slope.

Infer a force from the motion of a centroid

A ground-state packet in a harmonic trap is displaced by $0.200\,\mu\mathrm{m}$ and released from rest. The trap frequency is $500\,\mathrm{Hz}$. Use the mean-value equation to predict its initial acceleration and centroid after $0.500\,\mathrm{ms}$. What changes if the initial state is a broader mixture with the same means?

▼ **Calculate the acceleration, phase and width dependence**

With $V = m\omega^2 x^2/2$, linearity gives

$$\langle \ddot{x} \rangle = -\omega^2 \langle x \rangle .$$

Here $\omega = 2\pi(500\,\mathrm{s}^{-1})$, so

$$\langle \ddot{x} \rangle(0) = -(3141.6)^2 (2.00 \times 10^{-7}) = -1.97\,\mathrm{m/s^2}.$$

The zero initial velocity selects $\langle x \rangle(t) = 0.200\,\mu\mathrm{m}\cos(\omega t)$. After $0.500\,\mathrm{ms}$ the phase is $\pi/2$ and the centroid crosses the trap centre. A mixture with the same initial position and momentum means obeys the identical linear equation. Its single-shot outcomes may have a different spread.

Revisit chapter 1 · When does a wave packet follow Newton’s equation?.

Explore these relations

- For the unknown step, vary the energy-to-height ratio and compare reflected and transmitted probability currents. Open the live calculation.
- For the tunnelling decay length, compare finite-barrier transmission with energy and barrier width. Open the live calculation.

Chapter 3 · Operators and the quantum harmonic oscillator

Use commutators and ladder operators to derive an energy spectrum and connect its ground state with the position–momentum uncertainty bound.

FROM FIRST PRINCIPLES

Build the oscillator ladder from the canonical commutator

Near a stable minimum use the one-dimensional harmonic Hamiltonian $H = p^2/(2m) + m\omega^2 x^2/2$ and the canonical commutator $[x, p] = i\hbar$. Combine displacement and momentum into dimensionless operators so the two quadratic energy terms acquire the same coefficient:

$$a = \sqrt{\frac{m\omega}{2\hbar}}\, x + \frac{ip}{\sqrt{2m\hbar\omega}},$$
$$[a, a^\dagger] = 1.$$

Multiplying $a^\dagger a$ produces the two energy terms and a cross term proportional to the commutator. Retaining that cross term fixes the offset. Commuting the Hamiltonian with a then shows how it acts on an energy eigenstate:

$$H = \hbar\omega(a^\dagger a + \tfrac{1}{2}),$$
$$H(a|E\rangle) = (E - \hbar\omega)a|E\rangle.$$

The squared norm $\|a|E\rangle\|^2 = E/(\hbar\omega) - 1/2$ cannot be negative. Repeated lowering must end at a state annihilated by a, with energy $\hbar\omega/2$. Raising creates successive normalized states and gives

$$E_n = (n + 1/2)\hbar\omega.$$

The spacing follows from the algebra; the added half quantum reflects the impossibility of independently setting both noncommuting fluctuations to zero.

Background · Commutator

$$[A, B] = AB - BA$$

Order matters. The expectation of the commutator controls a lower bound on the product of uncertainties when the state lies in the relevant operator domains.

DERIVATION AND APPLICATIONS

Build position and momentum fluctuations from adjacent levels

A number-state oscillator has a definite energy and an uncertain position. How large are its fluctuations, and what survives when the mean position vanishes? Assume a one-dimensional harmonic trap with mass m and angular frequency ω. The ladder algebra gives

$$a|n\rangle = \sqrt{n}\,|n-1\rangle,$$
$$a^\dagger|n\rangle = \sqrt{n+1}\,|n+1\rangle.$$

Define the ground-state length $x_0 = \sqrt{\hbar/(2m\omega)}$ and momentum scale $p_0 = \sqrt{m\hbar\omega/2}$.

$$x = x_0(a + a^\dagger),$$
$$p = ip_0(a^\dagger - a).$$

Both operators take $|n\rangle$ into orthogonal neighbouring states. Their diagonal means vanish. To get a variance, compute the squared norm of the state after application. This avoids dropping the commutator when expanding the square:

$$\|x|n\rangle\|^2 = x_0^2\big(\|\sqrt{n}\,|n-1\rangle\|^2 + \|\sqrt{n+1}\,|n+1\rangle\|^2\big) = (2n+1)x_0^2.$$

The relative minus sign in $p|n\rangle$ disappears on taking the norm because the neighbours are orthogonal. It follows that

$$\Delta x = \sqrt{2n+1}\,x_0,$$
$$\Delta p = \sqrt{2n+1}\,p_0,$$
$$\Delta x\,\Delta p = (n+\tfrac{1}{2})\hbar.$$

The ground state attains the minimum uncertainty. Excited number states exceed that lower bound despite having no energy uncertainty. The same calculation gives equal kinetic and potential energy means, each $(n+1/2)\hbar\omega/2$. Each adjacent-level matrix element also specifies which spectral lines a weak displacement-dependent perturbation can drive.

WORKED EXAMPLE

Estimate a trapped ion's spatial spread

Change the inputs and investigate this derivation.

For $m = 6.64 \times 10^{-26}$ kg and $\omega = 2\pi(1.00\,\mathrm{MHz})$, the natural width is $x_0 = 11.2$ nm. The $n = 2$ preparation has $\Delta x = \sqrt{5}\,x_0 = 25.1$ nm and $\Delta x\Delta p = 2.50\hbar$. Dividing the latter by the position spread gives $\Delta p = 1.05 \times 10^{-26}$ kg m/s. The width would remain the same at later times in this stationary energy state.

SECOND APPLICATION

Distinguish two preparations with the same energy mean

Change the inputs and investigate this derivation.

Compare the pure state $|1\rangle$ with an equal incoherent mixture of $|0\rangle$ and $|2\rangle$. Both have mean energy $1.50\hbar\omega$ and $\langle x^2\rangle = 3x_0^2$, because these quantities are linear in n. In the mixture, the two energy outcomes differ from the mean by $\pm\hbar\omega$, so $\Delta E = \hbar\omega$; the pure number state has $\Delta E = 0$. An energy-resolved measurement distinguishes the preparations even when a width measurement cannot.

Try it · Oscillator uncertainty

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Occupation number n **2**

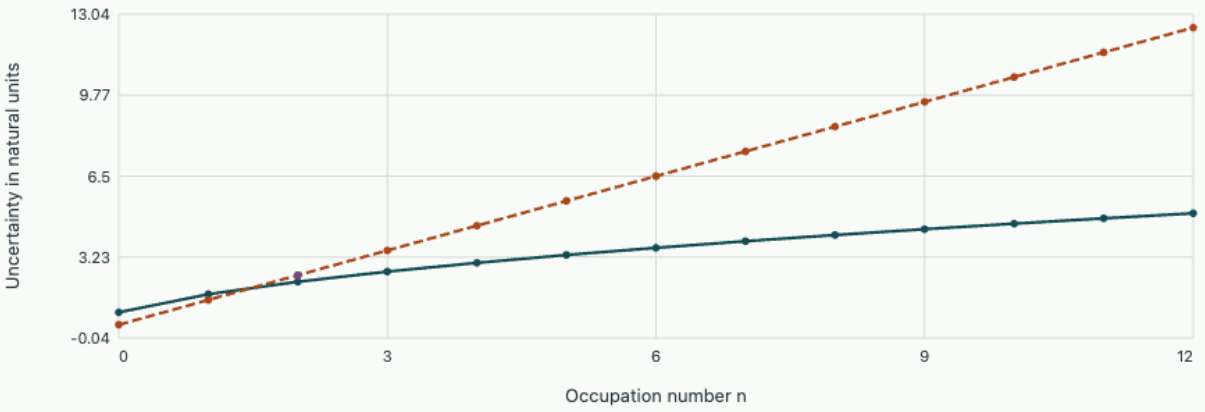

x0=√[ħ/(2mω)] and p0=√[mħω/2]. The n=0 state saturates the lower bound; number-state uncertainty grows with n.

Quantise near a stable minimum

A molecular bond stretched slightly from equilibrium has a potential $V(x) \simeq V(0) + kx^2/2$ if the equilibrium derivative vanishes and the curvature k is positive. For relative motion the relevant mass is the reduced mass of the two nuclei. Discard the additive constant and write $k = m\omega^2$. Quantum mechanics assigns the Hamiltonian $\hat{H} = \hat{p}^2/(2m) + m\omega^2\hat{x}^2/2$. Its two terms do not commute because position and momentum obey $[\hat{x}, \hat{p}] = i\hbar$.

For a normalised state, the expectation of an operator is $\langle A\rangle = \langle\psi|\hat{A}|\psi\rangle$. Applying the Cauchy–Schwarz inequality to $(\hat{x}-\langle x\rangle)|\psi\rangle$ and $(\hat{p}-\langle p\rangle)|\psi\rangle$ shows that the product of their norms is at least the magnitude of their inner product. The imaginary part of that inner product is controlled by the commutator. One obtains

$$\Delta x\,\Delta p \geq \hbar/2,$$

where the spreads are standard deviations of repeated measurements on equally prepared systems.

Factor a positive Hamiltonian

Define dimensionless operators $a = \sqrt{m\omega/(2\hbar)}\,\hat{x} + i\hat{p}/\sqrt{2m\hbar\omega}$ and $a^\dagger$ as its adjoint. Multiplying them in the indicated order and using the commutator yields $[a, a^\dagger] = 1$ and

$$\hat{H} = \hbar\omega(a^\dagger a + \tfrac{1}{2}).$$

For any state the expectation of $a^\dagger a$ is $\|a\psi\|^2$, which cannot be negative. The Hamiltonian therefore has a lower bound $\hbar\omega/2$. Interchanging operator order while treating x and p as ordinary numbers would miss this term. It accounts for the fact that a state cannot simultaneously place the particle at the bottom of the well and give it precisely zero momentum.

Now compute the commutator of $\hat{H}$ with a. The identity $[AB, C] = A[B, C] + [A, C]B$ gives

$$[\hat{H}, a] = -\hbar\omega a.$$

Operating with a on an energy eigenstate either gives zero or produces a new eigenstate lower by $\hbar\omega$. Successive lowering must terminate since the energy cannot become negative. The final state obeys $a|0\rangle = 0$. Raising from it gives $E_n = (n + 1/2)\hbar\omega$ and the normalisation factors $a|n\rangle = \sqrt{n}\,|n-1\rangle$ and $a^\dagger|n\rangle = \sqrt{n+1}\,|n+1\rangle$.

Find the shape behind the lowest level

Use $\hat{p} = -i\hbar\,d/dx$ in the ground-state condition. It reduces to $d\psi_0/dx = -(m\omega/\hbar)x\psi_0$, whose square-integrable solution is proportional to $\exp[-m\omega x^2/(2\hbar)]$. Normalising the squared Gaussian gives position variance $\hbar/(2m\omega)$. Its momentum variance is $m\hbar\omega/2$. The product of standard deviations is exactly $\hbar/2$, so the ground state saturates the general uncertainty bound.

The mean potential and kinetic energies in this state are both $\hbar\omega/4$. Increasing the stiffness narrows the position distribution and raises the momentum spread. This provides a physical interpretation of the increasing ground energy which is independent of the ladder notation. In an excited energy eigenstate the mean position still vanishes by parity, yet its variance is larger.

Read a vibrational line and a selection rule

Take effective mass 1.0×10^{-26} kg and harmonic frequency $f = 5.0$ THz. Then $\omega = 2\pi f$, adjacent levels differ by $hf = 20.7$ meV, and the zero-point energy is 10.3 meV. The ground-state position standard

deviation is 1.30×10^{-11} m. A spectral measurement of adjacent level spacings determines the curvature-to-mass ratio but cannot separate mass and stiffness without another measurement.

A weak oscillatory perturbation proportional to displacement has matrix elements of $\hat{x} \propto a + a^\dagger$. It therefore connects adjacent number states, giving the dipole selection rule $\Delta n = \pm 1$ within the harmonic model and a linear coupling. If a real molecular spectrum contains overtones or nonuniform spacings, anharmonicity or a nonlinear dependence of dipole moment on displacement can supply the missing physics.

It is possible to recover the zero-point energy without solving a differential equation. A centred packet of spatial variance σ_x^2 has potential energy $m\omega^2\sigma_x^2/2$. The uncertainty relation bounds its kinetic energy below by $\hbar^2/(8m\sigma_x^2)$ when mean momentum vanishes. Minimizing their sum over width gives $\sigma_x^2 = \hbar/(2m\omega)$ and an energy bound $\hbar\omega/2$. A Gaussian saturates the uncertainty inequality and reaches the bound. Making a classical oscillator sit identically at rest would require both energy contributions to vanish simultaneously, which the commutator forbids.

Displace an oscillator ground state

A harmonic oscillator can move almost like a classical oscillator while retaining its smallest quantum position uncertainty. Prepare the ground state and suddenly shift the centre of its trapping potential by distance d, without changing frequency. In the coordinates of the new well the old Gaussian starts displaced by $-d$ with mean momentum zero. Define a coherent state as an eigenstate of the lowering operator,

$$\hat{a}|\alpha\rangle = \alpha|\alpha\rangle.$$

Its displacement and momentum are encoded by $\alpha = \sqrt{m\omega/(2\hbar)}\langle x\rangle + i\langle p\rangle/\sqrt{2m\hbar\omega}$.

To find its energy distribution expand $|\alpha\rangle = \sum_n c_n|n\rangle$. Applying the lowering operator and equating the coefficient of $|n\rangle$ gives

$$\sqrt{n+1}\,c_{n+1} = \alpha c_n.$$

Recursion and normalization give $c_n = e^{-|\alpha|^2/2}\alpha^n/\sqrt{n!}$ up to a common phase. The energy index follows a Poisson distribution,

$$P_n = e^{-\bar{n}}\frac{\bar{n}^n}{n!},$$

$$\bar{n} = |\alpha|^2,$$

$$\langle E\rangle = \hbar\omega(\bar{n} + \tfrac{1}{2}).$$

Free evolution multiplies each coefficient by $e^{-in\omega t}$ apart from a common factor, turning α into $\alpha e^{-i\omega t}$. Position and momentum expectations execute the corresponding classical orbit, and the packet stays a fixed-width Gaussian because adjacent energy levels have exactly uniform spacing. A potential with appreciably unequal spacings instead allows these components to lose their phase alignment over time.

Estimate the number of excited quanta

Let an effective mass $m = 1.00 \times 10^{-26}$ kg oscillate at $\omega = 1.00 \times 10^{13}$ rad/s. Its ground-state position standard deviation is

$$x_{\mathrm{zpf}} = \sqrt{\hbar/(2m\omega)} = 2.30 \times 10^{-11}\,\mathrm{m}.$$

Shift the well by $d = 2x_{\mathrm{zpf}} = 46.0$ pm. In the new coordinate $\alpha = -d/(2x_{\mathrm{zpf}}) = -1$, giving $\bar{n} = 1$. Hence ground and first excited energies each occur with probability $e^{-1} = 0.368$, and the second excited energy occurs with probability $e^{-1}/2 = 0.184$.

The energy added above the ground state is $\hbar\omega = 6.58$ meV, identical to the classical change $m\omega^2 d^2/2$ immediately following the shift. Equality of mean energies does not make a single measurement deterministic: the energy standard deviation is also $\hbar\omega$ in this preparation. Moving the minimum quasistatically instead of suddenly can leave the system close to its instantaneous ground state and add essentially no excitation, provided the motion of the trap is slow on the oscillation timescale.

Number statistics provide a stringent check on the sudden-shift model. A coherent state with mean occupation one has probability $1 - 2.5/e = 0.0803$ of occupying level three or higher. Failure to observe those outcomes in a small sample would be unsurprising, but persistent absence in a well-resolved large sample would contradict the ideal state. Two displaced packets with opposite signs of α have the same number distribution and different later position records. Thus energy spectroscopy checks the magnitude of a displacement while a timed position measurement also checks its phase around the oscillator orbit.

A time sequence with the same mean added energy can still have different number statistics if the trap frequency changes along with its centre. A frequency jump changes the natural Gaussian width and can prepare a squeezed state. Observing a breathing position variance in addition to centre motion would reveal this extra change in the restoring curvature.

Recover oscillator nodes and parity from a position distribution

Apply the creation operator to a normalised Gaussian. The first two excited wavefunctions acquire the polynomial factors $\sqrt{2}x$ and $(2x^2 - 1)/\sqrt{2}$, where x is in oscillator lengths. Compute densities only after the amplitude has been formed.

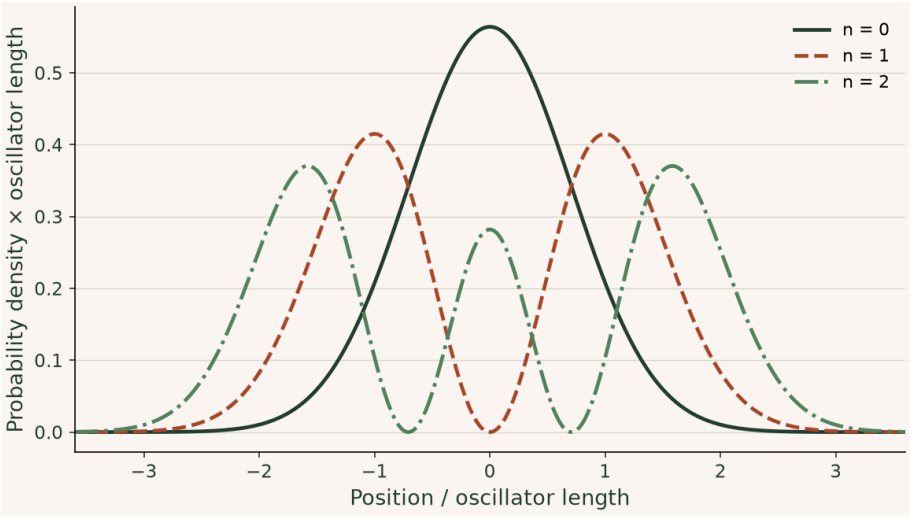

The odd first-excited state has a node at the origin, and the second-excited state has nodes at $\pm 1/\sqrt{2}$. Density reveals the nodes but loses the sign between them, which matters when building a coherent superposition.

Try it · Oscillator density

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Oscillator length (nm)	1
Excited-state admixture probability	0
Relative phase (degrees)	0

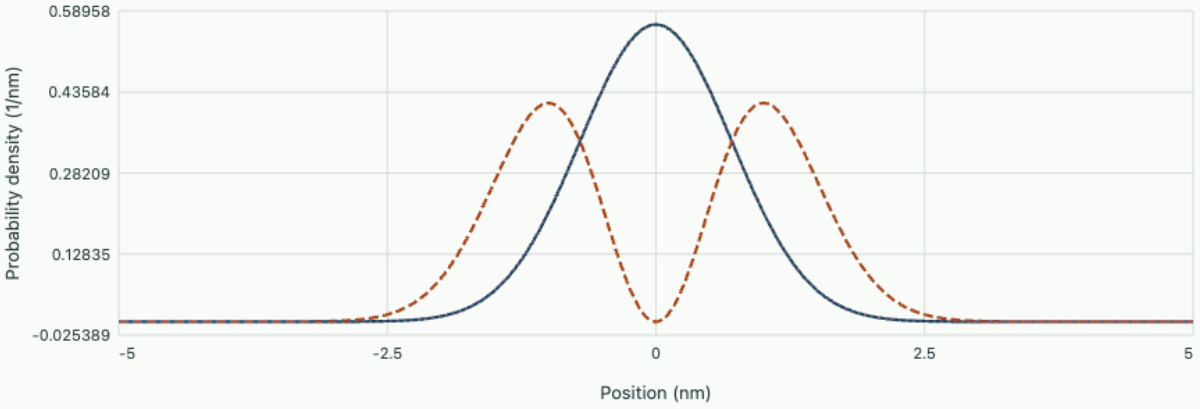

Ground and first-excited basis states are orthogonal. Changing the relative phase alters interference while preserving the integrated probability.

▶ **Read values from the plot**

Study the full derivation

Why an oscillator can move without its packet spreading · Quantum states, dynamics and measurement

Start with the measured mass and curvature of a potential. Factor its Hamiltonian, prove the energy ladder, construct a freely moving Gaussian from number states and find when squeezing or anharmonicity changes that prediction.

Chapter 4 · Angular momentum and hydrogenic atoms

Use rotational symmetry to organise atomic states and derive the hydrogen ground-state length and energy scales from the radial equation.

FROM FIRST PRINCIPLES

Derive the radial equation from three-dimensional motion

For a reduced mass μ in a central potential $V(r)$, rotational symmetry makes the angular-momentum operators compatible with the Hamiltonian. Write a stationary state as $\psi = R(r)Y_{\ell m}(\theta,\phi)$ with $L^2 Y_{\ell m} = \hbar^2 \ell(\ell+1) Y_{\ell m}$. In spherical coordinates the kinetic operator separates radial and angular changes:

$$-\frac{\hbar^2}{2\mu}\nabla^2 = -\frac{\hbar^2}{2\mu r^2}\partial_r(r^2\partial_r) + \frac{L^2}{2\mu r^2}.$$

Divide the eigenvalue equation by the angular factor. To work with a one-dimensional probability measure, substitute $u = rR$. Then $\int |R|^2 r^2 dr = \int |u|^2 dr$ and direct differentiation gives

$$(r^2 R')' = r u''.$$

The resulting equation is

$$-\frac{\hbar^2}{2\mu}u'' + \left[V(r) + \frac{\hbar^2 \ell(\ell+1)}{2\mu r^2}\right]u = Eu.$$

The extra positive term at nonzero ℓ is the angular kinetic energy, which inhibits penetration near the origin. For ordinary Coulomb bound states regularity imposes $u(0) = 0$ and normalizability requires decay at infinity. The radial half-line and its boundary conditions are part of the eigenvalue problem.

DERIVATION AND APPLICATIONS

What fraction of a ground-state atom lies outside a chosen radius?

The hydrogenic ground-state amplitude is largest at the nucleus, yet a measurement of radius is most likely to land away from it. To derive the distribution use a single electron bound to a point nucleus of charge $+Ze$, neglecting relativistic and finite-size corrections. Let μ be the reduced electron–nucleus mass, and write $\kappa = Ze^2/(4\pi\epsilon_0)$. For a spherically symmetric stationary state try $\psi = Ae^{-r/a}$. Inserting its radial derivatives into the three-dimensional Schrödinger equation gives

$$\frac{H\psi}{\psi} = -\frac{\hbar^2}{2\mu a^2} + \frac{1}{r}\left(\frac{\hbar^2}{\mu a} - \kappa\right).$$

The right side must be independent of radius. Cancel the inverse-radius term to fix $a = \hbar^2/(\mu\kappa)$. Normalization integrates probability over spherical shells, each of volume $4\pi r^2 dr$:

$$1 = 4\pi|A|^2 \int_0^\infty r^2 e^{-2r/a}\,dr = \pi a^3 |A|^2,$$

$$P_r(r) = \frac{4r^2}{a^3} e^{-2r/a}.$$

Thus $P_r(r)\,dr$ is the probability for a radius in an interval. Differentiating it places its maximum at $r = a$. To count the atoms found outside radius R, substitute $z = 2r/a$ and integrate the quadratic times an exponential twice by parts:

$$\Pr(r > R) = \tfrac{1}{2} \int_{2R/a}^\infty z^2 e^{-z}\,dz = e^{-2R/a}\left(1 + \frac{2R}{a} + \frac{2R^2}{a^2}\right).$$

At $R = 0$ the result is one, as required. The tail remains nonzero for every finite radius; a quote of the “atomic radius” must therefore specify the observable or convention used.

WORKED EXAMPLE

Quantify the extent of ordinary hydrogen

Change the inputs and investigate this derivation.

Take $a = 52.9$ pm. A radius $R = 105.8$ pm equals $2a$. Evaluating the tail gives

$$13e^{-4} = 0.238.$$

Nearly one quarter of radial outcomes fall beyond twice the most probable radius. Using only the relative density $e^{-4} = 0.0183$ at that point would omit both shell volume and integration over larger radii.

SECOND APPLICATION

Compare two ions through the same physical sphere

Change the inputs and investigate this derivation.

For a one-electron helium ion and a heavy nucleus approximation, doubling Z halves a to 26.45 pm. At a physical radius of 52.9 pm, the hydrogen tail is $5e^{-2} = 0.677$, while the ion tail is

$$13e^{-4} = 0.238.$$

The same fractional containment would be recovered by comparing radii in proportion to $1/Z$. This comparison assumes an isolated ion with no second electron to screen the nuclear field.

Try it · Ph coulomb radii

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Nuclear charge Z

1

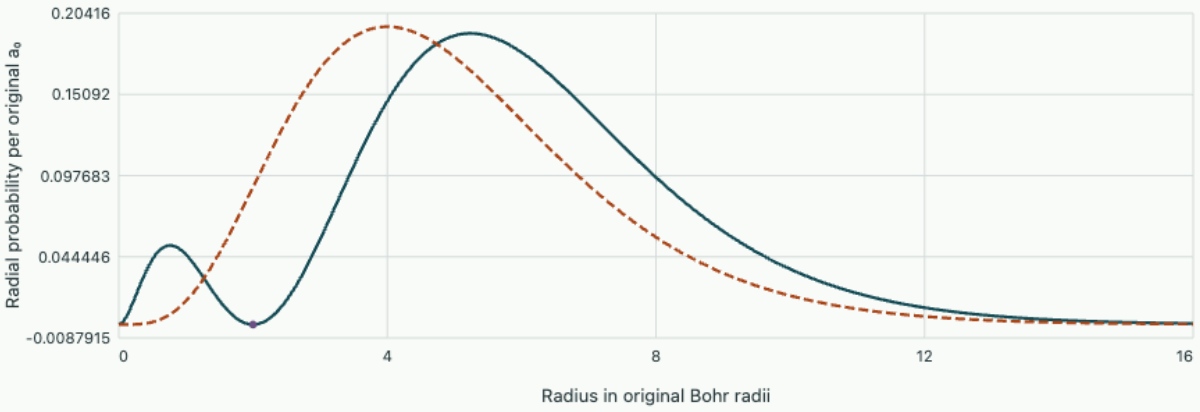

The 2s node is at 2 original Bohr radii. The most probable 2p radius is 4 original Bohr radii.

Identify the symmetry before choosing coordinates

An isolated electron and nucleus interacting through a central potential can be separated into centre-of-mass motion and relative motion. Let μ be the reduced mass of the two particles. The relative Hamiltonian is $\hat{p}^2/(2\mu) + V(r)$ with r their separation. Simultaneously rotating position and momentum does not change the energy. Angular momentum provides the observables that express this symmetry and organise the many possible three-dimensional wavefunctions.

Orbital angular momentum is $\hat{\mathbf{L}} = \hat{\mathbf{r}} \times \hat{\mathbf{p}}$. Its components obey $[L_x, L_y] = i\hbar L_z$ and cyclic permutations. The total magnitude squared $\hat{L}^2$ commutes with each component. We can therefore specify eigenvalues of energy, L^2 and one projection, conventionally L_z, together. The respective angular eigenvalues are $\ell(\ell+1)\hbar^2$ and $m\hbar$, with integer $\ell \geq 0$ and $m = -\ell, \ldots, \ell$. Spherical harmonics $Y_{\ell m}$ represent those angular states, with probability density proportional to their squared magnitude.

Separate angular and radial motion

Write $\psi(r,\theta,\varphi)=R(r)Y_{\ell m}(\theta,\varphi)$. Volume in spherical coordinates contains a factor r^2, so normalising the angular function leaves radial probability $r^2|R(r)|^2dr$. Defining $u(r)=rR(r)$ both simplifies the radial normalisation and removes the first derivative from the radial equation:

$$-\frac{\hbar^2}{2\mu}u''+\left[V(r)+\frac{\hbar^2\ell(\ell+1)}{2\mu r^2}\right]u=Eu.$$

For ordinary bound atomic states regularity at the origin gives $u(0)=0$ and normalisability gives decay at large radius. The angular momentum term penalises small radius when $\ell>0$. Although it resembles the effective centrifugal energy of classical central motion, here its coefficient comes from the angular eigenvalue rather than an adjustable classical angular momentum.

Derive the lowest Coulomb state

For hydrogen set $V=-e^2/(4\pi\epsilon_0 r)=-C/r$ with $C>0$. For $\ell=0$ try $R=Ae^{-r/a}$, hence $u=Are^{-r/a}$. The second derivative satisfies

$$u''/u=1/a^2-2/(ar).$$

Substitute this ratio in the radial equation. The terms proportional to $1/r$ cancel only when $a=\hbar^2/(\mu C)$, leaving $E_1=-\hbar^2/(2\mu a^2)$. This nodeless solution is the lowest bound state. Choosing another exponential length would leave a residual $1/r$ term, giving a useful check on the derivation.

For hydrogen with a heavy proton approximation, $a\simeq 0.0529$ nm and $E_1\simeq -13.6$ eV. The three-dimensional density of the ground state is largest at the origin. The probability per unit radial distance instead contains r^2 and peaks at $r=a$. A claim about the “most likely radius” must specify which measure is in use. Measuring a shell of equal thickness at a larger radius includes more volume.

Distinguish spectra from angular distributions

Solving the regular bound-state radial equation for general ℓ gives $E_n=-\mu C^2/(2\hbar^2 n^2)$ with $n=1,2,\dots$ and $0\le\ell\le n-1$. States with different m are degenerate by rotational symmetry. The additional degeneracy among different ℓ at fixed n is a special property of the Coulomb potential. Corrections such as fine structure and external fields can separate lines that the present model puts at one energy.

The transition from $n=3$ to $n=2$ releases $13.6(1/2^2-1/3^2)=1.89$ eV, giving wavelength $hc/\Delta E\simeq 656$ nm. The energy gap predicts the wavelength but not the strength of every possible transition. Electric dipole emission involves an angular matrix element of position, so it connects orbital angular momenta differing by one. An excited state can have the right energy gap while a particular dipole transition remains forbidden by symmetry. This is why both radial energies and angular labels belong in a spectroscopy record.

The reduced mass supplies a further spectroscopic test. Replacing a proton by a deuteron approximately doubles the nuclear mass while preserving its charge. Since $\mu=m_e/[1+m_e/M]$, the electronic reduced mass increases by a relative amount close to $m_e/(2m_p)=0.000272$. All gross transition frequencies rise by

the same fraction, while wavelengths fall. A protium line at 656 nm should have a deuterium counterpart roughly 0.179 nm to its short-wavelength side. Resolving the two lines requires sufficient spectral resolution and a source with both isotopes. Small additional shifts reveal corrections beyond the point-nucleus, nonrelativistic model.

Couple orbital and spin angular momenta

The angular labels of a central potential organize an electron's spatial and spin states separately. A weak spin–orbit interaction can split levels without breaking overall rotational symmetry. Model it within a fixed radial and orbital manifold by $H_{\mathrm{so}} = \lambda\,\mathbf{L}\cdot\mathbf{S}/\hbar^2$, with λ an energy representing the radial coupling average. The model assumes coupling small relative to separations from other radial or orbital terms. Let $\mathbf{J} = \mathbf{L} + \mathbf{S}$. The identity $J^2 = L^2 + S^2 + 2\mathbf{L}\cdot\mathbf{S}$ expresses the interaction entirely through simultaneous angular eigenvalues,

$$\Delta E_j = \frac{\lambda}{2}\left[j(j+1) - \ell(\ell+1) - s(s+1)\right].$$

The allowed total values run from $|\ell - s|$ to $\ell + s$ in unit steps. Each j has $2j+1$ possible magnetic projections. The interaction is a rotational scalar, so all projections of the same j remain degenerate here. The original labels m_ℓ and m_s taken separately no longer label eigenstates in general, since terms L_+S_- and L_-S_+ exchange projection while keeping their sum fixed.

For $\ell = 1$ and electron spin $s = 1/2$, the product space has three orbital choices and two spin choices, six states in all. The allowed $j = 3/2$ quartet shifts upward by $\lambda/2$ and the $j = 1/2$ doublet downward by λ for positive coupling. Weighted over all six states the shifts sum to zero, giving a useful check on energy bookkeeping.

Reconstruct a split spectral term

Suppose the two energies in such a p term are separated by 4.50 meV and the fourfold component is the upper one. The model gives splitting $3\lambda/2$, hence $\lambda = 3.00$ meV. The quartet lies 1.50 meV above the unsplit term and the doublet 3.00 meV below it. If the observed quartet energy is 2.00150 eV on a common reference, the inferred unsplit centre is 2.00000 eV and the doublet lies at 1.99700 eV. The corresponding difference in optical frequencies for transitions from a common lower level is $4.50\,\mathrm{meV}/h = 1.09$ THz.

Relative line intensities are generally not fixed solely by the quartet-to-doublet multiplicity, because lower-state populations and dipole matrix elements enter. An external magnetic field can additionally split different m_j values and eventually mix the angular couplings. In the special limit $\ell = 0$, there is only $j = s$ and the shift vanishes identically. Observing a doublet from an apparently orbital-singlet term would therefore point to another interaction or a mistaken assignment.

A measurement of angular momentum projection can distinguish the coupling assignments independently of transition wavelengths. The completely aligned state with $m_\ell = 1$ and $m_s = 1/2$ has total projection $3/2$ and can belong only to the $j = 3/2$ manifold. For smaller total projection a given j state generally contains more than one orbital–spin product. The small-field magnetic response then reflects the combined angular content rather than the separate labels of one arbitrarily chosen component. Preparing or identifying a

stretched state provides a particularly direct check on which member of a fine-structure doublet is the quartet.

Study the full derivation

Adding angular momentum and predicting which lines can appear · Quantum states, dynamics and measurement

A three-component orbital wavefunction combines with a two-component spin. Derive the allowed total angular momenta, construct their amplitudes explicitly and predict both forbidden transitions and relative strengths in a polarisation-resolved spectrum.

Problem sheet 2 · Quantised vibrations and atomic states

Combine the operator methods of Lecture 3 and central-potential structure of Lecture 4 to read spectral and spatial measurements.

Change an oscillator’s isotope

A bond with unchanged stiffness has reduced mass m and adjacent vibrational spacing 40 meV. An isotopic substitution doubles its reduced mass. Find the new spacing and the ratio of the ground-state position widths. Does measuring an overtone alone establish the isotope change?

▼ **Separate the effects of mass and frequency**

Since $\omega = \sqrt{k/m}$ the spacing falls by $1/\sqrt{2}$ to 28.3 meV. The width $\sqrt{\hbar/(2m\omega)}$ scales as $m^{-1/4}$ for constant stiffness, giving a ratio $2^{-1/4} = 0.841$. A single overtone may arise from anharmonicity or a nonlinear dipole coupling. Isotope identification requires enough independent information to distinguish those effects from a change in mass or stiffness.

Different definitions of a likely radius

For hydrogen’s normalised 1s spatial density $e^{-2r/a}/(\pi a^3)$, find the radius of largest shell probability and the probability inside a sphere of radius a. Explain why the probability to find the electron at one specified mathematical point is zero.

▼ **Integrate using the spherical measure**

Multiplying by $4\pi r^2$ gives radial density $4r^2 e^{-2r/a}/a^3$. Its nonzero maximum occurs at $r = a$ by differentiation. With $u = 2r/a$, the integral to a is

$$\tfrac{1}{2}\int_0^2 u^2 e^{-u} du = 1 - 5e^{-2} = 0.323.$$

Finite probability requires integration over a finite region. A maximum of the spatial density at one point is a statement about probabilities per unit volume, whereas a single point has zero volume.

Compare hydrogen and a one-electron ion

Neglect reduced-mass differences. For a helium nucleus with charge $+2e$ and one electron, derive the changes relative to hydrogen in ground-state length, binding energy and the wavelength of the $n = 3$ to $n = 2$ transition. Count the orbital states of the upper level before spin is included.

▼ **Rescale the Coulomb attraction**

Replacing C by $2C$ gives $a \propto 1/C$, so the length halves to about 0.0265 nm. Binding energy scales as C^2 and rises to 54.4 eV. The transition energy is four times larger, so its wavelength falls from about 656 to 164

nm. For $n = 3$, $\ell = 0, 1, 2$ contribute $1 + 3 + 5 = 9$ orbital states. The degeneracy count does not imply that all of their dipole transitions have the same intensity or are allowed.

Explore these relations

- For the isotope question, compare oscillator position and momentum widths with the oscillator length; then apply its mass dependence. Open the live calculation.
- For the atomic-radius questions, compare radial probability, expectation values and angular-momentum states. Open the live calculation.

Chapter 5 · Perturbations, degeneracy and level repulsion

Calculate an energy correction, identify where nondegenerate perturbation theory fails, and resolve an avoided crossing by diagonalising a two-state Hamiltonian.

FROM FIRST PRINCIPLES

Resolve an energy correction into virtual couplings

Take $H = H_0 + \lambda V$ with a known orthonormal spectrum $H_0|n\rangle = E_n^0|n\rangle$. Assume the target level is isolated and the coupling weak relative to the relevant energy gaps. Write its perturbed state in intermediate normalization as $|n\rangle + \lambda|n^{(1)}\rangle + \cdots$ with $\langle n|n^{(1)}\rangle = 0$. Insert this and the energy expansion in the eigenvalue equation. Terms at first order give

$$(H_0 - E_n^0)|n^{(1)}\rangle = (E_n^{(1)} - V)|n\rangle.$$

Projection onto n gives

$$E_n^{(1)} = V_{nn}.$$

Projection onto each $m \neq n$ gives the amplitude admixed from that other state. At second order the projection back onto n therefore gives

$$\langle m|n^{(1)}\rangle = \frac{V_{mn}}{E_n^0 - E_m^0},$$

$$E_n^{(2)} = \langle n|V|n^{(1)}\rangle = \sum_{m\neq n} \frac{|V_{mn}|^2}{E_n^0 - E_m^0}.$$

Higher unperturbed levels contribute negatively to the ground-state energy, while lower levels shift an excited state upward. A vanishing denominator signals that nearby states must be diagonalized together before expanding the effect of more distant states.

DERIVATION AND APPLICATIONS

Read a mixed eigenstate from the slope of its energy

A level diagram around an avoided crossing contains more information than the minimum gap. Its slopes can reveal how much of each underlying basis state is present. Consider two orthonormal states with Hamiltonian $H = (\Delta/2)\sigma_z + g\sigma_x$, real constant coupling g, and a tunable signed energy difference Δ. We call the state with $\sigma_z = +1$ the positive basis state. For a normalized, nondegenerate eigenstate, differentiate $H|n\rangle = E_n|n\rangle$ with respect to Δ. Left multiplication by $\langle n|$ gives

$$\langle n|\partial_\Delta H|n\rangle + \langle n|H|\partial_\Delta n\rangle = \partial_\Delta E_n + E_n\langle n|\partial_\Delta n\rangle.$$

The last terms cancel by Hermiticity. Since $\partial_\Delta H = \sigma_z/2$, the measured slope equals half the population difference. If P_+ is the probability of the positive basis state in this eigenstate, normalization gives

$$\partial_\Delta E_n = \tfrac{1}{2}\langle \sigma_z \rangle = P_+ - \tfrac{1}{2}.$$

The characteristic equation $E^2 = \Delta^2/4 + g^2$ supplies the eigenenergies. For the lower branch, evaluate its derivative explicitly:

$$E_- = -\tfrac{1}{2}\sqrt{\Delta^2 + 4g^2},$$
$$P_+^{(-)} = \tfrac{1}{2}\left(1 - \frac{\Delta}{\sqrt{\Delta^2 + 4g^2}}\right).$$

At zero detuning both basis states have equal weight. At large positive detuning the positive basis state belongs mainly to the upper branch. These are stationary-state probabilities at the chosen setting. Preparing that state by moving the setting additionally requires a sweep slow enough compared with the gap.

WORKED EXAMPLE

Estimate the minority component of a spectroscopic state

Change the inputs and investigate this derivation.

For $g = 0.150$ meV and $\Delta = 0.400$ meV, the measured branch separation is 0.500 meV. The lower slope is -0.400 meV per meV, so

$$P_+^{(-)} = 0.100.$$

A projective measurement resolving the original basis should yield about 100 positive outcomes in 1000 independent preparations of the lower state.

SECOND APPLICATION

Infer a coupling without reaching the crossing

Change the inputs and investigate this derivation.

At a second setting a fitted lower-branch slope is -0.300 while the energy gap is 0.800 meV. The measured slope requires $\Delta/\sqrt{\Delta^2 + 4g^2} = 0.600$, giving $\Delta = 0.480$ meV. Squaring the gap then yields

$$g = \tfrac{1}{2}\sqrt{0.800^2 - 0.480^2} = 0.320\,\mathrm{meV}.$$

The predicted minimum gap is 0.640 meV; a separate zero-detuning scan checks the inference.

Try it · Level population

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Coupling g (meV) **0.3**

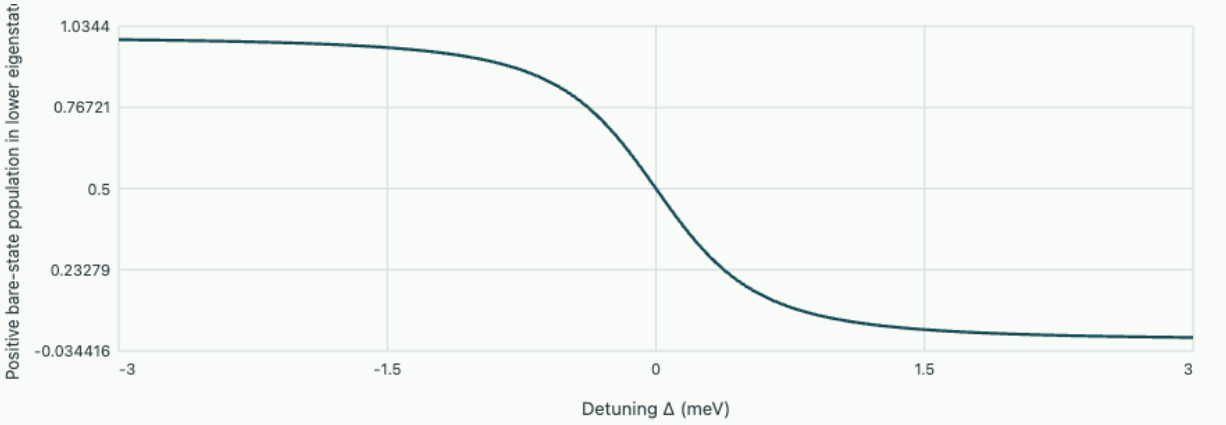

Lower-level bare-state weight

Hellmann–Feynman: ∂E−/∂Δ = P+−1/2. Population determines the measured energy slope even when a small-perturbation expansion fails.

Predict a weak response from a known spectrum

A tunable control can move two energy levels towards one another. Far apart, weak coupling shifts them only slightly. Near coincidence the observed spectrum can change substantially even if the same coupling is small on every other energy scale. Begin with a Hamiltonian $\hat{H}_0$ with known orthonormal eigenstates $|n\rangle$ and energies E_n^0. A weak time-independent perturbation $\lambda\hat{V}$ changes both energies and states, with λ a bookkeeping parameter eventually set to one.

Insert expansions $E_n = E_n^0 + \lambda E_n^{(1)} + \ldots$ and $|\psi_n\rangle = |n\rangle + \lambda|\delta n\rangle + \ldots$ into the eigenvalue equation. At first order $\hat{H}_0|\delta n\rangle + \hat{V}|n\rangle = E_n^0|\delta n\rangle + E_n^{(1)}|n\rangle$. Projection onto $\langle n|$ cancels the two terms containing the unperturbed Hamiltonian and gives

$$E_n^{(1)} = \langle n|\hat{V}|n\rangle.$$

This shift is a weighted average of a spatial perturbation in the original eigenstate.

Follow the change in the state

Projection onto a different eigenstate $\langle m|$ instead gives

$$\langle m|\delta n\rangle = V_{mn}/(E_n^0 - E_m^0).$$

The dimensionless ratio of the coupling matrix element to the level separation controls the amount of mixing. Substitution of this corrected state at the next order gives

$$E_n^{(2)} = \sum_{m\neq n} \frac{|V_{mn}|^2}{E_n^0 - E_m^0}.$$

A nondegenerate ground state has only higher-energy neighbours and therefore a nonpositive second-order shift. A vanishing diagonal element can make this second-order shift the leading effect. For example a symmetric bound state under a perturbation odd in position has zero first-order average. That fact alone does not establish that the state or its energy is unchanged.

As the energy separation approaches zero, the small-ratio condition fails. Apparent divergences signal that choosing one unperturbed vector as a unique starting state has become inappropriate. Within a degenerate subspace, diagonalise the perturbation first. Couplings to well-separated states outside the subspace can still be treated by ordinary perturbation theory if their matrix elements remain small compared with the external gaps.

Solve a pair exactly

Use two uncoupled levels $-\Delta/2$ and $+\Delta/2$, referenced to their midpoint, and take the off-diagonal coupling v real by a phase choice. The characteristic equation of the two-by-two Hamiltonian is

$$E^2 - (\Delta/2)^2 - v^2 = 0.$$

Thus $E_\pm = \pm\sqrt{(\Delta/2)^2 + v^2}$. At zero detuning the gap is $2|v|$. The equal-weight combinations of the two basis states diagonalise the Hamiltonian, even when the coupling would have been negligible far from coincidence.

Expanding the positive root for $|v| \ll |\Delta|$ gives $|\Delta|/2 + v^2/|\Delta|$ to leading correction; the lower level shifts down by the same amount. Their sum equals the original trace for every detuning. A measured change in the midpoint as well as the gap could therefore require a common diagonal shift in addition to this off-diagonal coupling.

Read coupling from an avoided crossing

Let $\Delta = 10.0$ meV and $v = 1.0$ meV. The unperturbed energies are ± 5.0 meV and the exact energies are ± 5.099 meV. The second-order prediction of a 0.100 meV shift is accurate to about 0.001 meV. At $\Delta = 1.0$ meV the exact energies are ± 1.118 meV and the perturbative correction $v^2/\Delta = 1.0$ meV is already larger than the original positive energy. Diagonalising the pair remains well conditioned there.

A spectrometer measuring the smallest separation in a detuning sweep would infer $|v| = 1.0$ meV from a minimum gap of 2.0 meV. Level-dependent transition strengths can supply a separate test: as the eigenstates exchange their dominant basis character across the avoided crossing, a probe coupling mainly to one basis state will redistribute intensity between the two observed lines.

An alternative experimental strategy fits states well away from the mixing region first. There the slopes with respect to the control parameter reveal the approximate unperturbed levels. Extrapolating those lines locates the crossing that would occur if the coupling vanished. Measuring the two actual levels near that point then separately determines the energy origin and the magnitude of the off-diagonal interaction. A phase change of either basis state reverses the sign of a real coupling, leaving both energies unchanged. Its sign relative to another amplitude requires an interference measurement.

Sweep through an avoided crossing

A spectrum measured at a fixed control setting reveals a gap. To follow a state while turning the control, we must also know how quickly its instantaneous eigenvector rotates. Consider smooth $H(t)$ and two nondegenerate eigenstates. Differentiate $H|n\rangle = E_n|n\rangle$ and take the inner product with another eigenstate m. The term containing $\dot{E}_n$ vanishes by orthogonality, giving

$$\langle m|\dot{n}\rangle = \frac{\langle m|\dot{H}|n\rangle}{E_n - E_m}.$$

When the state is expanded in the instantaneous basis with its dynamical phases, this coupling drives the amplitude of m. The relative phase oscillates at angular frequency $|E_n - E_m|/\hbar$. A slowly changing coupling then averages out over many oscillations. The scale $\eta = \hbar|\langle m|\dot{H}|n\rangle|/|E_m - E_n|^2$ compares basis rotation with that phase rate. For a smooth isolated sweep, keeping it small throughout is a useful adiabaticity test. The final probability also depends on the full drive and preparation.

Take $H = \Delta(t)\sigma_z/2 + g\sigma_x$ with fixed $g > 0$ and detuning increased linearly through zero. The gap is $\sqrt{\Delta^2 + 4g^2}$. At the centre the eigenstates point along the x direction in spin space, so the off-diagonal matrix element of $\dot{H} = (\dot{\Delta}/2)\sigma_z$ has magnitude $|\dot{\Delta}|/2$. This is also where the sweep is most demanding:

$$\eta(0) = \frac{\hbar|\dot{\Delta}|}{8g^2}.$$

With $g = 0.100\,\mathrm{meV}$, $\hbar = 0.6582\,\mathrm{meV\,ps}$ and $|\dot{\Delta}| = 0.00100\,\mathrm{meV/ps}$, the minimum splitting is $0.200\,\mathrm{meV}$ and $\eta(0) = 0.00823$. Starting sufficiently far away in the lower eigenstate, a slow sweep transfers the occupation to the opposite uncoupled basis state as the lower eigenvector exchanges character. Multiplying the rate by twenty gives $\eta(0) = 0.165$ and a weaker separation of time scales. At $g = 0$ the crossing is exact and this finite-gap argument breaks down.

Practise on sheet 3 · Set the speed of a control ramp. Further reading · Massachusetts Institute of Technology · 8.06 Quantum Physics III · Chapter 6, Adiabatic approximation, §2 Adiabatic Evolution and §4 Landau-Zener transitions.

Apply degenerate perturbation theory to hydrogen in an electric field

A nondegenerate atomic parity eigenstate has zero permanent electric dipole along z. It is tempting to infer that a weak uniform field always shifts an energy only at second order. For hydrogen in the nonrelativistic Coulomb model, however, the $2s$ and $2p$ levels are degenerate. A field $\mathcal{E}$ along z adds $V = e\mathcal{E}z$ for an electron, with positive e denoting the elementary charge. The perturbation is odd under inversion. Diagonal parity-basis elements vanish, but an off-diagonal matrix element can survive between opposite parities.

Choose $|2s\rangle$ and $|2p, m=0\rangle$, since $z = r\cos\theta$ preserves magnetic projection. With conventional real hydrogen phases, the angular integral $\langle Y_{00}|\cos\theta|Y_{10}\rangle = 1/\sqrt{3}$. Writing $\rho = r/a_0$, the normalized radial functions are proportional to $(2-\rho)e^{-\rho/2}$ and $\rho e^{-\rho/2}$. Including radial volume and the extra r from z gives

$$\int_0^\infty R_{20}R_{21}r^3 dr = \frac{a_0}{8\sqrt{3}}\int_0^\infty (2-\rho)\rho^4 e^{-\rho}d\rho = -3\sqrt{3}\,a_0.$$

Thus the off-diagonal position element is $-3a_0$. Its sign depends on orbital phases, while the magnitude and energy shifts do not. The perturbation matrix in this degenerate subspace has zero diagonal and equal off-diagonal entries $-3e\mathcal{E}a_0$. Its eigenvalues are $\pm 3e\mathcal{E}a_0$ with orthogonal equal superpositions of the two original orbitals. The remaining $m = \pm 1$ states have no first-order partners here and remain unshifted in this model.

Assess a resolvable Stark splitting

Let $\mathcal{E} = 2.00 \times 10^{6}$ V/m and $a_0 = 5.29 \times 10^{-11}$ m. The field-induced voltage across one Bohr radius is 1.058×10^{-4} V, giving each mixed level an energy shift magnitude 0.317 meV. Their separation is 0.635 meV or 0.154 THz. This is far smaller than the electronvolt-scale separation from other principal shells, supporting restriction to the $n = 2$ manifold. The absolute laboratory field, line widths and finer internal splittings still set the domain where the ideal doublet picture is appropriate.

If the two unperturbed orbitals are split by a small residual gap δ, the new separation is $\sqrt{\delta^2 + 4(3e\mathcal{E}a_0)^2}$. At extremely weak field the change from $|\delta|$ is then quadratic, crossing to approximately linear splitting once the off-diagonal coupling dominates. Varying the field over this crossover can distinguish truly degenerate first-order splitting from a barely resolved avoided crossing without assigning a permanent dipole to either original parity orbital.

The slopes of the split energies also have an observable interpretation. By differentiating a normalized eigenvalue with respect to the applied field, the Hellmann–Feynman relation gives $dE/d\mathcal{E} = e\langle z\rangle$ in the same sign convention. The two mixed states therefore have opposite mean electron displacements of magnitude $3a_0$ in the degenerate first-order model. Their charge distributions are oriented along or against the applied field although neither original parity state had an orientation. Following a field reversal requires care about which eigenbranch or prepared state is being tracked when a much smaller unresolved zero-field splitting is present.

Measuring slopes on both sides of field reversal tests the branch assignment independently of any common instrumental drift in the energy reference.

Read a coupling from an avoided crossing

Set the two uncoupled energies to $\pm d/2$ and switch on a real constant matrix element 0.30 in the plotted energy units. The characteristic equation gives

$$E^2 = (d/2)^2 + 0.30^2.$$

Follow the eigenvalue as detuning changes sign.

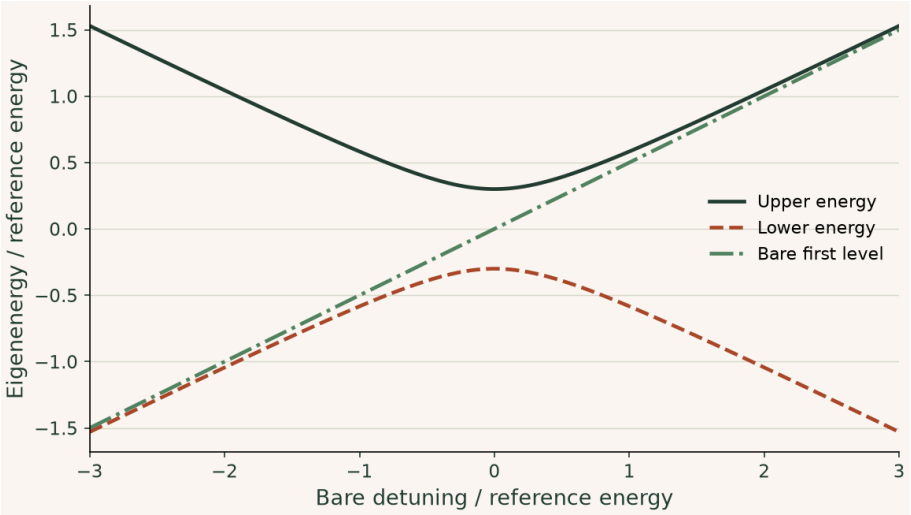

The closest energy separation is 0.60, twice the coupling. Far from the crossing ordinary second-order shifts are small. At zero detuning a perturbation formula with the bare level separation in the denominator breaks down.

Try it · Level mixing

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Coupling (energy units) **0.3**

Selected detuning **0**

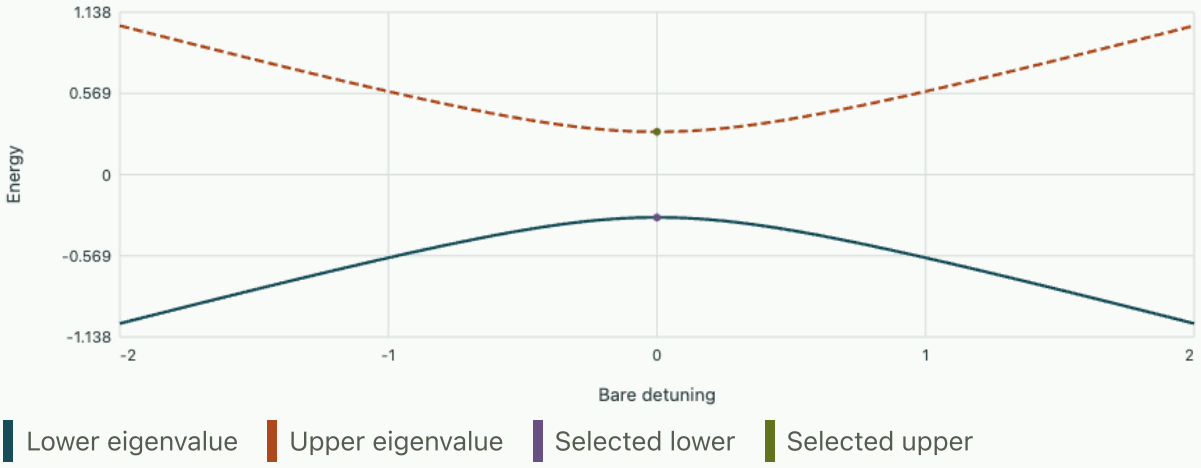

Minimum gap = 2g = 0.6. At nonzero coupling the eigenvectors exchange bare-state character through the avoided crossing.

▶ **Read values from the plot**

Study the full derivation

From a Coulomb orbit to the fine-structure multiplet · Quantum states, dynamics and measurement

Separate the hydrogenic Schrödinger equation, identify the origin of its large degeneracy and evaluate the three dominant relativistic corrections. The cancellation among them predicts an equality that spectroscopy can then test.

Chapter 6 · Composite states and spin correlations

Predict joint and conditional spin measurements for two particles and distinguish an entangled state from a mixture with the same one-particle statistics.

FROM FIRST PRINCIPLES

Build a joint spin probability from a two-particle amplitude

Two distinguishable spin-one-half measurements have a product basis $|\uparrow\uparrow\rangle, |\uparrow\downarrow\rangle, |\downarrow\uparrow\rangle, |\downarrow\downarrow\rangle$. In a singlet the normalized state is $(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)/\sqrt{2}$. Measure the first spin along z and the second along an axis tilted by θ in the xz plane. A positive eigenstate along the new axis is

$$|+\theta\rangle = \cos(\theta/2)|\uparrow\rangle + \sin(\theta/2)|\downarrow\rangle.$$

The first positive outcome selects the singlet's first product term. Its remaining second-particle amplitude along the tilted axis is the overlap with $|\downarrow\rangle$. The Born rule for the joint projector now gives

$$P(+z,+\theta) = |\langle\uparrow,+\theta|\psi^-\rangle|^2 = \tfrac{1}{2}\sin^2(\theta/2) = \tfrac{1}{4}(1-\cos\theta).$$

For the other tilted outcome the squared overlap is $\cos^2(\theta/2)$. Adding the two joint outcomes gives $P(+z) = 1/2$. Dividing the joint result by this marginal yields the conditional probability for the partner. The same single-particle count occurs at all analyser angles even though the coincidence pattern changes.

DERIVATION AND APPLICATIONS

Obtain the two-spin coincidence law from projectors

A pair prepared in a spin singlet has no preferred axis, yet coincidence counts change when one analyser rotates. Start with two spin-one-half particles in $|\psi^-\rangle = (|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)/\sqrt{2}$. Ideal spin measurements along unit vectors $\mathbf{a}$ and $\mathbf{b}$ have recorded outcomes $s, t = \pm 1$, in units of $\hbar/2$. Neglect detector losses. Since the eigenvalues of $\mathbf{a}\cdot\boldsymbol{\sigma}$ are ± 1, its projector onto result s is

$$\Pi_s(\mathbf{a}) = \frac{I + s\,\mathbf{a}\cdot\boldsymbol{\sigma}}{2},$$

$$P(s,t) = \langle\psi^-|\Pi_s(\mathbf{a})\otimes\Pi_t(\mathbf{b})|\psi^-\rangle.$$

Tracing over either particle gives $I/2$, so terms linear in one spin have zero expectation. To evaluate the remaining product, choose both axes initially along z. The singlet basis components always have opposite outcomes, giving $\langle\sigma_z\otimes\sigma_z\rangle = -1$. Rotational invariance of the singlet requires its two-index correlation matrix to be proportional to the identity. The zz result fixes the coefficient:

$$\langle\sigma_i\otimes\sigma_j\rangle = -\delta_{ij},$$

$$\langle(\mathbf{a}\cdot\boldsymbol{\sigma})\otimes(\mathbf{b}\cdot\boldsymbol{\sigma})\rangle = -\mathbf{a}\cdot\mathbf{b}.$$

Expand the product of projectors and substitute this correlation. For angle θ between the analysers the full coincidence probabilities and their sum over the second analyser become

$$P(s,t) = \tfrac{1}{4}(1 - st\cos\theta),$$
$$\sum_t P(s,t) = \tfrac{1}{2}.$$

Parallel axes forbid equal outcomes. Perpendicular axes give four equally likely outcomes despite the pure preparation. Local counts alone reveal neither the relative angle nor the phase coherence of the joint state.

WORKED EXAMPLE

Predict a coincidence histogram

Change the inputs and investigate this derivation.

With 1200 independently prepared pairs and a 60.0° angle, the equal-outcome cells each have probability $(1 - 0.5)/4 = 0.125$ and the opposite-outcome cells each have probability 0.375. The expected counts, ordered as $++, +-, -+, --$, are $150, 450, 450, 150$. Given a positive result for the first analyser, the probability of a negative result at the second is

$$450/(150 + 450) = 0.750.$$

SECOND APPLICATION

Separate a singlet from a classical preparation

Change the inputs and investigate this derivation.

A source instead produces an equal incoherent mixture of the two product states anti-aligned along z. Both sources are perfectly anti-correlated along that axis. Align both analysers along x. Each product preparation now gives independent, equally probable outcomes, so $P(\text{equal}) = 0.500$. The singlet predicts zero equal outcomes. For 200 ideal pairs the predictions are 100 versus zero equal coincidences, while either local detector still reports each sign about half of the time.

Try it · Singlet

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Analyser angle (degrees) **60**

Independent pairs recorded **1200**

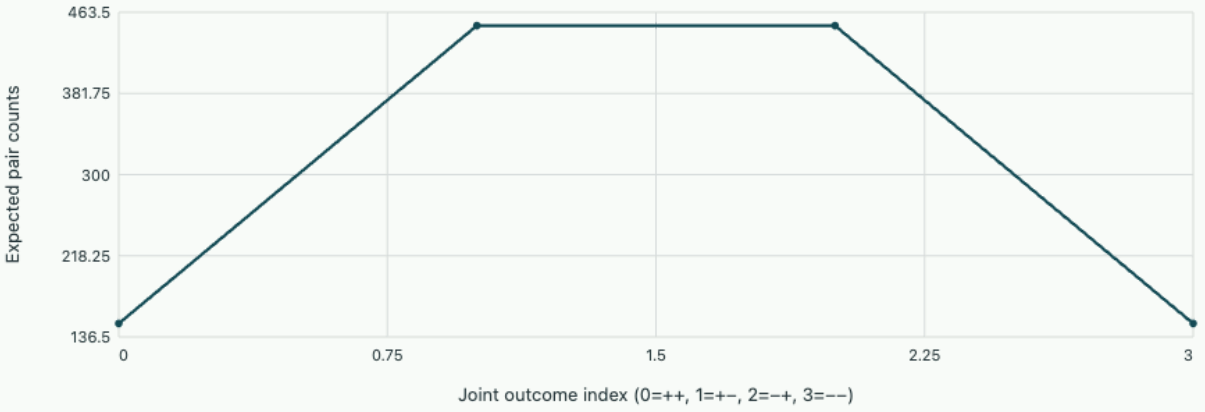

Spin singlet

Each single-analyser marginal is one half; joint correlation = −cosθ = -0.5. Compare rotated analysers to distinguish the singlet from an incoherent anti-aligned mixture.

Build the state space of a pair

A spin-½ measurement along a chosen axis yields two possible outcomes. Label its dimensionless results $+1$ and -1, corresponding to spin angular momenta $\pm\hbar/2$. For two separately addressable particles A and B, a joint state belongs to the tensor product of their two spin spaces. In the z basis we need four product vectors: both up, A up and B down, A down and B up, or both down. A general normalised state contains four complex amplitudes whose squared magnitudes sum to one.

An uncorrelated pure preparation can be written $|\psi_A\rangle \otimes |\psi_B\rangle$. Its product-structure forces a relation among the four amplitudes. For example the state $(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)/\sqrt{2}$ has two nonzero amplitudes for opposite spins and zero amplitude for equal spins. No product of two individual pure states has that pattern. This entangled state is the spin singlet. Spin labels here refer to the internal states at the two measurement locations.

Project along the same and different axes

Measure both particles along z. The two opposite outcomes have probability one half each, and the product of their dimensionless outcomes always equals -1. Alice's result alone is random. Conditional on her obtaining $+1$, the corresponding outcome for Bob along the same axis is definitely -1. A classical mixture of the two opposite z product states with equal weights would give exactly the same statistics in this one measurement arrangement.

To distinguish the preparations, rotate Bob’s analyser through angle θ in the xz plane while Alice continues to measure z. A spin initially up along z has probability $\cos^2(\theta/2)$ of being up along Bob’s axis and $\sin^2(\theta/2)$ of being down along it. For the singlet, conditional on Alice's $+1$ outcome, Bob starts in the opposite state. His probabilities are therefore interchanged. Summing over Alice's equally likely outcomes gives the joint correlation $E(z,b) = -\cos\theta$. Its single-party marginals remain one half.

Find the coherence that an ensemble can hide

Rotational invariance of the singlet gives $E(a,b) = -\hat{\mathbf{a}}\cdot\hat{\mathbf{b}}$ for any chosen axes. In particular if both measure x, their results are again perfectly anticorrelated. Verify this directly by expressing $|{\uparrow_z}\rangle = (|{\uparrow_x}\rangle + |{\downarrow_x}\rangle)/\sqrt{2}$ and $|{\downarrow_z}\rangle = (|{\uparrow_x}\rangle - |{\downarrow_x}\rangle)/\sqrt{2}$. The two equal-$x$ amplitudes cancel in the superposition.

In the classical mixture of the two opposite z states those amplitudes cannot cancel across different ensemble members. Each particle measured along x is individually random and the two x outcomes are independent conditional on either preparation. The mixture consequently has $E(x,x) = 0$. Measuring repeated pairs in two bases reveals an off-diagonal coherence invisible to either single-party measurement. Local outcome statistics alone stay random regardless of the distant analyser setting.

Predict counts and a Bell combination

At $\theta = 60^\circ$ the singlet gives $E = -1/2$. The probability of opposite outcomes is

$$(1 - E)/2 = 0.75.$$

For 400 ideal detected pairs one therefore expects about 300 opposite and 100 equal outcomes, with ordinary statistical fluctuations. Rotating both analysers together preserves these fractions because only their relative angle matters for a singlet.

The distinction can be sharpened with four settings. In one plane choose Alice’s angles 0° and 90° and Bob’s 45° and −45°. Form $S = E(a,b) + E(a,b') + E(a',b) - E(a',b')$. The three first correlations are $-1/\sqrt{2}$ and the last is $+1/\sqrt{2}$, giving $|S| = 2\sqrt{2}$. A model assigning local predetermined binary outcomes to all four settings gives for each emitted pair a sum $A(B + B') + A'(B - B')$ of magnitude two; averaging bounds $|S|$ by two if the sampled ensemble is independent of the choices of settings. A laboratory interpretation must also account for detection statistics and the sampling protocol. The calculation concerns joint correlations and grants no control over the individual random outcomes.

If one observer discards the other observer’s records, the local spin outcomes remain equally likely for each measurement axis. Sorting the same data into subsets according to the partner’s outcomes reveals the

correlations. This sorting requires comparing the two records through an ordinary communication channel.

Classify a pure bipartite state by its Schmidt weights

The spin singlet has maximally mixed individual spins even though its joint state is pure. A partially entangled pair allows a continuous test of that distinction. Any pure state of two finite systems can be written as $\sum_{ij} A_{ij}|i\rangle|j\rangle$. Tracing out the second system gives

$$\rho_A = AA^\dagger.$$

Singular-value decomposition rotates the two local bases independently and writes the same state as $\sum_k \sqrt{\lambda_k}|u_k\rangle|v_k\rangle$, where $\lambda_k \geq 0$ and their sum is one. The reduced states on both sides have the same nonzero eigenvalues λ_k.

A pure product state has just one nonzero Schmidt weight. Two or more weights imply entanglement, because separate rotations of the local bases cannot remove singular values. The reduced purity is $\sum_k \lambda_k^2$, and the entanglement entropy for a pure pair is $-\sum_k \lambda_k \ln \lambda_k$. A nonzero reduced entropy describes uncertainty in local observations. It can occur jointly with zero entropy for the complete pure state. If the joint preparation is already mixed, reduced entropy alone is insufficient to quantify its entanglement.

Choose a measurement basis for the partner

Prepare $|\psi\rangle = \sqrt{0.90}|00\rangle + \sqrt{0.10}|11\rangle$. Its local populations are 0.90 and 0.10, giving reduced purity $0.90^2 + 0.10^2 = 0.820$ and entropy $-0.90\ln 0.90 - 0.10 \ln 0.10 = 0.325$ in natural-log units. Measuring the second spin in the computational basis prepares the first in $|0\rangle$ or $|1\rangle$ with those respective probabilities. Each branch has a definite computational outcome.

Now measure the partner in $|\pm\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}$. Both outcomes have probability one half. Conditional on the result, the first spin is in $\sqrt{0.90}|0\rangle \pm \sqrt{0.10}|1\rangle$, already normalized. A subsequent plus-basis measurement on the first spin gives plus with probability $[1 \pm 2\sqrt{0.90 \times 0.10}]/2$, namely 0.800 or 0.200. Without receiving the partner's result the first observer averages these two values and sees 0.500, exactly as predicted by the unchanged reduced matrix.

As the smaller Schmidt weight tends to zero the state approaches a product and the two conditional plus probabilities coalesce. Equal weights give maximal local entropy $\ln 2$. A classical mixture of $|00\rangle$ and $|11\rangle$ with the same 0.90 and 0.10 weights produces identical computational-basis statistics but gives 0.500 in both conditional plus measurements. The pair of rotated analysers therefore tests coherence between the alternatives rather than merely confirming their occupation probabilities.

Experiments verify the conditional prediction by collecting both detector outcomes and matching them trial by trial. If one sorts first-spin data without correctly identifying the partner from the same pair, the two 0.800 and 0.200 conditional probabilities both tend back toward the unconditional value of one half. This loss of contrast can thus result from timing or pairing errors even when each local detector has excellent single-spin fidelity. Repeating the computational-basis correlation measurement supplies an independent pairing check before interpreting the more delicate rotated-basis signal as a property of the prepared wavefunction.

Study the full derivation

Distinguishing entanglement from ordinary shared randomness · Quantum states, dynamics and measurement

Two distant analysers sometimes agree even when individual outcomes look random. Derive which correlations shared classical information can produce, compute the quantum prediction and identify what a finite table of coincidence counts can establish.

Problem sheet 3 · Small perturbations and paired measurements

Apply the eigenstate and measurement methods from Lectures 5 and 6 to spectral shifts and two-spin counts.

Diagnose an avoided crossing

Two uncoupled levels are tuned symmetrically about zero energy. Their observed minimum gap is 0.80 meV. Find the coupling magnitude, then the exact positive eigenvalue at a bare separation of 2.0 meV and its second-order estimate. Identify the observable that tests whether an additional common energy shift is present.

▼ **Extract the coupling at zero detuning**

The minimum splitting is $2|v|$, hence $|v| = 0.40$ meV. At bare energies ± 1.0 meV the upper eigenvalue is

$$\sqrt{1.0^2 + 0.40^2} = 1.077\,\mathrm{meV}.$$

The perturbative result is

$$1.0 + (0.40)^2/2.0 = 1.080\,\mathrm{meV}.$$

Coupling alone preserves the sum of the two energies; measuring the midpoint as well as their difference can reveal a diagonal shift common to the two states.

Distinguish an incoherent source

A pair source produces a singlet with probability 0.70 and an equal classical mixture of opposite z spins otherwise. Predict $E(z,z)$ and $E(x,x)$, then the expected fraction of opposite x outcomes. Can either one-particle z histogram reveal the admixture?

▼ **Average over preparations**

Both preparations always give opposite z outcomes, hence $E(z,z) = -1$. The singlet gives $E(x,x) = -1$ and the incoherent mixture gives zero. Their weighted correlation is -0.70, so opposite x outcomes occur with probability $(1 + 0.70)/2 = 0.85$. Both one-particle marginals are evenly split in the z basis for each preparation. Detecting the admixture therefore requires appropriate joint measurements.

Mix a pair of distant levels weakly

An unperturbed lower state has energy 0 and an upper state 50 meV. A perturbation has zero diagonal elements and off-diagonal magnitude 2.0 meV. Estimate the lower-state energy shift and the probability that a measurement of the original energy on the perturbed ground state finds the upper state. How do both estimates change if the perturbation doubles?

▼ **Calculate mixing before squaring it**

The first-order diagonal energy shift is zero. The lower state has a first-order admixture of upper state of magnitude $2/50 = 0.040$, giving upper-outcome probability about 0.0016. The lower energy shifts by $-2^2/50 = -0.080$ meV at second order. Doubling coupling doubles the amplitude, hence quadruples the probability to approximately 0.0064 and the downward shift to approximately 0.320 meV. Their common quadratic scaling holds only while the mixing ratio remains small.

Set the speed of a control ramp

A coupled pair has minimum splitting $0.400\,\mathrm{meV}$. Set a constant detuning speed that keeps $\eta(0) \le 0.0200$. Estimate the shortest time for a $2.00\,\mathrm{meV}$ ramp at this speed. How must the speed change if the minimum gap is halved?

▼ **Determine the limiting speed and ramp duration**

The coupling is half the minimum splitting, $g = 0.200\,\mathrm{meV}$. Solving the centre criterion gives

$$|\dot{\Delta}| \le 8(0.200)^2(0.0200)/0.6582 = 0.00972\,\mathrm{meV/ps}.$$

A constant-speed ramp through $2.00\,\mathrm{meV}$ therefore takes at least $206\,\mathrm{ps}$. Halving the gap halves g and reduces the allowed speed to a quarter, $0.00243\,\mathrm{meV/ps}$. These numbers set a local adiabaticity margin. An exact transition probability would require specifying the start, end and switching of the sweep.

Revisit chapter 5 · Sweep through an avoided crossing.

Explore these relations

- For the avoided crossing, set coupling to 0.40 and detuning to 2.00 in consistent meV units and compare the exact eigenvalues. Open the live calculation.
- For the control ramp, vary sweep rate and coupling and compare adiabatic following with transition probability. Open the live calculation.

Chapter 7 · Estimating a ground-state energy

Construct a variational upper bound, optimize the length of a trial orbital and evaluate what the bound says about its shape.

FROM FIRST PRINCIPLES

Why minimizing a trial energy gives an upper bound

Take a self-adjoint Hamiltonian bounded below, with ground energy E_0. Choose a nonzero trial state obeying the physical boundary conditions and having finite energy expectation. For a discrete spectrum, expand it in normalized eigenstates $|n\rangle$. Orthonormality removes cross terms when the Hamiltonian acts on this expansion:

$$|\psi\rangle = \sum_n c_n |n\rangle,$$

$$\frac{\langle\psi|H|\psi\rangle}{\langle\psi|\psi\rangle} = \frac{\sum_n |c_n|^2 E_n}{\sum_n |c_n|^2}.$$

Subtract E_0. Each term now contains a nonnegative weight times a nonnegative excitation energy, giving

$$E[\psi] - E_0 = \frac{\sum_n |c_n|^2 (E_n - E_0)}{\sum_n |c_n|^2} \geq 0.$$

A continuum contribution gives the same inequality with an integral. Optimizing a trial length or mixing coefficient decreases the smallest accessible value of this quotient. Equality requires all weight to lie in the ground eigenspace; a low estimated energy alone does not specify which local features of the wave function are accurate. Checking a width, contact density or other independent observable tests the adequacy of the chosen trial family.

Background · Variational energy

$$E[\psi] = \langle\psi|H|\psi\rangle / \langle\psi|\psi\rangle$$

For an admissible nonzero trial state and a Hamiltonian bounded below, this is an upper bound on the lowest energy. A good energy can hide a less accurate state.

DERIVATION AND APPLICATIONS

Optimize a superposition when the trial functions overlap

A single variational length can improve the size of a trial orbital, but it cannot reproduce every possible shape. A second trial function introduces a mixing coefficient. Frequently two convenient functions are individually normalized yet overlap. Let the linearly independent functions be ϕ_1, ϕ_2, let $H_{ij} = \langle\phi_i|H|\phi_j\rangle$, and define the overlap matrix $S_{ij} = \langle\phi_i|\phi_j\rangle$. In a trial superposition $\psi = \sum_i c_i \phi_i$, normalization depends on the cross term:

$$\langle \psi | \psi \rangle = c^{\dagger} S c,$$

$$E[c] = \frac{c^{\dagger} H c}{c^{\dagger} S c}.$$

To find stationary values, minimize $c^{\dagger} H c - E(c^{\dagger} S c - 1)$ with respect to the complex conjugates c_i^{*}. Differentiation of each bilinear form removes a single row factor, giving two homogeneous equations for the same coefficients:

$$\sum_j (H_{ij} - E S_{ij}) c_j = 0,$$

$$\det(H - E S) = 0.$$

Because S is positive definite, the quotient is well defined for every nonzero trial vector. Its smallest root is an upper bound to the exact ground energy. Adding an independent trial function cannot raise this optimum, since setting its coefficient to zero recovers the original trial state. For two equivalent sites with $H_{11} = H_{22} = h$, $H_{12} = H_{21} = v$ and real $S_{12} = s$, add and subtract the rows to obtain

$$E_{\mathrm{sym}} = (h + v)/(1 + s),$$

$$E_{\mathrm{anti}} = (h - v)/(1 - s).$$

The signs label symmetry, not a universal ordering. The denominator records the actual norm of each combination. As $|s| \to 1$ one combination becomes nearly redundant; estimating its energy then becomes sensitive to small matrix-element errors.

WORKED EXAMPLE

Estimate two trial energies from overlapping orbitals

Change the inputs and investigate this derivation.

Take $h = 3.00$ eV, $v = -0.500$ eV and $s = 0.250$. The symmetric combination yields $2.50/1.25 = 2.00$ eV; the antisymmetric combination gives

$$3.50/0.750 = 4.67\,\mathrm{eV}.$$

The first improves the single-orbital bound by 1.00 eV. Normalization of its coefficient vector gives $c_1 = c_2 = 1/\sqrt{2(1+s)} = 0.632$, even though the basis functions are individually normalized.

SECOND APPLICATION

Change the energy reference to test an implementation

Change the inputs and investigate this derivation.

Shift the Hamiltonian by CI with $C = 2.00$ eV. In the overlapping basis this changes the matrix to $H + CS$, because the matrix elements of the identity are S_{ij}. The new diagonal is 5.00 eV and the off-diagonal is zero. The two roots become

$$5.00/1.25 = 4.00,$$

$$5.00/0.750 = 6.67\,\mathrm{eV}.$$

Both shifted by exactly 2.00 eV; leaving the off-diagonal unchanged would fail this energy-reference test.

Try it · Overlap eigen

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Orbital overlap s **0.25**

Energy-reference shift C (eV) **0**

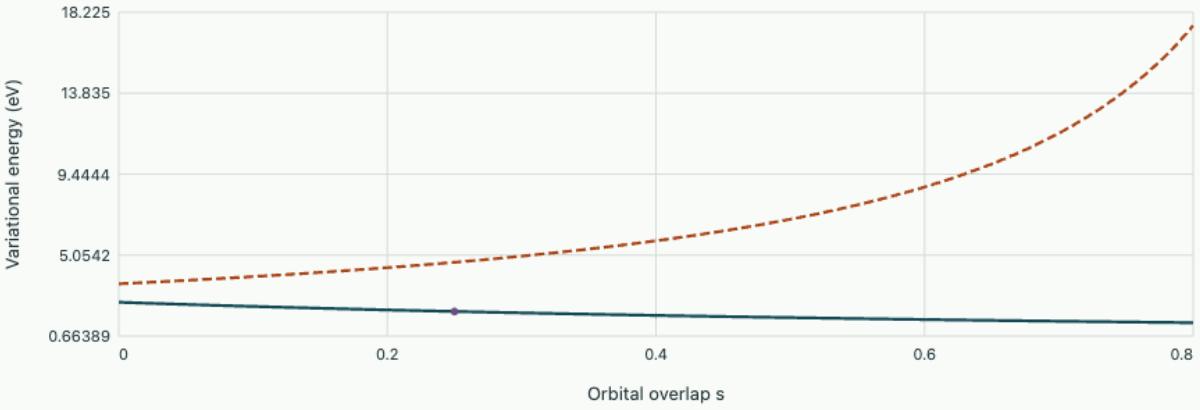

Symmetric generalised eigenstate

Antisymmetric generalised eigenstate

Selected symmetric energy

Unshifted H diagonals 3 eV and off-diagonal −0.5 eV. An energy shift is H→H+CS, so both generalised eigenvalues shift by exactly C.

Start from this chapter. The lowest eigenvalue minimises the Rayleigh quotient. Compare refinements for a polygonal domain and convert the converged spatial eigenvalue to energy using the particle mass and a physical length.

OneSimulation laboratory · Turn spatial eigenvalues into confinement energies

Can changing a Gaussian width repair a missing cusp?

A good variational parameter can optimise the extent of a bound state without correcting its local shape. A one-dimensional attractive contact has a precisely known derivative jump. Compare two trial families and test which aspect of their shape the energy notices.

Start from the jump that the potential enforces

Use length and energy units in which $\hbar = m = 1$, and take $V(x) = -g\delta(x)$ with $g > 0$. Integrating the stationary Schrödinger equation between $-\epsilon$ and $+\epsilon$ and taking the limit gives

$$\psi'(0^+) - \psi'(0^-) = -2g\psi(0).$$

The wavefunction itself is continuous. Outside the contact a bound eigenstate decays as $e^{-\kappa|x|}$, and the derivative jump fixes $\kappa = g$. Its exact energy is $-g^2/2$.

For a normalised admissible trial function the quadratic form is $E[\psi] = \frac{1}{2}\int |\psi'|^2 dx - g|\psi(0)|^2$. This form permits smooth trials that do not obey the eigenstate jump: they have finite energy and give legitimate upper bounds. Consider the normalised Gaussian $\psi_b = (b^2/\pi)^{1/4} e^{-b^2x^2/2}$ with $b > 0$. Its mean x^2 is $1/(2b^2)$, and its derivative is $-b^2 x\psi_b$. The two contributions yield

$$E_G(b) = \frac{b^2}{4} - \frac{gb}{\sqrt{\pi}},$$

$$b_{\mathrm{opt}} = \frac{2g}{\sqrt{\pi}},$$

$$E_G^{\min} = -\frac{g^2}{\pi}.$$

Compare with $\psi_a = \sqrt{a}e^{-a|x|}$. Here the two half-lines both contribute to a total kinetic term $a^2/2$ and the contact samples $|\psi_a(0)|^2 = a$. Minimising $a^2/2 - ga$ gives $a = g$, the exact solution. For $g = 1$ the Gaussian optimum is about -0.31831, whereas the exact value is -0.5. The negative Gaussian energy suffices to prove binding below the free continuum. Its magnitude still underestimates the binding energy by about 36 per cent.

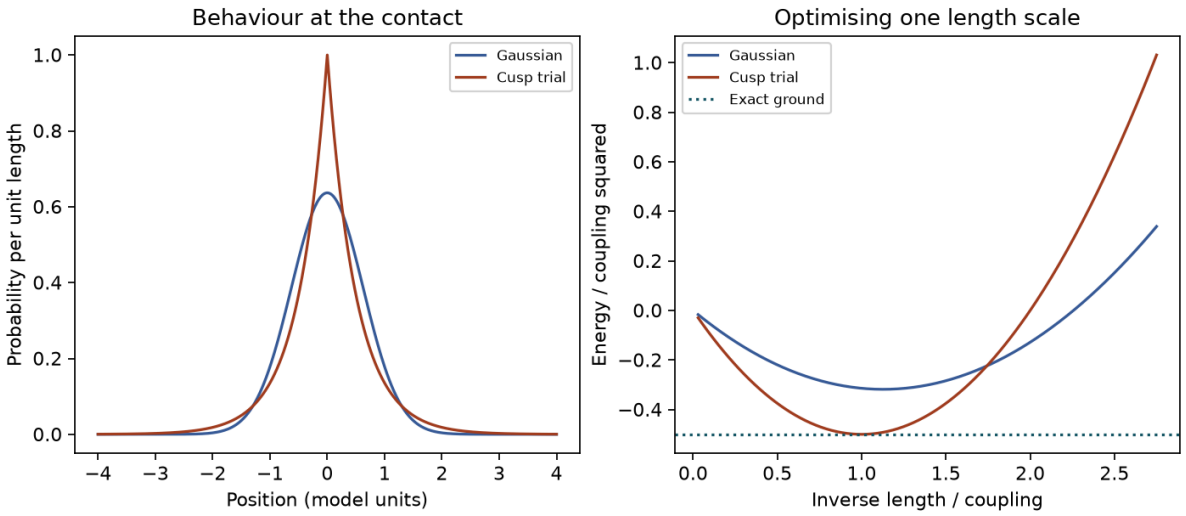

Both energy curves are upper bounds. Only the family capable of reproducing the correct derivative jump reaches the exact ground energy.

Try it · Cusp trial

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Attractive contact strength 1

Inverse-length scale / optimum 1

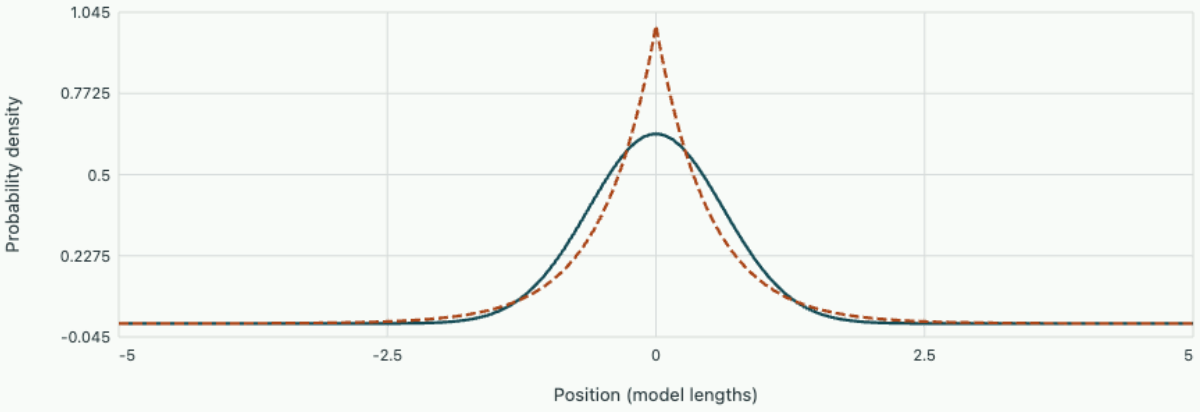

Gaussian / cusp trial energies -0.31831 / -0.5. Exact ground energy $-g^2/2$ = -0.5. Only the cusp family can satisfy the contact derivative jump.

Download the corresponding editable notebook. The browser recalculates the reduced equations directly; the notebook exposes the full calculation and its checks.

The notebook separates trial-width error from the limitation of the trial family. Retune both widths together and observe the quadratic rise near each minimum. At the exact cusp width, a ten per cent fractional width error raises the cusp energy by just one per cent of the binding energy. That energetic insensitivity does not mean that contact observables are equally accurate: the central probability density changes linearly with a. A variational energy can therefore look well converged before a highly local observable has converged.

An optimiser gives $E = -0.315$ at $g = 1$ within the Gaussian family. Could a better scalar width alone bring it within one per cent of the true binding energy? What further check would distinguish optimisation error from model error?

▼ **Work through the solution**

The minimum possible within the family is $-1/\pi$, already far above -0.495, the target upper bound for one-per-cent accuracy. Optimising b can recover only the difference between -0.315 and -0.31831. Check the

analytic minimum or sample the curvature around it. Then enrich the family with a cusp degree of freedom and compare both the energy and the derivative jump or contact density.

Open editable notebook Download .ipynb

A calculation that can be wrong in one direction

An exact time-independent Schrödinger equation can be difficult to solve even when the Hamiltonian is known. The variational principle allows a trial state to make a precise claim about the lowest energy. Suppose the Hamiltonian is self-adjoint and bounded below. Expand any normalized trial state of finite energy in its energy eigenbasis. For a discrete spectrum, with continuum terms replaced by integrals when needed,

$$\langle H\rangle_\psi = \sum_n |c_n|^2 E_n \geq E_0 \sum_n |c_n|^2 = E_0.$$

Lowering an expectation value inside the trial family improves the upper bound. Trial functions must obey the physical boundaries and give finite relevant energy integrals. A smooth function confined to the wrong symmetry sector can converge to the lowest state of that sector while missing the true ground state. To estimate an excited level one needs additional orthogonality or a suitable subspace principle.

OneSimulation · How much can a larger trial space lower an energy?

Confine an electron to an infinite 10 nm well and apply a linear potential rising by twelve untilted ground-state energies from left to right. Expand the wavefunction in the first N sine functions of the untilted box. They satisfy both boundary conditions exactly. A larger basis contains every earlier trial function, so the lowest Ritz energy can only decrease. The min–max principle gives the corresponding bound for the ordered excited levels.

$$H_{mn} = E_0 n^2 \delta_{mn} + \Delta V \int_0^L \frac{2}{L} \sin\frac{m\pi x}{L}\,\frac{x}{L}\,\sin\frac{n\pi x}{L}\,dx,$$

$$E_0 = \frac{\hbar^2\pi^2}{2m_e L^2}.$$

OneSimulation diagonalises nested sections of this matrix. The two errors are measured against a 64-function reference, itself a variational approximation. Successive functions improve the ability to concentrate probability near the low-potential wall. Change the tilt to see how the basis size required for a given energy accuracy changes.

```
"""Rayleigh–Ritz bounds for an electron in a tilted infinite square well."""
import matplotlib.pyplot as plt
import numpy as np
from onesim.numerics.linear.matrices.dense_linear_algebra import symmetric_eigenvalues

width_nm = 10.0
hbar, mass, electronvolt = 1.054571817e-34, 9.1093837e-31, 1.602176634e-19
energy_unit = 1e3 * np.pi**2 * hbar**2 / (2 * mass * (width_nm * 1e-9)**2 * electronvolt)
tilt = 12.0  # end-to-end potential rise in units of the untilted ground energy
basis = np.arange(1, 65)
i, j = basis[:, None], basis[None, :]
```

```
position = np.eye(len(basis)) / 2  # matrix of x/L in the sine basis
odd = (i - j) % 2 == 1
position[odd] = 2 / np.pi**2 * (1 / (i+j)[odd]**2 - 1 / (i-j)[odd]**2)
hamiltonian = np.diag(basis**2) + tilt * position
reference = symmetric_eigenvalues(hamiltonian)[:2] * energy_unit
sizes = np.arange(2, 13)
energies = np.array([symmetric_eigenvalues(hamiltonian[:n, :n])[:2] for n in sizes]) * energy_unit
errors = energies - reference
assert np.all(errors > 0) and np.all(np.diff(energies, axis=0) <= 1e-10)
print("64-state reference energies (meV):", reference.round(4))
fig, ax = plt.subplots(figsize=(6.6, 3.5), layout="constrained")
for index, label in enumerate(("Ground state", "First excited state")):
    ax.semilogy(sizes, errors[:, index], "o-", label=label)
ax.set(xlabel="Number of sine functions", ylabel="Excess above 64-state reference (meV)")
ax.grid(alpha=0.2)
ax.legend()
fig.savefig("quantum-mechanics.png", dpi=180)
```

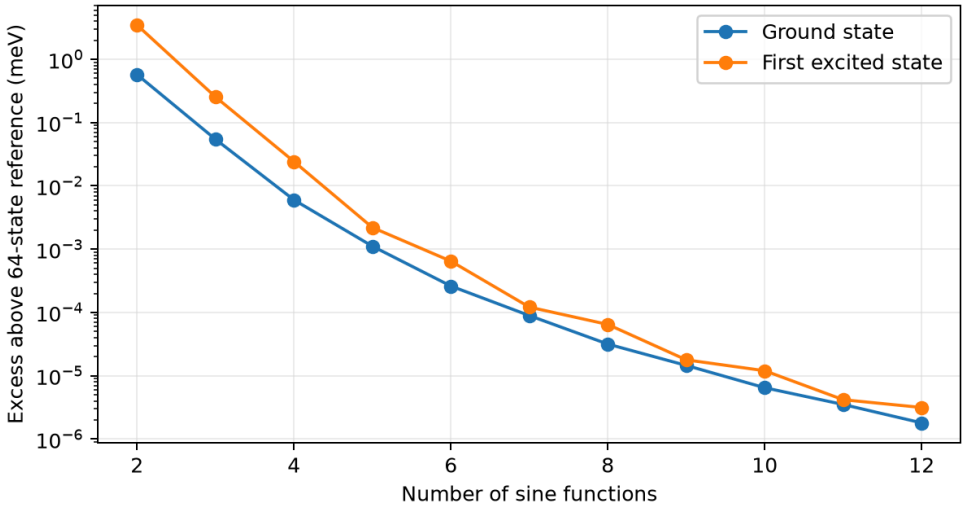

The reference energies are 21.495 and 38.379 meV. The excited state initially needs more basis functions to reach the same absolute accuracy.

Try it · Ritz basis

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Potential rise / untilted ground energy

12

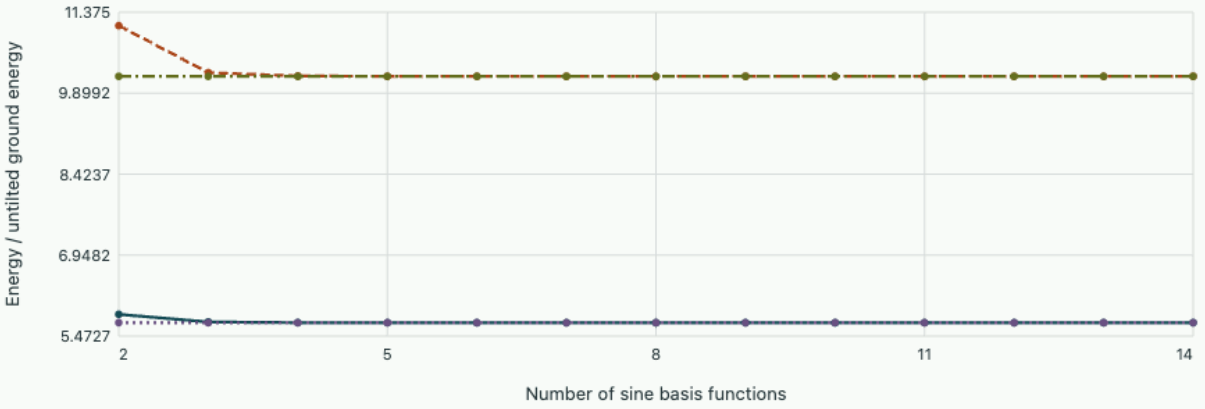

Nested sine spaces preserve the variational bounds. Browser reference has 24 functions; the supplied OneSimulation script uses 64. Untilting recovers n² exactly.

Download the Python example

Choose a length and compute two energies

As an example estimate a hydrogen-like electron bound to a fixed proton by $V = -\kappa/r$, with $\kappa = e^2/(4\pi\epsilon_0)$. The exact ground state has a cusp at the origin; deliberately choose the smooth, normalized Gaussian $\psi = (\pi a^2)^{-3/4}\exp[-r^2/(2a^2)]$. The parameter a permits expansion and contraction but cannot repair the cusp or the large-radius shape. It therefore tests how much energy a simple length-scale argument can recover.

For a decaying state integration by parts expresses kinetic energy as $\hbar^2\int|\nabla\psi|^2/(2m)$. Since $\nabla\psi = -\mathbf{r}\psi/a^2$ and $\langle r^2\rangle = 3a^2/2$, the result is $3\hbar^2/(4ma^2)$. For potential energy the radial measure $4\pi r^2 dr$ cancels one power of $1/r$, leaving $\langle 1/r\rangle = 2/(\sqrt{\pi}a)$. Thus

$$E(a) = \frac{3\hbar^2}{4ma^2} - \frac{2\kappa}{\sqrt{\pi}a}.$$

As a becomes small, kinetic localization costs energy faster than attraction can reduce it. At very large a the state becomes weakly bound. These limiting trends guarantee an intermediate minimum.

Optimize and assess the answer

With $a_0 = \hbar^2/(m\kappa)$, differentiation gives $a_* = (3\sqrt{\pi}/4)a_0 = 1.329a_0$ or 0.0703 nm. The minimized energy is $-4m\kappa^2/(3\pi\hbar^2)$, about -11.55 eV compared with the fixed-proton result -13.61 eV. It lies above the

exact energy by about 2.06 eV, as required. At the optimum the two terms satisfy

$$2\langle T\rangle = -\langle V\rangle.$$

That scaling relation follows from extremization over a length and can hold even with an imperfect wavefunction.

Energy is stationary near an optimum, so a good energy estimate alone cannot certify other observables. For example, sampling the contact probability $|\psi(0)|^2$ is sensitive to the cusp missing here. One can enlarge the basis using several Gaussian widths and diagonalize the Hamiltonian in their span. The old optimum remains an allowed choice, so the improved lowest eigenvalue can only fall or stay equal.

To see why the energy can be forgiving, resolve a normalized trial state as $\sqrt{1-\eta^2}|0\rangle + \eta|1\rangle$ using two exact eigenstates. Its energy error is $\eta^2(E_1 - E_0)$, quadratic in the admixture. An observable with a nonzero off-diagonal matrix element between these levels can already change at first order in η. Comparing several independent observables therefore tests more of the trial wavefunction than repeating the energy comparison.

Vary an effective nuclear charge for two electrons

Helium presents a variational problem in which the main missing physics can itself be identified. Fix a nucleus of charge $+Ze$ and use atomic units, where $\hbar = m_e = e^2/(4\pi\epsilon_0) = 1$, length is measured in Bohr radii and energy in hartrees. Ignoring relativity, use $H = \sum_{i=1}^{2}[-\nabla_i^2/2 - Z/r_i] + 1/r_{12}$. Try two identical normalized exponential orbitals $\phi = (\zeta^3/\pi)^{1/2}e^{-\zeta r}$ with a spin singlet so that the complete electronic state is antisymmetric. The trial parameter ζ lets both orbitals expand in response to repulsion.

For either electron the kinetic expectation is $\zeta^2/2$ and $\langle 1/r\rangle = \zeta$. The two one-body terms together give $\zeta^2 - 2Z\zeta$. To compute the interaction, spherical averaging the Coulomb kernel for the isotropic densities gives $1/\max(r_1, r_2)$, the shell-potential result. Split the radial domain at $r_1 = r_2$ and double the part with $r_1 > r_2$. With radial probability $P(r) = 4\zeta^3 r^2 e^{-2\zeta r}$ this gives

$$\left\langle \frac{1}{r_{12}} \right\rangle = 2\int_0^\infty \frac{P(r_1)}{r_1}\,dr_1 \int_0^{r_1} P(r_2)\,dr_2 = \frac{5\zeta}{8}.$$

For the last integration, set $x = 2\zeta r_1$. The inner integral is $1 - e^{-x}(1 + x + x^2/2)$, so the remaining integral involves only exponential moments. The repulsion is positive and linear in ζ: squeezing both clouds together increases typical electrostatic repulsion while the faster quadratic kinetic cost eventually prevents collapse.

Estimate the first ionization energy of helium

For $Z = 2$, minimize $E(\zeta) = \zeta^2 - 4\zeta + 5\zeta/8$. Setting its derivative to zero gives $\zeta_* = 2 - 5/16 = 1.6875$ and $E_{\mathrm{trial}} = -\zeta_*^2 = -2.84766$ hartrees, or -77.49 eV relative to a bare nucleus and two free electrons. An unscreened choice $\zeta = 2$ gives -2.750 hartrees. Allowing the cloud to expand lowers the trial energy by about 2.66 eV.

The one remaining electron in He^+ is a hydrogenic ion with exact fixed-nucleus energy $-Z^2/2 = -2.000$ hartrees in this model. The trial estimate for removing one electron is therefore 0.84766 hartree, or 23.1 eV.

Because the final one-electron energy is exact while the neutral trial energy is an upper bound, this difference is a lower bound on the model's ionization energy. The trial state allows average screening but leaves the electrons statistically independent in position. Multiplying the product orbital by a factor depending explicitly on r_{12} can favour configurations where the electrons avoid one another and lower the bound further.

The effective size can be compared directly with the unscreened hydrogenic model. For one exponential orbital $\langle r \rangle = 3a_0/(2\zeta)$, so the optimum mean radius for either electron is $0.889a_0$, compared with $0.750a_0$ at $\zeta = 2$. The increased radius and reduced energy come from the same variation, but they have different error guarantees. Variational minimization bounds energy, while no analogous inequality requires this mean radius to be either above or below its exact value. Comparing the inferred radius with scattering or density data therefore probes trial-wavefunction quality beyond the energy test.

Diagnose the missing electron correlation locally

The optimum product state gives a useful upper bound yet treats the positions of the two electrons as statistically independent. A short-distance diagnostic follows by asking what happens when the electrons approach each other while remaining away from the nucleus. Introduce relative coordinate $\mathbf{r} = \mathbf{r}_1 - \mathbf{r}_2$ and a centre coordinate. In atomic units the relative part of their kinetic Hamiltonian has the term $-\nabla_r^2$, since the reduced mass is one half. The Coulomb repulsion is $1/r$. For a singlet with nonzero amplitude at coincidence, expand its spherical average near $r = 0$ as $\psi(r) = \psi(0)(1 + cr + \cdots)$.

The radial Laplacian of the linear term is $2c\psi(0)/r$. If the Schrödinger equation is to admit a finite eigenvalue without an uncancelled $1/r$ term, the relative kinetic and repulsive singularities must cancel:

$$-2c\,\psi(0)/r + \psi(0)/r = 0,$$

$$c = \tfrac{1}{2}.$$

Thus the spherical average of the amplitude initially rises away from pair coincidence. The positive slope quantifies a correlation hole. A simple product of smooth one-electron orbitals has no such first-order slope at an ordinary coincident position. Optimizing its single length scale can reproduce average screening without satisfying this local condition.

For a separation of $0.10a_0$, the leading cusp term predicts a relative amplitude roughly five percent above its coincident value at fixed electronic centre, before quadratic corrections. Squared probability density therefore rises by about ten percent to first order. This compares local densities, not the count per radial shell, which also contains the shell area factor. A trial factor $1 + r_{12}/(2a_0)$ supplies the required initial slope while leaving a separate orbital scale available for long-distance screening. Its potential to improve the energy is motivated by a local equation, rather than simply by the fact that adding another variational parameter cannot worsen a minimum.

Choose a variational length in a linear trap

Take $H = -\frac{1}{2}d^2/dx^2 + |x|$ and a normalised trial amplitude proportional to $\exp[-x^2/(2s^2)]$. Integration of the gradient and potential terms gives $1/(4s^2)$ and $s/\sqrt{\pi}$. Use their sum as a ground-state upper bound.

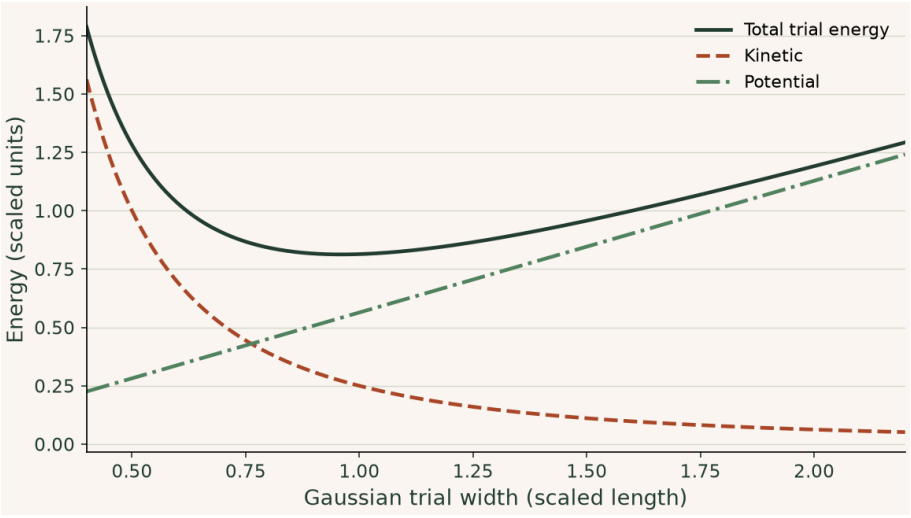

The minimum occurs at $s^3 = \sqrt{\pi}/2$, or $s \approx 0.961$, giving a variational energy near 0.813. At this optimum twice the kinetic contribution equals the potential contribution, a useful independent check on the scaling calculation.

Try it · Linear variational

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Trial width s **0.96**



Optimal width (√π/2)^(1/3) = 0.96054. Selected energy 0.81289. This is a variational upper bound for a linear |x| potential.

▶ **Read values from the plot**

Study the full derivation

Two electrons, exchange splitting and the cost of localisation · Quantum states, dynamics and measurement

Build a two-electron state that respects exchange, derive its charge and spin correlations, and then solve a two-site model exactly. The two examples reveal when a triplet is favoured and when an apparently weak hopping lowers the singlet instead.

Chapter 8 · Semiclassical quantization

Use a local wavelength to construct the quantization rule for a smooth well and recognize where that construction breaks down.

FROM FIRST PRINCIPLES

Build the semiclassical phase and the turning-point condition

In a classically allowed interval define $p(x) = \sqrt{2m[E - V(x)]} > 0$ and seek a one-dimensional stationary wave $\psi = A(x)e^{iS(x)/\hbar}$. Assume p varies little over a local wavelength. Substitute into the Schrödinger equation and separate terms by powers of $\hbar$. The leading real and next imaginary balances give

$$(S')^2 = p^2,$$
$$2A'S' + AS'' = 0 \quad \Longrightarrow \quad A^2 S' = \text{constant}.$$

Thus the phase increment is $\int p\,dx/\hbar$ and the travelling-wave amplitude varies as $p^{-1/2}$. This approximation fails where $p = 0$. Near each smooth simple turning point the potential is locally linear; the decaying exterior Airy solution matches an interior sine with phase shift $\pi/4$. The waves entering from the left and right must describe the same standing wave. Combining the two end phases gives

$$\int_{x_-}^{x_+} p(x)\,dx = \pi\hbar(n + \tfrac{1}{2}),$$
$$n = 0, 1, \ldots.$$

The offset comes from the two soft turning points. A hard wall imposes its own node condition and changes the quantization offset. The construction becomes more reliable when the state samples many local wavelengths inside the well.

DERIVATION AND APPLICATIONS

Read a quantum level spacing from a classical clock

The semiclassical quantization rule uses the area swept in phase space. Its change with energy also carries information about time. Consider a smooth one-dimensional binding potential with two simple turning points $x_\pm(E)$, mass m, and many accessible levels. Write the half-cycle action using the positive momentum branch $p(x, E) = \sqrt{2m[E - V(x)]}$:

$$I(E) = \int_{x_-}^{x_+} p(x, E)\,dx,$$
$$I(E_n) \simeq \pi\hbar(n + \tfrac{1}{2}).$$

The endpoints move with energy, but their boundary contributions when differentiating the integral contain $p(x_\pm) = 0$. Inside the well $\partial_E p = m/p$. Its divergence near a simple turning point behaves as an inverse square root and is integrable. Using $p/m = \dot{x}$ on the forward leg now gives

$$\frac{dI}{dE} = \int_{x_-}^{x_+} \frac{m}{p(x,E)}\,dx = \int_{x_-}^{x_+} \frac{dx}{\dot{x}} = \frac{T_{\mathrm{cl}}(E)}{2}.$$

Two neighbouring quantized actions differ by $\pi\hbar$. If the classical period varies little over a level spacing, linearizing the action about one level leads to

$$E_{n+1} - E_n \simeq \frac{2\pi\hbar}{T_{\mathrm{cl}}(E_n)} = \frac{h}{T_{\mathrm{cl}}(E_n)}.$$

Thus a densely spaced part of the spectrum is associated with a long classical excursion. For a harmonic well the period is constant and the result is exactly the known spacing; near a separatrix the rapidly varying period makes a one-step linearization less reliable. Only the period of a complete return to the starting phase-space point belongs in this relation.

WORKED EXAMPLE

Estimate a frequency from adjacent lines

Change the inputs and investigate this derivation.

Consider an approximately harmonic vibrational well with nearby levels separated by 2.50 meV. With $h = 4.136$ meV ps, the inferred classical period is 1.65 ps and the ordinary oscillation frequency is 0.604 THz. A transit from the left turning point to the right should take about 0.827 ps. Equating this one-way transit with the period would overestimate the energy spacing by two.

SECOND APPLICATION

Predict the scaling change for a quartic well

Change the inputs and investigate this derivation.

In $V(x) = \lambda x^4$ the turning point is $x_t = (E/\lambda)^{1/4}$. Substituting $x = x_t u$ into $T_{\mathrm{cl}} = 2\int_{-x_t}^{x_t} m\,dx/p$ leaves an energy-independent integral and gives

$$T_{\mathrm{cl}} \propto E^{-1/4}.$$

Therefore local level spacing grows as $E^{1/4}$. At sixteen times the energy a spacing of 0.180 meV should become about 0.360 meV, assuming both regions contain enough levels for semiclassical counting. The quadratic trap would show no comparable growth.

Try it · Power law spacing

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Potential exponent p in |x|^p

4

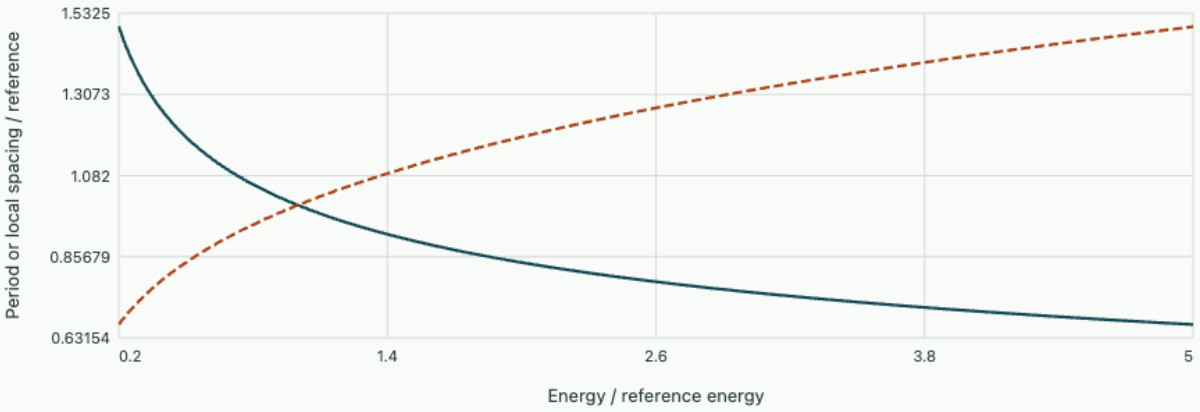

Classical period Semiclassical spacing

The action scales as E^(1/2+1/p). Differentiation gives T∝E^(1/p−1/2) and ΔE≈h/T. This local-spacing estimate improves for high quantum numbers.

A locally oscillating solution

Consider one-dimensional bound motion of mass m in a smoothly varying potential. In allowed regions the stationary equation is $\psi'' + p(x)^2\psi/\hbar^2 = 0$ with $p = \sqrt{2m(E-V)}$. If the potential changes little over a local wavelength, a small interval resembles a free-particle region. Accumulating its local phase produces the action integral $\int p\,dx/\hbar$. Yet a wave with changing local momentum must also change amplitude in order to conserve current.

Substitute $\psi = A(x)\exp[iS(x)/\hbar]$. The leading real terms give

$$S'^2 = p^2.$$

The next imaginary terms give $2A'S' + AS'' = 0$, or $(A^2S')' = 0$. Therefore each travelling contribution has amplitude proportional to $1/\sqrt{p}$. Equivalently a classical particle dwells longer where it is slow, and the smoothed semiclassical probability mirrors this dwell time. The neglected derivative terms are small away from turning points when $\hbar|p'| \ll p^2$.

Match through a turning point

At an ordinary turning point $V(x_t) = E$ with nonzero slope, the local wavelength diverges and the previous criterion fails. Linearizing V there gives the Airy equation. Its decaying forbidden-side solution continues to an oscillatory allowed-side standing wave with a phase offset of $\pi/4$. Matching decaying exterior tails on both sides of a well bounded by two such turning points imposes

$$\int_{x_-}^{x_+} p(x)\,dx = \pi\hbar(n + \tfrac{1}{2}),$$

$$n = 0, 1, 2, \ldots.$$

The offset encodes the two smooth reflections. An infinite hard wall instead imposes a node directly and has a different reflection phase. Applying this rule unchanged to a hard box would count its levels incorrectly. For slowly varying barriers the forbidden action also controls exponential attenuation, but the two-turning-point integral above determines the energies of a classically allowed well.

A trap with measurable spacings

For $V = m\omega^2 x^2/2$ let $x_t = \sqrt{2E/(m\omega^2)}$. Put $x = x_t \sin u$ in the action integral. The integrand becomes $(2E/\omega)\cos^2 u\,du$ between $-\pi/2$ and $\pi/2$, yielding $\pi E/\omega$. The rule gives $E_n = \hbar\omega(n + 1/2)$, exceptionally exact even for the oscillator’s lowest levels.

For a mechanical trap at frequency 100 kHz, $E_4 = 4.5hf = 2.98 \times 10^{-28}$ J, equivalent to 21.6 μK on division by k_B. The spacing is hf or 4.80 μK. These temperatures express energy units; no equilibrium ensemble was assumed. As a transfer test consider $V = gx^4$. Rescaling $x = (E/g)^{1/4} y$ makes the action proportional to $E^{3/4}$. Quantization gives $E_n \propto (n + 1/2)^{4/3}$, so successive spacings grow with excitation. At sufficiently high level number the action changes slowly over a wavelength, making the predicted trend more trustworthy than a low-level numerical prefactor.

An independent check on the energy spacing follows by differentiating the action with respect to energy. The integrand vanishes at the turning points, so their moving endpoints add no first-order term. Inside the integral $\partial p/\partial E = m/p = 1/v$. Consequently the derivative of the one-way action is the time for the classical particle to travel from one turning point to the other. Adjacent quantum levels at high excitation then have separation approximately h divided by the full classical period. Measuring a variation of level spacing with excitation is therefore also probing the dependence of the classical orbit period on energy.

Estimate tunnelling between two bound wells

A nearly symmetric double well can have pairs of quantum states concentrated mainly in classically separate regions. For one such pair take approximately orthogonal localized states $|L\rangle$, $|R\rangle$ with equal isolated energy E_0 and a small real tunnelling matrix element $-J$. The sign of either localized basis state can be chosen so $J > 0$. The two-level Hamiltonian has symmetric and antisymmetric eigenstates with energies $E_0 - J$ and $E_0 + J$, respectively. A local preparation $|L\rangle$ is a superposition of both and acquires a relative phase at rate $2J/\hbar$.

On evolution the probability to be in the other well is $P_R(t) = \sin^2(Jt/\hbar)$. The first complete transfer occurs at $t_{\mathrm{tr}} = \pi\hbar/(2J) = h/(2\Delta E)$ with $\Delta E = 2J$. In a high smooth barrier the localized wave tails overlap

weakly. A semiclassical estimate gives

$$J \simeq J_0 e^{-\mathcal{S}/\hbar},$$

with $\mathcal{S} = \int_{x_1}^{x_2} \sqrt{2m[V(x) - E_0]}\, dx$ between inner turning points. This is an amplitude action. The through-barrier transmission probability of a scattering experiment instead contains twice that exponent. Confusing the two quantities gives an incorrect predicted change in the oscillation time.

Predict how a gate changes a tunnelling period

Suppose a symmetric electron double well shows a spectroscopic splitting $\Delta E = 40.0$ μeV. Its predicted first full transfer time is

$$h/(2\Delta E) = 51.7\,\mathrm{ps}.$$

Use the measured splitting to eliminate the unknown semiclassical prefactor. A gate modification adds a nearly flat forbidden segment of width 0.300 nm where $V - E_0 = 0.150$ eV and otherwise changes the wells negligibly. Using $\hbar^2/(2m_e) = 0.0381$ eV nm² gives decay wavenumber $\kappa = \sqrt{0.150/0.0381} = 1.98$ nm⁻¹. The added dimensionless action is 0.595.

If the prefactor and individual well frequencies remain similar, the new splitting is $40.0 e^{-0.595} = 22.1$ μeV, and transfer time rises to 93.6 ps. Doubling the width change multiplies the splitting by the square of the original ratio, not by twice the original ratio. These estimates concern coherent dynamics before energy relaxation or phase noise destroys the relative phase.

A gate can also shift the two well bottoms differently. With detuning δ the splitting is $\sqrt{\delta^2 + 4J^2}$, while maximum transfer falls to $4J^2/(\delta^2 + 4J^2)$. A larger measured splitting can therefore accompany poorer population transfer if symmetry is lost. Comparing spectroscopy and time-resolved populations distinguishes a changed tunnel coupling from an unintended potential offset.

The small tunnel coupling also sets a time criterion for disorder: a fluctuating left–right offset comparable with the splitting can produce substantial changes in both transfer amplitude and oscillation frequency. A static distribution of offsets across repeated preparations can wash out ensemble oscillations even when each individual preparation evolves coherently. Repeating the experiment after compensating the measured detuning separates this ensemble averaging from intrinsic energy relaxation. The semiclassical action comparison is most reliable when the two localized levels remain nearly resonant throughout the modification and no higher intrawell level enters the same small energy window.

Independent measurements of the individual well minima supply a practical detuning calibration.

Study the full derivation

How a turning point supplies the missing semiclassical phase · Quantum states, dynamics and measurement

Far from a turning point the wavelength follows the classical momentum. At the turning point that approximation becomes singular. Derive the Airy connection, quantise a smooth well, distinguish a hard wall and estimate a tunnelling exponent with a defensible error check.

Problem sheet 4 · Energy bounds and semiclassical actions

Construct a bound from a normalized trial state, compare two reflection conditions and infer a potential's scaling from level spacings.

A polynomial trial state

In an infinite well $0 < x < L$, take $\psi = Nx(L - x)$. Determine N, compute its energy expectation and compare with the exact lowest level. Use $\int_0^L x^2(L-x)^2 dx = L^5/30$. What does the comparison test?

▼ **Show the variational calculation**

Normalization gives

$$N = \sqrt{30/L^5}.$$

Since the trial state vanishes at the walls, the kinetic expectation is $\hbar^2 \int |\psi'|^2 dx/(2m)$. The derivative is $N(L - 2x)$, whose squared integral is

$$N^2 L^3/3 = 10/L^2.$$

Thus the bound is $10\hbar^2/(2mL^2)$ and the exact ground energy $\pi^2\hbar^2/(2mL^2)$. The bound is $10/\pi^2 - 1 = 1.32\%$ high. Correct endpoints and normalization yield a valid test even though the polynomial has a different shape from a sine.

One wall and one smooth turning point

Consider a particle with an infinite wall at $x = 0$ and $V(x) = Fx$ for positive x. Evaluate the one-way action to the turning point. Using $\int_0^{x_t} p\,dx = \pi\hbar(n + 3/4)$ derive the energy law and the ratio E_1/E_0 in this approximation.

▼ **Show the mixed-boundary quantization**

The turning point is $x_t = E/F$. The integral is

$$\int_0^{E/F} \sqrt{2m(E - Fx)}\,dx = 2\sqrt{2m}E^{3/2}/(3F).$$

Hence $E_n = [3\pi\hbar F(n + 3/4)/(2\sqrt{2m})]^{2/3}$, giving $E_1/E_0 = (7/3)^{2/3} = 1.76$. The hard wall fixes a node while the smooth turning point contributes one Airy phase. The low-level ratio is approximate because semiclassical assumptions fail near the turning point and there are few oscillations in the allowed region.

Identify the stiffness of a potential

A symmetric smooth well behaves as $V = g|x|^s$ and high-level energies scale approximately as $n^{1.50}$. Derive the scaling for general s, identify the exponent and predict how energies at fixed high n change if g is multiplied by four.

▼ **Show the action rescaling**

Set $x = (E/g)^{1/s} y$. The action becomes a constant times $\sqrt{m}\, g^{-1/s} E^{1/2+1/s}$. Quantization gives $E_n \propto g^{2/(s+2)} (n + 1/2)^{2s/(s+2)}$ at fixed mass. Solving $2s/(s+2) = 1.50$ gives $s = 6$. A fourfold increase in g multiplies energies by $4^{1/4} = \sqrt{2}$. This inference uses the spatial region sampled by the measured high levels; the central part of an actual well may follow a different law.

Explore these relations

- For the polynomial trial state, compare the variational energy estimate with the exact ground level; the adjustable trial family in the plot is different from the polynomial in this sheet. Open the live calculation.
- For the potential inference, vary its power law and compare the semiclassical level-spacing trend. Open the live calculation.

Notebook practice · Can changing a Gaussian width repair a missing cusp?

Chapter 9 · Transitions under a timed drive

Derive the line shape of a weak finite-duration drive and distinguish Fourier width from a system's intrinsic lifetime.

FROM FIRST PRINCIPLES

How the time profile of a drive becomes a transition amplitude

Assume an initially occupied eigenstate $|i\rangle$ of a time-independent Hamiltonian H_0, perturbed by a weak time-dependent operator $V(t)$. Expand the state in its unperturbed energy eigenvectors and factor out their known phases. Inserting the expansion into the time-dependent Schrödinger equation cancels the unperturbed energy terms:

$$|\psi(t)\rangle = \sum_n c_n(t) e^{-iE_n t/\hbar}|n\rangle,$$

$$i\hbar \dot{c}_f = \sum_n V_{fn}(t) e^{i(E_f - E_n)t/\hbar} c_n.$$

At first order the right side uses the initial amplitudes $c_i = 1, c_{n\neq i} = 0$. Integrating from zero through pulse duration T gives

$$c_f^{(1)}(T) = -\frac{i}{\hbar}\int_0^T V_{fi}(t) e^{i\omega_{fi}t} dt,$$

$$\omega_{fi} = (E_f - E_i)/\hbar.$$

A square nearly resonant sinusoidal drive has a slowly varying integrand proportional to $e^{i\Delta t}$ plus a rapidly oscillating component. Keeping the slow term produces amplitude proportional to $T\,\mathrm{sinc}(\Delta T/2)$ and probability proportional to its square. The approximation requires small population transfer; a longer resonant drive eventually requires coupled state amplitudes.

DERIVATION AND APPLICATIONS

Why a strong resonant pulse stops increasing the transition probability

Perturbation theory predicts a probability growing as pulse duration squared on resonance. To follow the subsequent return of population, keep both amplitudes coupled. Consider an isolated pair of levels driven by a monochromatic field with a nearly resonant frequency. Use the rotating-wave approximation, justified when the detuning and coupling are both small compared with the carrier frequency. Let Δ be the angular detuning and Ω the resonant angular Rabi frequency. During a square pulse the rotating-frame Hamiltonian, up to an identity term, is

$$H_R = \tfrac{\hbar}{2}(\Omega\sigma_x + \Delta\sigma_z),$$

$$H_R^2 = \tfrac{\hbar^2}{4}(\Omega^2 + \Delta^2)I.$$

The square follows because the two Pauli matrices anticommute. In the exponential propagator the even powers of H_R therefore multiply the identity and the odd powers multiply H_R. Sum the two Taylor series with $\Omega_e = \sqrt{\Omega^2 + \Delta^2}$:

$$U(t) = I\cos(\Omega_e t/2) - i\,\frac{\Omega\sigma_x + \Delta\sigma_z}{\Omega_e}\sin(\Omega_e t/2).$$

Starting in the lower bare state, only σ_x contributes to the upper-state amplitude. Its squared magnitude is

$$P_e(t) = \frac{\Omega^2}{\Omega^2 + \Delta^2}\sin^2\left(\frac{\sqrt{\Omega^2 + \Delta^2}\,t}{2}\right).$$

On resonance a pulse area $\Omega t = \pi$ transfers the entire population; twice that duration returns it to the lower state. At short times the leading term is $\Omega^2 t^2/4$, recovering perturbation theory. Detuning both speeds up the oscillation and reduces its attainable maximum. These two signatures help distinguish mistuning from a change in drive strength.

WORKED EXAMPLE

Design a population inversion pulse

Change the inputs and investigate this derivation.

If $\Omega/(2\pi) = 2.00$ MHz, the resonant inversion time is $\pi/\Omega = 0.250$ microseconds. A 0.125 microsecond pulse gives $P_e = 0.500$. Holding the drive on for 0.500 microseconds instead gives zero in the ideal model. The source supplies and subsequently receives energy during a reversible oscillation.

SECOND APPLICATION

Infer the mistuning from an incomplete oscillation

Change the inputs and investigate this derivation.

Suppose the same drive produces maxima of 0.640 and full population oscillations repeat every 0.400 microseconds. The period gives

$$\Omega_e/(2\pi) = 2.50\,\text{MHz}.$$

The maximum gives $\Omega/\Omega_e = \sqrt{0.640} = 0.800$, recovering the calibrated 2.00 MHz drive. The detuning magnitude is $\sqrt{2.50^2 - 2.00^2} = 1.50$ MHz in ordinary-frequency units. A population trace alone leaves its sign undetermined.

How a time-dependent perturbation transfers population

For a time-independent Hamiltonian H_0 the energy components of a state accumulate phases without changing populations. Let a weak applied perturbation $V(t)$ act for a known interval. Prepare the unperturbed

eigenstate $|i\rangle$ at $t = 0$ and write $|\psi(t)\rangle = \sum_n c_n(t) e^{-iE_n t/\hbar}|n\rangle$. The interaction-picture amplitudes isolate changes beyond the free phases. Substitution into Schrödinger's equation and projection onto $\langle f|$ give

$$i\hbar \dot{c}_f = \sum_n V_{fn}(t) e^{i\omega_{fn} t} c_n,$$

$$\omega_{fn} = (E_f - E_n)/\hbar.$$

At first order retain $c_i = 1$ on the right-hand side. For $f \neq i$ the resulting amplitude is the Fourier component of the perturbation at the transition frequency,

$$c_f(T) = -(i/\hbar) \int_0^T V_{fi}(t) e^{i\omega_{fi} t}\, dt.$$

An energetically resonant drive still fails to populate a state if the coupling matrix element vanishes by symmetry. The operator describing the drive is as important as its frequency.

A rectangular pulse

Let $V_{fi}(t) = V_{fi}\cos(\omega t)$ for $0 < t < T$, with an upward transition so $\omega_{fi} > 0$. Splitting the cosine into exponentials gives frequencies $\omega_{fi} - \omega$ and $\omega_{fi} + \omega$. Near resonance, for a pulse lasting many carrier cycles, the rapidly oscillating sum-frequency term contributes little. With detuning $\Delta = \omega_{fi} - \omega$ the remaining integral gives

$$P_{i\to f} \simeq \frac{|V_{fi}|^2 T^2}{4\hbar^2} \left[\frac{\sin(\Delta T/2)}{\Delta T/2}\right]^2.$$

At exact resonance all contributions build with the same phase, so for a weak short pulse probability grows quadratically in time. The first zeros occur at $|\Delta| = 2\pi/T$, or ordinary-frequency detuning $1/T$. The line shape measures coherent addition over the finite observation time. It can be broad even when the isolated energy levels have negligible natural width.

Calibrate a pulse

Take a 2.00 ms rectangular pulse with $|V_{fi}|/\hbar = 200$ rad/s. The resonant transition probability is

$$(200 \times 0.002/2)^2 = 0.0400.$$

At 250 Hz detuning the sinc argument is $\pi/2$, leaving probability 0.0162. The first zeros lie 500 Hz on either side. Doubling pulse duration at the same coupling predicts four times the peak probability and half the first-zero spacing within the weak-drive approximation. A measurement can test both the area dependence and frequency scaling.

Once the predicted population is no longer small, depletion of the initial level removes the first-order approximation. For an isolated resonantly driven two-state system the rotating-wave solution is $\sin^2(|V_{fi}|T/(2\hbar))$, describing coherent population transfer and return. Lifetime decay or field fluctuations

introduce still other envelopes. Varying pulse duration at fixed drive strength distinguishes finite-time broadening from linewidths that persist as the pulse is lengthened.

The pulse envelope provides an additional independent control. Replacing the sudden rectangular turn-on and turn-off with smooth edges reduces spectral sidelobes because the Fourier transform of the smoother envelope falls faster away from resonance. Peak transfer at resonance, in the weak-drive regime, depends on the time integral of the drive envelope. Two pulses of different shapes can therefore have the same low-power resonant transition probability yet excite a nearby unwanted level with different probabilities. Specifying duration alone would be insufficient to reproduce such a measurement.

From a finite pulse to Fermi's golden rule

An isolated resonant final state initially collects amplitude proportional to time. A dense set of final states can instead produce a nearly constant total transition rate. To see the change, let a weak drive connect one initial state to levels around energy $E_\star = E_i + \hbar\omega$. Denote the matrix element of the resonant Fourier component by M. For a cosine drive it contains the factor of one half. First-order time integration over a square pulse gives a phase times $(Mt/\hbar)\,\mathrm{sinc}[(E-E_\star)t/(2\hbar)]$, where $\mathrm{sinc}(u) = \sin(u)/u$.

The different final states are orthogonal, so their probabilities add. If the density $\rho(E)$ counts accessible final states per unit energy, and ρ and M vary slowly over the line width,

$$P(t) \simeq \frac{|M|^2 t^2}{\hbar^2} \int \rho(E)\,\mathrm{sinc}^2\!\left(\frac{(E-E_\star)t}{2\hbar}\right) dE.$$

Set $u = (E-E_\star)t/(2\hbar)$. Away from a band edge one can extend the integral across the line and use $\int_{-\infty}^{\infty} \mathrm{sinc}^2 u\, du = \pi$. One power of time cancels against the narrowing energy interval, leaving $P(t) \simeq \Gamma t$ with $\Gamma = 2\pi|M|^2\rho(E_\star)/\hbar$. This intermediate-time result assumes many final levels within the line and little depletion of the initial state. For a small isolated system the long-time dynamics can resolve individual levels again.

As a numerical example, use $M = 0.500\,\mu\mathrm{eV} = 0.000500\,\mathrm{meV}$, $\rho = 1000\,\mathrm{meV}^{-1}$ and $\hbar = 0.6582\,\mathrm{meV\,ps}$. The rate is $0.00239\,\mathrm{ps}^{-1}$, or a transition probability of 0.0239 after $10.0\,\mathrm{ps}$. The first zero lies $2\pi\hbar/t = 0.414\,\mathrm{meV}$ from resonance, spanning about 414 level spacings on each side for a locally uniform density. The small transition probability and dense spectrum are distinct checks on the approximation.

Practise on sheet 5 · Compare two continua competing for one initial state. Further reading · Massachusetts Institute of Technology · 8.06 Quantum Physics III · Chapter 2, Time-dependent approximation methods, §§2.1–2.3, Incoherent light, Spontaneous emission and The photoelectric effect.

Separate phase accumulation from the drive pulses

A single spectroscopy pulse ties spectral resolution to its own duration. The Ramsey method separates two short coherent pulses by an interval of free evolution. Work with two isolated states $|g\rangle$, $|e\rangle$ in the rotating frame of a stable reference. Approximate each pulse as short compared with the free interval and strong

enough to rotate the Bloch vector by $\pi/2$ about the same transverse axis. Ignore relaxation and pulse detuning during the rotations. Starting in $|g\rangle$, the first pulse prepares $(|g\rangle - i|e\rangle)/\sqrt{2}$ for the chosen pulse phase.

Over free interval T the excited component gains relative phase $e^{-i\Delta T}$, with $\Delta = \omega_{eg} - \omega_{\mathrm{ref}}$. The second pulse mixes both paths to $|e\rangle$. Using the rotation matrix $R_x(\pi/2) = (I - i\sigma_x)/\sqrt{2}$ gives

$$c_e = -\frac{i}{2}(1 + e^{-i\Delta T}),$$

$$P_e = \frac{1 + \cos(\Delta T)}{2}.$$

A controllable phase on the second pulse shifts these fringes; on exact resonance two aligned pulses complete a π rotation and transfer all population. For a statistical mixture with the same half populations after pulse one, the two contributions to the final probability have no mutual phase and give one half regardless of delay. Thus the separated-pulse signal measures phase retained between pulses, even when no energy transfer takes place during most of the sequence.

Measure a small reference-frequency error

Take a free interval of 2.00 ms. Adjacent maxima occur at ordinary detunings separated by $1/T = 500$ Hz. The first minimum on either side occurs at 250 Hz. The separation of the two half-height points immediately surrounding resonance is

$$1/(2T) = 250\,\mathrm{Hz}.$$

These three figures refer to different linewidth or spacing conventions and should be identified before comparing experimental reports.

Set the second pulse phase so that the calibrated signal near zero detuning is $P_e = [1 + \sin(2\pi\delta\nu T)]/2$. A measured $P_e = 0.600$ gives $\delta\nu = \arcsin(0.200)/(2\pi T) = 16.0$ Hz on the branch nearest the reference. This setting is more sensitive to small signed frequency errors than the top of a fringe, where the first derivative vanishes. A first scan or a shorter free interval is needed to identify which repeated fringe is being tracked.

If the transverse coherence decays during free evolution, observed fringe amplitude is multiplied by the coherence factor, often approximated by e^{-T/T_2}. Extending the delay narrows fringe spacing but eventually reduces contrast. Measuring two phase settings in quadrature can separate an actual phase shift from a loss of visibility due to decoherence or pulse-area errors.

The frequency sensitivity per detected particle can also be estimated. At half population an ensemble of N independent measurements has standard error $1/(2\sqrt{N})$ in the estimated excited fraction. On the central quadrature slope $dP_e/d\delta\nu = \pi T$, giving frequency uncertainty approximately $1/(2\pi T\sqrt{N})$ for ideal contrast. With the 2.00 ms delay and ten thousand independent detections this is about 0.80 Hz for a single ensemble. Imperfect contrast reduces the slope, and an unknown phase offset adds a systematic error that collecting more particles alone cannot remove.

Separate detuning from shortened observation time

Start a two-level system in its ground state, apply a monochromatic drive of resonant Rabi frequency Ω, and neglect relaxation. In the rotating-wave approximation detuning δ changes the oscillation frequency to $\sqrt{\Omega^2 + \delta^2}$ and reduces its maximum transfer.

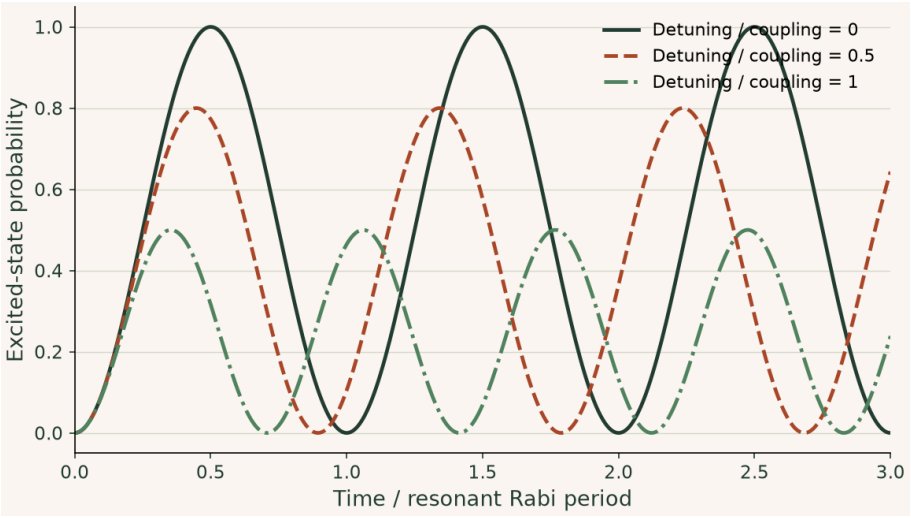

At detuning equal to the resonant Rabi frequency the largest excited population is one half and its oscillation is √2 times faster. Sampling only at the nominal resonant π-pulse time confounds both effects.

Try it · Rabi

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Detuning / resonant Rabi frequency 1

Pulse length in resonant π times 1

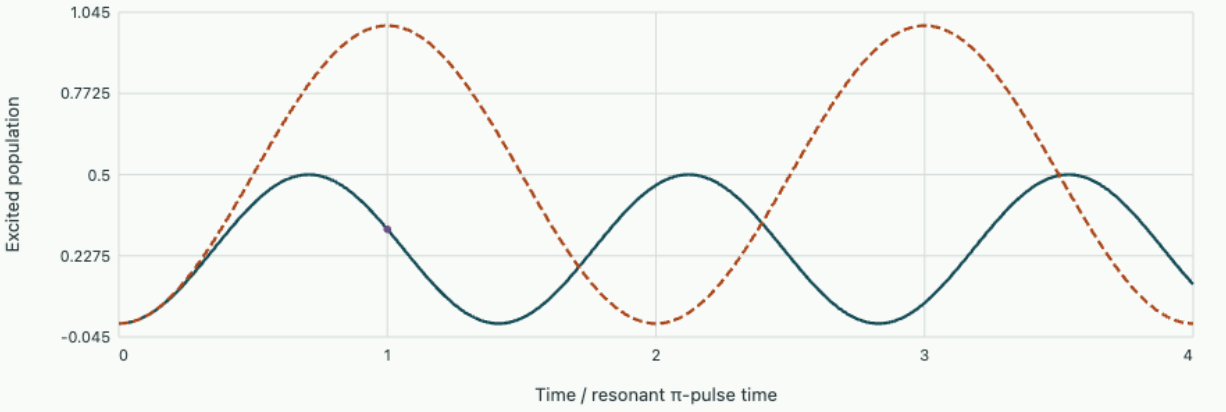

Maximum excited population 0.5; pulse-end population 0.31656. Rotating-wave, coherent two-level model.

▶ **Read values from the plot**

Study the full derivation

From a finite pulse to a continuum transition rate · Quantum states, dynamics and measurement

Switch a weak perturbation on for a measured duration. Its Fourier transform determines the amplitude of a discrete transition. Summing over a dense set of final states reveals where a time-independent transition rate comes from and where it fails.

Chapter 10 · Exchange and two-particle energies

Build allowed two-electron states and calculate how exchange symmetry changes the expectation of a repulsive interaction.

FROM FIRST PRINCIPLES

Construct the allowed two-electron spatial wave functions

Consider identical electrons prepared from two orthonormal spatial orbitals a, b. A product $|ab\rangle$ labels which coordinate slot contains which orbital. Exchanging the slots maps it to $|ba\rangle$; the same exchange squared is the identity, so its eigenvalues are ± 1. Form the normalized eigenvectors

$$\Phi_\pm = (|ab\rangle \pm |ba\rangle)/\sqrt{2},$$
$$P_{12}\Phi_\pm = \pm\Phi_\pm.$$

The fermionic exchange postulate applies to the full state. A spatially symmetric pair therefore combines with an antisymmetric spin state, and a spatially antisymmetric pair with symmetric spin. For a symmetric interaction W, expand its expectation in both components. The two diagonal terms coincide and the off-diagonal terms are complex conjugates, so

$$\langle \Phi_\pm | W | \Phi_\pm \rangle = J \pm K,$$
$$J = \langle ab | W | ab \rangle,$$
$$K = \mathrm{Re}\langle ab | W | ba \rangle.$$

The difference of $2K$ is an interference effect in the two-particle amplitude. It can affect the spin-dependent spectrum even when W is itself independent of spin. If the orbitals cease to be orthogonal, their two-particle normalization must be recomputed before applying this expectation value.

DERIVATION AND APPLICATIONS

Exchange a pair of particles prepared in overlapping orbitals

The familiar factor of $1/\sqrt{2}$ for a two-particle superposition assumes orthogonal one-particle states. Localized orbitals centred near each other can overlap. Let $a(x), b(x)$ be normalized orbitals and $S = \langle a | b \rangle$. Temporarily label the two coordinate slots 1 and 2 to evaluate inner products. The overlap of $|a(1)b(2)\rangle$ with its exchange is the product of two one-particle inner products:

$$\langle ab | ba \rangle = \langle a | b \rangle \langle b | a \rangle = |S|^2.$$

Symmetrize or antisymmetrize the coordinate wave function and normalize only after including this overlap. For $|S| < 1$ the two allowed candidates are

$$\Psi_\pm(x_1, x_2) = \frac{a(x_1)b(x_2) \pm b(x_1)a(x_2)}{\sqrt{2(1 \pm |S|^2)}}.$$

The appropriate spatial symmetry depends on the particles and their spin preparation. To connect the result to a measurable energy, consider a weak one-dimensional short-range interaction modelled by $W = g\delta(x_1 - x_2)$. First-order perturbation theory evaluates the probability density on the coincidence line. With $D = \int |a(x)|^2 |b(x)|^2 dx$ it gives

$$\langle W \rangle_+ = \frac{2gD}{1 + |S|^2},$$

$$\langle W \rangle_- = 0.$$

At equal coordinates the two exchanged terms have the same value, so the antisymmetric amplitude cancels exactly. This cancellation applies to an ideal contact interaction. A force with finite range can still sample nonzero amplitude on either side of the coincidence line. In the limit $|S| = 1$ the antisymmetric candidate itself disappears.

WORKED EXAMPLE

Estimate the interaction splitting in separate traps

Change the inputs and investigate this derivation.

Two localized orbitals have overlap magnitude 0.500 and measured density overlap $D = 0.120\,\mu\mathrm{m}^{-1}$. For a weak repulsion $g = 2.00\,\mathrm{meV}\,\mu\mathrm{m}$ the symmetric spatial state acquires a shift $2(2.00)(0.120)/1.25 = 0.384$ meV. The antisymmetric state has zero first-order contact shift. A distinguishable pair in the unsymmetrized product would give $gD = 0.240$ meV. All estimates assume orbital shapes can be held fixed as the interaction is turned on.

SECOND APPLICATION

Put both particles in the same orbital

Change the inputs and investigate this derivation.

For two bosons occupying the same normalized orbital, $S = 1$ and $\Psi_+ = a(x_1)a(x_2)$. A box orbital of uniform density over $L = 5.00$ micrometres has $D = 1/L = 0.200\,\mu\mathrm{m}^{-1}$. At the same coupling the shift is $g/L = 0.400$ meV. Using the orthogonal-orbital normalization here would double this energy. Two spin-polarized identical fermions cannot prepare this shared-orbital state.

Try it · Exchange contact

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Contact strength g (meV µm) **2**

Density overlap D (1/µm) **0.12**

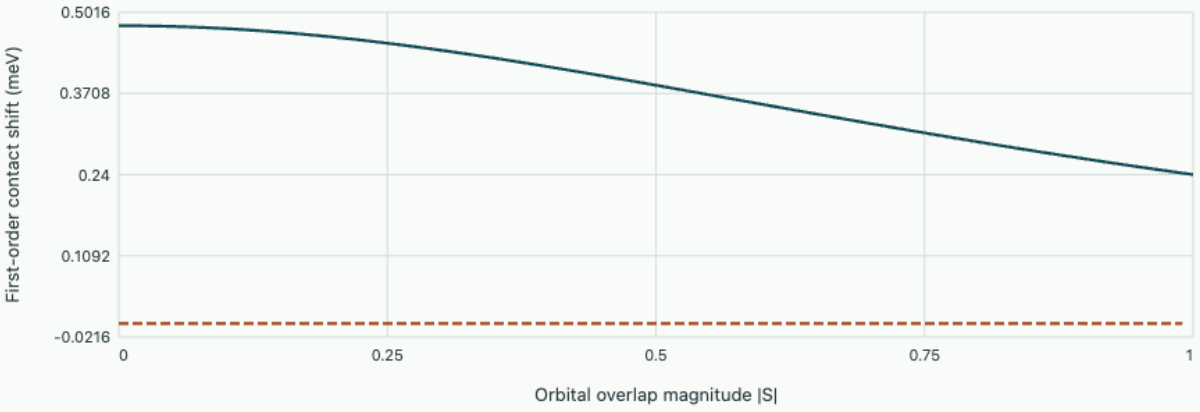

Symmetric spatial state Antisymmetric spatial state

The antisymmetric wavefunction vanishes at coincidence. Its state is undefined when the two orbitals become identical; the symmetric S=1 limit is well defined.

Exchange acts on the entire state

Two identical particles observed at different locations cannot be assigned enduring physical labels by a measurement of position alone. The operator P_{12} swaps all coordinates of the two particles, including spin. Since a second swap restores the original state, $P_{12}^2 = 1$ and exchange eigenvalues are ± 1. In three dimensions identical electrons are fermions, so their total state is antisymmetric under exchange. For a spin-independent Hamiltonian, spatial and spin factors can often be treated separately while enforcing the required product symmetry.

Take two normalized orthogonal spatial orbitals a and b. Both products $a(1)b(2)$ and $b(1)a(2)$ represent the same occupied orbitals. Their symmetric and antisymmetric combinations are

$$\Phi_{\pm} = \frac{a(1)b(2) \pm b(1)a(2)}{\sqrt{2}}.$$

The spin singlet is antisymmetric, requiring Φ_+ for the spatial part. All three spin triplets are symmetric, requiring Φ_-. Two electrons can share one spatial orbital in a singlet because they occupy different spin orbitals. In contrast the antisymmetric spatial combination vanishes if one sets $b = a$. The physical constraint applies to the full one-electron state.

Direct and exchange matrix elements

Suppose a, b are different eigenstates of a single-particle Hamiltonian with energies ϵ_a, ϵ_b. For their two-particle sum both spatial combinations have energy $\epsilon_a + \epsilon_b$. Now include a weak symmetric interaction $W(\mathbf{r}_1, \mathbf{r}_2)$. Expanding $\langle \Phi_\pm | W | \Phi_\pm \rangle$ gives two identical diagonal terms and two conjugate off-diagonal terms. Their combination is

$$\langle W \rangle_\pm = J \pm K,$$

$$J = \langle ab | W | ab \rangle,$$

$$K = \mathrm{Re} \langle ab | W | ba \rangle.$$

For the Coulomb interaction K is nonnegative. The triplet spatial state has suppressed coincidence probability, since $\Phi_-(\mathbf{r}, \mathbf{r}) = 0$. A repulsive interaction can therefore cost less in the triplet than in the corresponding singlet. This splitting is produced by interference of exchanged amplitudes in the interaction expectation, even when the interaction does not directly depend on spin.

A measured splitting

For two orbitals suppose the independent-particle energy sum is -8.00 eV, the direct integral is 1.20 eV and exchange is 0.070 eV. In the fixed-orbital estimate the singlet has energy -6.73 eV and the triplet -6.87 eV. The separation is $2K = 0.140$ eV, corresponding to frequency 3.39×10^{13} Hz. This energy difference need not produce a strong electric-dipole line between the two states. A spin-independent electric dipole preserves the orthogonality of singlet and triplet spin factors.

Orbitals localized far apart can still have appreciable direct electrostatic interaction because charge densities at different sites interact. The exchange integral also requires orbital overlap at a given coordinate and falls rapidly as overlap is removed. Holding sites apart while monitoring the singlet–triplet splitting can consequently distinguish the exchange contribution from the shared direct energy shift. Strong correlation or substantial orbital rearrangement would require a larger basis than the two fixed configurations used here.

Thermal population measurements require one more factor. With negligible magnetic field the three triplet spin projections share the same energy. At temperature well above the splitting, their combined population can approach three times that of the singlet. To extract an energy splitting from population ratios, include this multiplicity before taking a Boltzmann logarithm.

Interfere two identical bosons at a beam splitter

Exchange symmetry has measurable consequences even when there is no interaction energy. A balanced lossless optical beam splitter mixes two incident spatial modes a, b into outputs c, d. Choose phases so creation operators transform as $a^\dagger \to (c^\dagger + d^\dagger)/\sqrt{2}$ and $b^\dagger \to (c^\dagger - d^\dagger)/\sqrt{2}$. The relative sign enforces

orthogonality of the two output combinations, as required for a unitary transformation. Prepare one photon in each input with identical spectrum, time profile and polarization.

The input is $a^\dagger b^\dagger|0\rangle$. Applying the two transformations yields terms for both photons leaving c, both leaving d, or one leaving each output. The two one-per-output terms have opposite signs. Boson creation operators commute, so they cancel exactly:

$$|\psi_{\text{out}}\rangle = \frac{(c^\dagger)^2 - (d^\dagger)^2}{2}|0\rangle = \frac{|2,0\rangle - |0,2\rangle}{\sqrt{2}}.$$

The alternatives that cancel are both photons transmitted and both photons reflected into the one-per-output event. Photon number is conserved, but ideally the coincidence probability vanishes. If the photons have orthogonal polarization or arrive at distinguishable times, the corresponding internal states retain information about the alternatives. Their one-per-output probabilities then add to one half. This is a two-particle interference measurement, distinct from ordinary interference fringes in a single detector intensity.

Predict a delay scan for Gaussian wave packets

Let normalized temporal field amplitudes be $f(t)$ and $g(t)$ and let their overlap be $s = \int f^*(t)g(t)\,dt$. Retaining these internal labels while adding the two exchange amplitudes gives coincidence probability $P_{\text{coin}} = (1 - |s|^2)/2$ for an ideal balanced splitter with one photon in each input. Choose otherwise identical transform-limited Gaussian pulses whose individual arrival-time intensity distributions have standard deviation σ_t. Shifting one by delay τ makes $|s|^2 = \exp[-\tau^2/(4\sigma_t^2)]$, obtained by completing the square in the overlap integral.

Take $\sigma_t = 100$ fs. At $\tau = 200$ fs the expected coincidence probability is $(1 - e^{-1})/2 = 0.316$, between zero for perfect overlap and 0.500 for widely separated pulses. For 10,000 incident pairs and ideal detection one expects about 3160 coincidences. At a delay $2\sigma_t\sqrt{\ln 2} = 167$ fs the coincidence probability has recovered halfway from its minimum to the distant baseline. The full width between those two half-depth delays is 333 fs. Dark counts, multiple-pair generation or unequal beam-splitter transmission can raise the measured minimum even at zero delay. Recording the far-delay baseline separates overall detector efficiency from the normalized depth of the dip.

It is important here that indistinguishability includes every degree of freedom capable of preserving information about the two input alternatives. Two photons with the same central wavelength and similar arrival times may still have orthogonal polarizations or distinguishable frequency profiles. Including polarization overlap multiplies the temporal overlap by the appropriate inner product for separable polarization and time states. Rotating one linear polarization continuously therefore provides an independent control of coincidence depth, even with the delay held at zero. It can test whether residual coincidences come from distinguishability or from an unequal optical splitting ratio.

Study the full derivation

Two electrons, exchange splitting and the cost of localisation · Quantum states, dynamics and measurement

Build a two-electron state that respects exchange, derive its charge and spin correlations, and then solve a two-site model exactly. The two examples reveal when a triplet is favoured and when an apparently weak

hopping lowers the singlet instead.

Problem sheet 5 · Driven transitions and exchange symmetry

Interpret pulse spectra, apply a parity rule, and distinguish orbital occupation from complete fermion-state occupation.

Read the pulse duration

A rectangular weak pulse gives first spectral zeros 2.00 kHz apart across resonance. The on-resonance transfer probability is 0.0100. Find the pulse duration, coupling $|V_{fi}|/\hbar$ and probability halfway from the peak to the first zero.

▼ **Show the spectral calibration**

The first zeros are at $\pm 1/T$, giving $2/T = 2000$ Hz or $T = 1.00$ ms. Resonance gives

$$|V_{fi}|/\hbar = 2\sqrt{P}/T = 200\,\mathrm{rad/s}.$$

Halfway to a zero the argument is $\pi/2$, giving $0.0100(4/\pi^2) = 0.00405$. These numbers assume a stable transition, a nearly resonant carrier and sufficiently many cycles for the counter-rotating term to average away.

A forbidden line and an allowed one

A one-dimensional well is symmetric about zero. An initial even ground orbital is driven by $V(t) = Fx\cos\omega t$. Determine which parity of final state can be reached to first order. Compare to driving with $Kx^2\cos\omega t$ and explain why increasing pulse length fails to enable the forbidden first-order integral.

▼ **Show the matrix-element symmetry**

For an even final state the product $\psi_f^* x \psi_i$ is odd and its integral over symmetric limits vanishes. Only odd final parity can have a nonzero dipole matrix element. With x^2 the operator is even, so even final parity is allowed and odd parity excluded. Pulse duration changes the frequency integral multiplying the spatial matrix element. It cannot remove a spatial zero. A broken potential symmetry, a different driving operator or a higher-order transition can change an observed selection rule.

Extract an exchange integral

Two electrons occupy distinct orbitals. Measurements relative to the same energy reference give singlet energy 1.36 eV and triplet energy 1.24 eV. In the fixed-orbital model determine K and the common nonexchange contribution. Would three electrons fit into the same spatial orbital in this nonrelativistic model?

▼ **Show the state counting**

The energies are $E_{\mathrm{common}} \pm K$, with the plus for the singlet, so $K = (1.36 - 1.24)/2 = 0.060$ eV and $E_{\mathrm{common}} = 1.30$ eV. The shared energy includes the single-particle sum and direct interaction and cannot be

separated by this splitting alone. One spatial orbital affords two orthogonal electron spin states, up and down. Three occupied spin orbitals drawn from this two-dimensional space give a vanishing antisymmetric state, so a third electron must use a different spatial orbital.

Compare two continua competing for one initial state

Suppose two orthogonal channels have $(M, \rho) = (1.00\,\mu\mathrm{eV}, 500\,\mathrm{meV}^{-1})$ and $(0.500\,\mu\mathrm{eV}, 1500\,\mathrm{meV}^{-1})$. Find the total rate and the fraction of transitions into each channel. At $5.00\,\mathrm{ps}$, estimate how often the initial state has made either transition.

▼ **Add rates and normalize branching fractions**

In meV units the products $|M|^2\rho$ are $0.000500\,\mathrm{meV}$ and $0.000375\,\mathrm{meV}$. Multiplication by $2\pi/\hbar$ gives 0.00477 and $0.00358\,\mathrm{ps}^{-1}$, hence $\Gamma_{\mathrm{total}} = 0.00835\,\mathrm{ps}^{-1}$. The branching fractions are $4/7$ and $3/7$. Multiplying the total rate by $5.00\,\mathrm{ps}$ gives a probability of 0.0418 to leading order. Larger state density partly offsets the smaller amplitude in the second channel.

Revisit chapter 9 · From a finite pulse to Fermi’s golden rule.

Explore these relations

- For a rectangular pulse, compare duration and detuning in the coherent two-level solution; the weak-transfer limit gives the perturbative spectral envelope. Open the live calculation.
- For exchange symmetry, compare the spatial contact probability of symmetric and antisymmetric states. Open the live calculation.

Chapter 11 · Density matrices and loss of phase

Distinguish a coherent superposition from a statistical mixture and reconstruct a dephasing rate from measurements in two bases.

FROM FIRST PRINCIPLES

Turn a preparation ensemble into a density matrix

Suppose a preparation delivers normalized state $|\psi_j\rangle$ with classical probability w_j, where the weights sum to one. The Born probabilities for a measurement outcome with projector Π must be averaged over the unrecorded preparation. Use the identity $\langle\psi_j|\Pi|\psi_j\rangle = \mathrm{Tr}(|\psi_j\rangle\langle\psi_j|\Pi)$ and linearity of the trace:

$$P(\Pi) = \sum_j w_j \langle\psi_j|\Pi|\psi_j\rangle = \mathrm{Tr}(\rho\Pi),$$

$$\rho = \sum_j w_j |\psi_j\rangle\langle\psi_j|.$$

In an energy basis the diagonal entries determine energy populations. To measure coherence, project onto $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$. Multiplication and unit trace give

$$P(+) = \tfrac{1}{2}(\rho_{00} + \rho_{11} + \rho_{01} + \rho_{10}) = \tfrac{1}{2} + \mathrm{Re}\,\rho_{01}.$$

Thus equal energy populations can correspond either to perfectly coherent superposition or to an incoherent mixture. For unitary individual evolutions $\rho(t) = U\rho(0)U^\dagger$; random relative phases across the ensemble can suppress the average off-diagonal entry while preserving both populations. Tomography in a further phase-shifted basis accesses the imaginary component.

Background · Density operator

$$\rho = \sum_j p_j |\psi_j\rangle\langle\psi_j|$$

A physical state is positive and has trace one. Different statistical preparations can give identical density operators and measurement predictions.

DERIVATION AND APPLICATIONS

Separate reversible phase dispersion from irreversible decay with an echo

An ensemble of spins can lose average phase coherence even when every individual spin evolves unitarily. Start each spin in an equal-amplitude superposition. In a frame rotating at the reference frequency, assume each spin has a time-independent angular detuning δ drawn from a Gaussian with zero mean and standard

deviation σ. Ignore energy relaxation during the measurement. Under free evolution, the two components accumulate a phase difference δt. The normalized transverse coherence is the characteristic function of the detuning distribution:

$$C(t) = \int_{-\infty}^{\infty} \frac{d\delta}{\sqrt{2\pi}\sigma} \exp[-\delta^2/(2\sigma^2)] e^{-i\delta t}.$$

Completing the square in the Gaussian exponent or Fourier transforming the normalized Gaussian gives

$$C_{\mathrm{free}}(t) = e^{-\sigma^2 t^2/2},$$
$$P_{+x}(t) = \tfrac{1}{2}[1 + \mathrm{Re}\, C_{\mathrm{free}}(t)].$$

Now apply an ideal instantaneous π rotation about x at time τ. It exchanges the two levels, reversing the sign of subsequent relative phase accumulation in the toggling frame. An individual spin accumulates $\delta\tau - \delta(t-\tau)$ before readout at $t \geq \tau$. Gaussian averaging yields

$$|C_{\mathrm{echo}}(t)| = e^{-\sigma^2(2\tau - t)^2/2}.$$

At $t = 2\tau$ the static phase dispersion refocuses for every detuning. A separate independent exponential homogeneous decay factor e^{-t/T_2} would survive in the echo amplitude. Detunings changing appreciably between the two halves also spoil the cancellation. The echo therefore measures which part of an apparent decay remembers its earlier phase.

WORKED EXAMPLE

Predict recovery in a broadened sample

Change the inputs and investigate this derivation.

Take $\sigma/(2\pi) = 1.00$ kHz. At $t = 0.500$ ms the unrefocused coherence is $e^{-\pi^2/2} = 0.00719$, so

$$P_{+x} = 0.5036.$$

With a π pulse at 0.250 ms, the same static ensemble instead has ideal echo magnitude one at 0.500 ms. The two experiments distinguish rapid loss of ensemble contrast from loss of the individual phase information.

SECOND APPLICATION

Infer a homogeneous decay and a timing tolerance

Change the inputs and investigate this derivation.

Suppose echoes at 2.00 and 4.00 ms have normalized magnitudes 0.800 and 0.640. Their ratio is 0.800, giving $T_2 = -2.00\,\mathrm{ms}/\ln 0.800 = 8.96$ ms. If the readout is displaced by 0.100 ms from the true echo, static broadening with the previous σ reduces its contrast by another factor $e^{-(2\pi\,0.100)^2/2} = 0.821$. Timing and homogeneous decay enter separately in this inference.

Try it · Static echo

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Rms cyclic detuning (kHz)	1
Refocusing pulse time τ (ms)	0.25
Homogeneous T2 (ms)	10

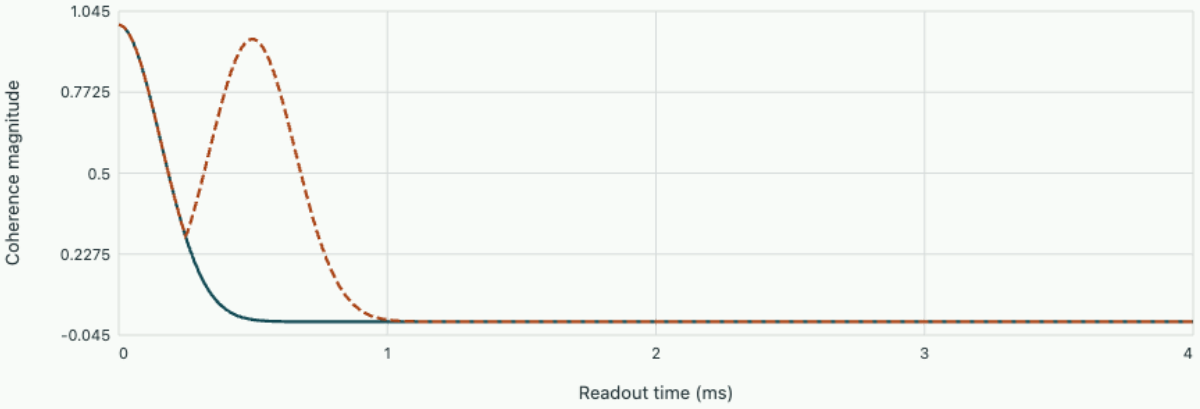

Echo centre at t=2τ=0.5 ms; its height exp(−2τ/T2)=0.95123. Static detuning reverses; homogeneous decay remains.

Two preparations with identical energy populations

Prepare a two-level system with states $|0\rangle$ and $|1\rangle$. One source always prepares $(|0\rangle + |1\rangle)/\sqrt{2}$. Another prepares $|0\rangle$ or $|1\rangle$ at random with equal probability. Measuring in the displayed basis gives identical frequencies for the two sources. After projection onto $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, however, the first source always gives plus and the second gives plus on half the trials. An adequate description must store more than the populations seen in one basis.

Define $\rho = \sum_j p_j |\psi_j\rangle\langle\psi_j|$ for a preparation ensemble. Probabilities follow from $P(Q) = \mathrm{Tr}(\rho Q)$ for a projector Q. The trace is one, Hermiticity makes expectation values real and positivity ensures probabilities are nonnegative. The first preparation has all matrix entries equal to $1/2$. The second has diagonal entries $1/2$ and off-diagonal entries zero. Although a density matrix can be decomposed into ensembles in multiple ways, every measurement prediction depends only on the matrix itself.

Coherence can be lost to an unobserved system

Consider a pure joint state $(|0\rangle|e_0\rangle + |1\rangle|e_1\rangle)/\sqrt{2}$ of system and environment. Average over any complete environment basis, equivalently take the partial trace. The system retains equal populations, and its off-diagonal element is $\rho_{01} = \langle e_1|e_0\rangle/2$. If the environment states become orthogonal, they carry enough information to distinguish the two alternatives and system-only interference disappears. The combined system can still be pure.

A simple dephasing model multiplies off-diagonal elements by e^{-t/T_2}, leaving populations fixed. In a rotating frame that removes the known level splitting, initial plus preparation then gives

$$P_+(t) = \frac{1 + e^{-t/T_2}}{2},$$

$$\operatorname{Tr}\rho^2 = \frac{1 + e^{-2t/T_2}}{2}.$$

The first expression follows by projecting onto plus; the second is the sum of diagonal entries of ρ^2. Purity ranges from one for a pure state to one half for the completely mixed two-level state. Population relaxation would change the diagonal entries too. Measuring both the energy basis and the interference basis helps decide whether the simpler model is adequate.

An interference record

Following a calibrated plus preparation, a measurement in the rotating frame gives $P_+ = 0.700$ after 30.0 µs. The contrast is $2P_+ - 1 = 0.400$, giving $T_2 = -30.0/\ln(0.400) = 32.7$ µs. The inferred purity is

$$(1 + 0.400^2)/2 = 0.580.$$

At 60.0 µs an exponential model gives contrast 0.160 and $P_+ = 0.580$. A different next value can signal nonexponential noise or a misestimated coherent phase.

A single interference quadrature cannot distinguish reduced visibility from uncompensated precession. To test the rotating frame assumption, vary the final projection phase and fit the full fringe. Its phase reports the coherent rotation; its amplitude reports visibility. Combining this with an energy-basis measurement determines the full Bloch vector and makes the purity estimate meaningful even when populations drift.

Purity estimates from finite counts can also be misleading when each fitted probability is treated independently. The reconstructed Bloch vector must have length at most one for a physical density operator. Longer vectors indicate sampling fluctuations or inconsistent calibrations. Increasing repetitions narrows counting uncertainty, while rotating a known reference state helps diagnose systematic errors among measurement axes.

Energy relaxation places a bound on phase memory

Pure dephasing leaves populations fixed. A two-level system coupled to a cold environment can also lose energy by relaxing from excited state $|1\rangle$ to ground state $|0\rangle$. Suppose the probability for an initially excited system to survive time t is $s = e^{-t/T_1}$. A simple system–environment evolution sends $|1, E_0\rangle$ to

$\sqrt{s}|1, E_0\rangle + \sqrt{1-s}|0, E_1\rangle$ and leaves $|0, E_0\rangle$ unchanged. The two environment states represent absence or presence of a released excitation and are orthogonal.

For an initial superposition only its unrelaxed excited amplitude shares the original environment state with the ground amplitude. Partial tracing consequently multiplies the off-diagonal system element by $\sqrt{s}$, while the excited-state probability is multiplied by s. Additional independent Markovian phase noise can multiply the off-diagonal element by e^{-t/T_ϕ}. In a rotating frame the resulting evolution is

$$\rho_{11}(t) = e^{-t/T_1}\rho_{11}(0),$$
$$\rho_{01}(t) = e^{-t/T_2}\rho_{01}(0),$$
$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi}.$$

Thus even a perfectly stable transition frequency cannot make coherence last indefinitely when population relaxes. In this two-level model $T_2 \le 2T_1$, with equality if there is no additional pure dephasing. At finite bath temperature upward jumps also occur and populations approach a nonzero equilibrium excitation. A multilevel environment or nonexponential noise can require a different fitting model.

Separate two decay mechanisms using independent preparations

First prepare the excited state and measure population decay; suppose it gives $T_1 = 40.0$ µs. Prepare $|+\rangle$ independently and fit fringe contrast to obtain $T_2 = 50.0$ µs. The pure-dephasing rate is $1/50.0 - 1/80.0 = 0.00750$ µs⁻¹, corresponding to $T_\phi = 133$ µs. At $t = 40.0$ µs after plus preparation, the excited population is $e^{-1}/2 = 0.184$, and the real coherence in the calibrated rotating frame is

$$e^{-0.8}/2 = 0.225.$$

Then

$$P_+ = 1/2 + \mathrm{Re}\rho_{01} = 0.725.$$

The reconstructed purity is $\rho_{00}^2 + \rho_{11}^2 + 2|\rho_{01}|^2 = 0.801$ to rounding. As time grows, coherence vanishes but the cold-bath model returns the system to the pure ground state, so purity ultimately rises back to one. This differs qualitatively from pure dephasing of an initial plus state, whose limiting purity is one half. An independently measured $T_2 > 2T_1$ would call for checking the data and assumptions about temperature, state assignment, fit functions and time-dependent noise before assigning an unphysical negative pure-dephasing rate.

The two independent preparations can be checked against a single density matrix at an additional time. Starting in plus, a general measurement along a transverse direction gives a cosine fringe whose offset is one half regardless of the energy populations, while its amplitude measures the transverse coherence. A separate energy-basis measurement supplies the population imbalance. Together they determine the full Bloch-vector length, which should never exceed one. If independently fitted transverse and longitudinal decay curves imply a greater length at some time, the combined fitted model cannot represent a physical quantum state there even if each measured curve looks plausible in isolation.

Checking positivity across the fitted interval also constrains long extrapolations from records that cover only a small fraction of a relaxation time.

Study the full derivation

How an unobserved environment changes a quantum state · Quantum states, dynamics and measurement

Model a qubit and a two-state environment by a joint unitary. Derive channels, the continuous-time master equation and the relation between longitudinal and transverse relaxation. A recorded event and an unread detector produce different conditional predictions.

Chapter 12 · Confinement and spectral geometry

Map a two-dimensional boundary-value problem to quantum energies and determine which information an energy spectrum can reveal about a well.

FROM FIRST PRINCIPLES

Derive a rectangular billiard spectrum from its boundary conditions

Take a free particle of mass m confined by infinite potential walls on the rectangle $0 < x < a$, $0 < y < b$. The potential is zero inside and the wave function vanishes at each edge. Use $\psi = X(x)Y(y)$ in the stationary Schrödinger equation and divide by the product. A function of x plus a function of y equals a constant everywhere, so the contributions must separately be constant:

$$-X''/X = k_x^2,$$
$$-Y''/Y = k_y^2,$$
$$E = \hbar^2(k_x^2 + k_y^2)/(2m).$$

The boundary at zero in each coordinate selects sine solutions. At the opposite walls those sines must vanish again, allowing only an integer number of half-wavelengths:

$$k_x = n\pi/a,$$
$$k_y = j\pi/b,$$
$$E_{nj} = \frac{\hbar^2\pi^2}{2m}\left(\frac{n^2}{a^2} + \frac{j^2}{b^2}\right), \quad n, j \geq 1.$$

A zero index would give the identically zero function here. Scaling both lengths by the same factor s reduces every energy by s^{-2} while preserving degeneracy patterns. For more complicated domains, direct separation is generally unavailable, but the same Dirichlet eigenvalue problem determines the allowed squared wave numbers.

DERIVATION AND APPLICATIONS

How do the walls modify the leading state count?

The area term in a quantum billiard counts allowed wave vectors as though they were continuous. The finite perimeter removes a systematic number of states from this estimate. Take a rectangular two-dimensional domain of sides a, b, an infinitely high wall on each side, and one spin species. The separated standing waves vanish at both ends; consequently their positive integer labels start at one. Let K be a cutoff in wave number. The exact count is a lattice-point problem:

$$N(K) = \#\{(n, m) : n, m \geq 1,\ (\pi n/a)^2 + (\pi m/b)^2 \leq K^2\}.$$

For clarity mirror the quarter ellipse into all four quadrants in the integer (n, m) plane. Every interior positive pair then contributes four points. A point on either axis belongs to none of the physical Dirichlet modes, and

the origin also has to be removed. Denoting the full-ellipse count by N_{full} gives the exact identity

$$4N = N_{\mathrm{full}} - 2\lfloor aK/\pi \rfloor - 2\lfloor bK/\pi \rfloor - 1.$$

A smooth local average at large K replaces the full-ellipse count by area $\pi(aK/\pi)(bK/\pi)$ to leading order. Retain the leading subtractions from the axes, whose lengths grow linearly with K. With physical area $A = ab$ and perimeter $P = 2(a + b)$ this yields

$$\overline{N}(K) \simeq \frac{AK^2}{4\pi} - \frac{PK}{4\pi}.$$

The bar signifies a count smoothed over nearby thresholds. Fluctuations in exact lattice-point counts and constant corner contributions remain. This is a cumulative estimate for many states; it cannot resolve individual degeneracies. For a quadratic dispersion relate wave number to energy by $E = \hbar^2 K^2/(2m^*)$ at effective mass m^*.

WORKED EXAMPLE

Compare the area estimate with exact enumeration

Change the inputs and investigate this derivation.

Take a square of side L with $KL/\pi = 5$. Positive pairs satisfy $n^2 + m^2 \le 25$. For $n = 1, 2, 3, 4$ the allowed counts are $4, 4, 4, 3$, giving 15 in all. The area estimate alone gives $25\pi/4 = 19.63$, while the perimeter subtracts 5.00, leaving 14.63. Agreement near this cutoff illustrates the systematic correction; it is not an exact equality.

SECOND APPLICATION

Distinguish two billiards with the same area

Change the inputs and investigate this derivation.

At fixed area 100 nm², compare a 10×10 nm square and a 20×5 nm rectangle. Their perimeters are 40 and 50 nm. At $K = 5.00$ nm⁻¹ both leading area counts are 198.9. The wall terms reduce them to about 183.0 and 179.0 states. Thus the longer perimeter reduces the smooth count by an estimated 3.98 states at the same cutoff. Spin degeneracy would multiply both predictions and their difference by two.

Try it · Weyl

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Rectangle aspect ratio at fixed area

1

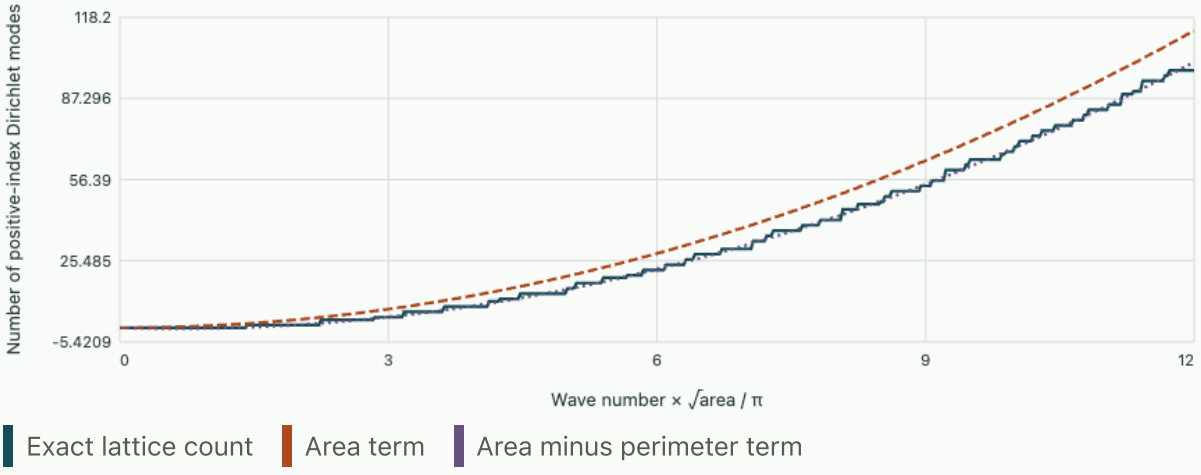

At KL/π=5 for a square: 15 modes. The perimeter correction is an asymptotic smoothed count, not an exact staircase at low frequency.

Start from this chapter. Two shapes can have the same Dirichlet spectrum. Compare their numerical ground energies at fixed scale, inspect their different mode shapes, then double the length scale without changing the shape.

OneSimulation laboratory · Turn spatial eigenvalues into confinement energies

A domain becomes part of the Hamiltonian

Constrain a nonrelativistic spinless particle of mass m to a flat region Ω with perfectly impenetrable walls. There is no potential inside. The wavefunction vanishes on the boundary, and square-integrability fixes its overall norm. The stationary equation becomes $-\nabla^2\psi = k^2\psi$ with $E = \hbar^2k^2/(2m)$. Thus the Laplacian, together with its domain and boundary condition, determines the spectrum. An acoustic pressure mode at a rigid reflecting wall generally has a different boundary condition and cannot be transferred unchanged.

For a rectangle with sides a, b, write $\psi = X(x)Y(y)$. Separation and the four zero-value conditions give $X \propto \sin(n_x\pi x/a)$, $Y \propto \sin(n_y\pi y/b)$ with positive integers. The energies are proportional to

$n_x^2/a^2 + n_y^2/b^2$. If $a = b$, exchanged index pairs give a degeneracy when unequal. Mildly stretching one side splits such a pair because nodes along the stretched direction cost less gradient energy.

Scaling and counting states

Under uniform dilation $\mathbf{r} = s\mathbf{r}'$ the gradient becomes

$$\nabla_{\mathbf{r}} = s^{-1}\nabla_{\mathbf{r}'}.$$

All eigenvalues of $-\nabla^2$ and all energies fall by s^2 in the denominator while nodal patterns simply stretch with the boundary. This result holds for any shape and any level of the empty hard-wall well. A potential with a fixed external length scale would add another term whose scaling requires separate analysis.

For highly excited modes, count the allowed points inside a wavenumber disk. In a large rectangle the spacing is π/a by π/b, and only the positive quadrant contributes. Dividing its area by one state’s area gives

$$N(k) \simeq \frac{ab\,k^2}{4\pi} = \frac{Ak^2}{4\pi}.$$

The leading counting law extends to smooth or piecewise smooth domains, and the Dirichlet boundary gives a subleading term $-Pk/(4\pi)$ for perimeter P under the usual asymptotic conditions. It predicts the general growth of the level count. At small k, discreteness and geometry dominate and an individual ground level cannot be inferred from area alone.

Test an inference about shape

For an electron in a 2.00 nm square, $\hbar^2/(2m_e) = 0.03810$ eV nm² and the ground energy is

$$0.03810 \times 2\pi^2/(2.00)^2 = 0.188\,\mathrm{eV}.$$

Doubling just one side gives

$$0.03810\pi^2(1/4 + 1/16) = 0.118\,\mathrm{eV}.$$

Doubling both gives 0.0470 eV. These outcomes distinguish directional stretching from uniform dilation.

Now compare the two irregular wells in the Dirichlet eigenmode experiment. With a physical length unit of 1 nm, their first dimensionless eigenvalues are both about 2.538, giving electron ground energies about 0.0967 eV. Predict their energies with a 2 nm unit, then verify the lowest stored level drops to 0.0242 eV for the rescaled well while its dimensionless eigenvalue stays fixed. Agreement of a few computed levels within numerical error alone cannot establish exact isospectrality. Known noncongruent isospectral domains show why even a full list of spectral lines need not uniquely specify a boundary.

The mass scaling offers a separate diagnostic on the conversion from spatial eigenvalues. Replacing the electron by a hypothetical particle of double its mass halves every mapped energy at fixed shape and physical size. No spatial eigenproblem has to be solved again because the mass occurs only in the conversion factor.

Open one direction and obtain a threshold spectrum

A closed rectangular well has discrete levels in both directions. A long conducting strip with hard side walls and open longitudinal ends has confinement in only one. Assume a nonrelativistic particle of effective mass m^*, width W in y and unrestricted position x, with zero interior potential. A separated solution is

$$\psi = e^{ik_x x}\sin(n\pi y/W).$$

The energy follows directly from the Laplacian,

$$E_n(k_x) = \frac{\hbar^2 n^2\pi^2}{2m^*W^2} + \frac{\hbar^2 k_x^2}{2m^*} = E_n^{\perp} + \frac{\hbar^2 k_x^2}{2m^*}.$$

Each transverse eigenfunction supplies a propagation channel only when the total energy exceeds its threshold $E_n^{\perp}$. Below threshold k_x is imaginary and a corresponding component excited at an interface decays along the strip. The group velocity in an open channel is $\hbar k_x/m^*$ and tends to zero at threshold.

For a normalization length L_x, periodic boundary conditions place successive k_x values $2\pi/L_x$ apart. Counting both propagation directions gives $N_n/L_x = k_x/\pi$ below a chosen energy above threshold, per spin state. Differentiating with respect to energy yields a density of states per unit length proportional to $(E - E_n^{\perp})^{-1/2}$. Slow longitudinal motion near an opening creates an enhanced number of states per energy interval. The ideal divergence is integrable: the total number added in a small energy window remains finite. Finite temperature, disorder and instrumental resolution round the singular feature.

Count channels in a semiconductor strip

Take $W = 20.0$ nm and $m^* = 0.0670 m_e$. Then $\hbar^2/(2m^*) = 0.0381/0.0670 = 0.569$ eV nm². The first threshold is

$$E_1^{\perp} = 0.569\pi^2/(20.0)^2 = 14.0\,\mathrm{meV}.$$

The second is 56.1 meV and the third 126 meV. At 70.0 meV two transverse modes carry travelling waves. The second has kinetic energy in the longitudinal direction of 13.9 meV, giving $k_x = \sqrt{0.0139/0.569} = 0.156$ nm⁻¹. Its wavelength along the strip is $2\pi/k_x = 40.3$ nm, considerably longer than the width.

Reduce the width to 16.0 nm while keeping the reservoir energy fixed. Every transverse threshold rises by $(20/16)^2 = 1.5625$, so the second moves to 87.7 meV and closes. A wave launched from the original wide section may then reflect or populate an evanescent component at the narrowing. Resolving those reflection amplitudes requires matching fields at the interface. The threshold count alone predicts how many modes remain available far beyond it, furnishing a useful test before doing a full scattering calculation.

A threshold can be recognized without measuring absolute particle energy if the width is varied systematically. Plot an identified onset energy against inverse squared width. In a strip with a constant, spatially uniform band-edge offset the slopes are in the ratio of squared transverse quantum numbers while the intercept gives that common offset. A finite confining barrier permits wavefunction penetration beyond the nominal geometric boundary and lowers the transverse energy compared with a hard wall of the same

width. Checking several mode onsets helps distinguish effective width changes from a shift of the entire interior potential.

Such comparisons should use the same effective mass throughout the fitted energy range.

Try it · Scale the recorded polygon spectrum

Keep the polygon shape and hard-wall boundary condition fixed while changing its uniform length unit or the particle mass. The reference is the actual OneSimulation polygon-A ground energy, 0.09669586 eV at 1 nm for an electron. This calculation converts that stored eigenvalue into a physical energy; it does not recompute the spatial mesh. Doubling the length should give 0.02417397 eV. Doubling the mass instead should give 0.04834793 eV.

Open the OneSimulation polygon eigenproblem, code and output

Open this live scaling calculation

Try the calculation

Uniform length scale (nm) 1

Particle mass / electron mass 1

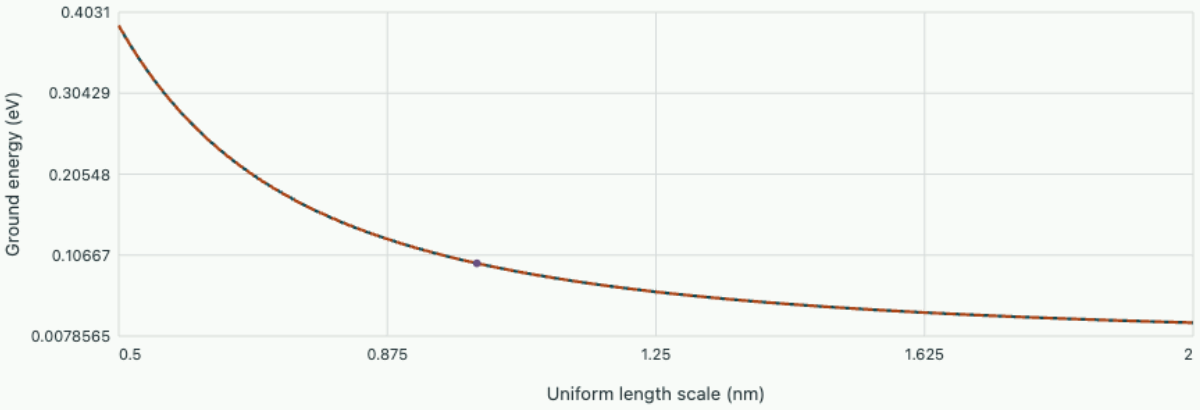

Ground energy 0.096696 eV. All levels scale by 1 relative to the recorded 1 nm, electron-mass polygon. The plotted baseline 0.09669586 eV is the recorded OneSimulation polygon-A ground energy; this algebraic rescaling preserves its shape and hard-wall condition and does not re-solve the 2D mesh.

Lift a square cavity's degeneracy at fixed area

Stretch a hard-wall two-dimensional box without changing its area $A = L_x L_y$. In units $\pi^2\hbar^2/(2mA)$, separation of variables gives $n_x^2/r + n_y^2 r$. The overall length scale is fixed, so all shifts here arise from shape.

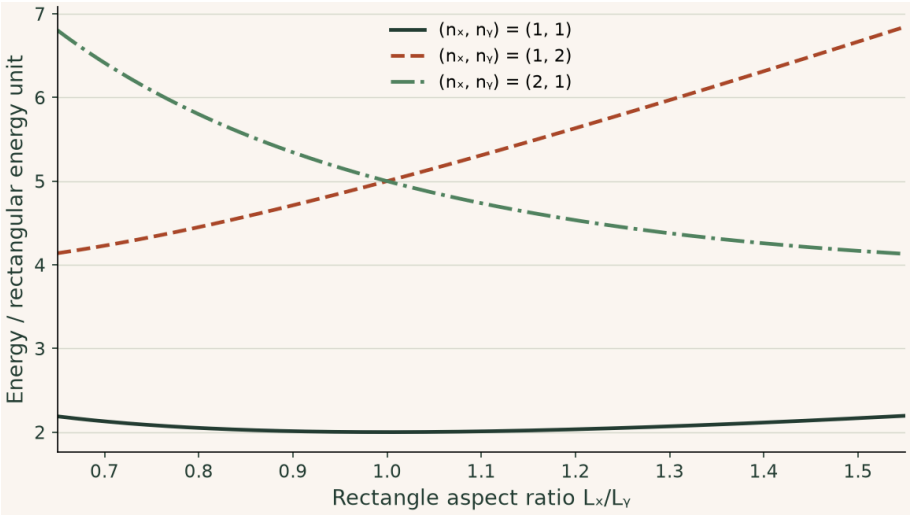

At the square the (1,2) and (2,1) states share energy 5. At aspect ratio 1.25 the levels move to 5.8 and 4.45. These independent separable branches cross exactly at the square; a generic symmetry-breaking coupling can instead cause level repulsion.

Try it · Billiard

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Aspect ratio Lx/Ly at fixed area

1

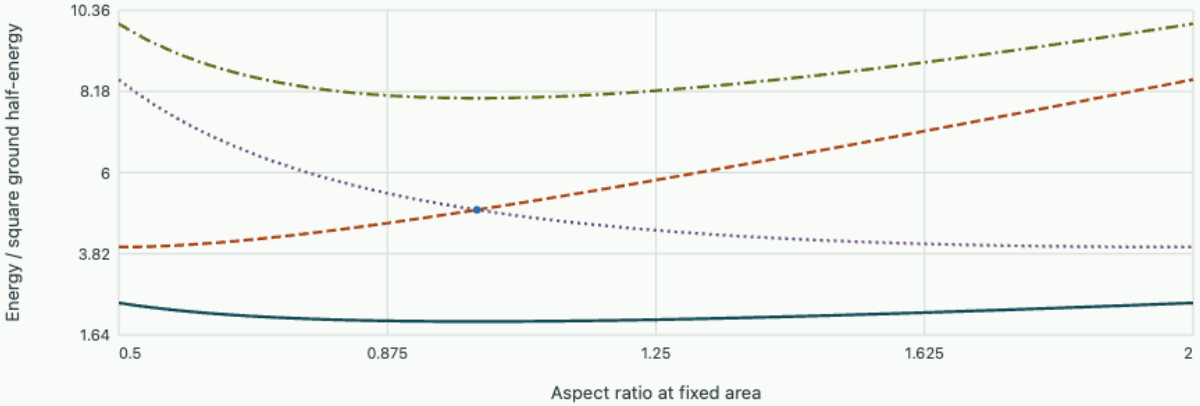

E12 = 5, E21 = 5. Their crossing is exact for the separable hard-wall rectangle.

▶ **Read values from the plot**

Study the full derivation

What a momentum eigenstate means in a finite experiment · Quantum states, dynamics and measurement

A detector never counts a perfectly sharp momentum eigenket. Derive continuum normalisation from a finite box, change the spectral measure from momentum to energy and predict the full evolution of a normalisable Gaussian packet.

Problem sheet 6 · Coherence and the confinement spectrum

Reconstruct a two-state density matrix, identify a degeneracy splitting and test a geometric inference from high-energy levels.

Tomography of an imperfect preparation

For independently repeated preparations the probabilities for positive Pauli x, y, z outcomes are $0.70, 0.40, 0.60$. Determine the Bloch vector, density matrix and purity. What would dephasing in the z basis do to each probability at long times?

▼ **Show the reconstruction**

Each Bloch component is $2P_{+} - 1$, giving $(0.40, -0.20, 0.20)$. From $\rho = (I + \mathbf{r}\cdot\boldsymbol{\sigma})/2$, diagonal entries are $0.60, 0.40$ and the upper off-diagonal entry is

$$(r_x - i r_y)/2 = 0.20 + 0.10i.$$

Purity is

$$(1 + |\mathbf{r}|^2)/2 = (1 + 0.24)/2 = 0.62.$$

Pure z dephasing removes the two transverse components while preserving populations, so the final positive probabilities are $0.50, 0.50, 0.60$. The nonzero y component would be missed by measuring only one transverse basis.

Split a square-well doublet

For an electron in a 3.00 nm square well compare the $(2,1)$ and $(1,2)$ energies. Stretch the x side to 3.30 nm while holding y fixed and compute both levels. Take $\hbar^2/(2m_e) = 0.03810$ eV nm².

▼ **Show the directional energy shifts**

In the square both have $E = 0.03810\pi^2(5/9) = 0.209$ eV. After stretching, the $(2,1)$ level has $E = 0.03810\pi^2(4/3.30^2 + 1/3.00^2) = 0.180$ eV. For $(1,2)$ the factor is $1/3.30^2 + 4/3.00^2$, yielding 0.202 eV. Modes with more rapid oscillations along the stretched axis lose more gradient energy. Uniform scaling of both dimensions would preserve the degeneracy.

Estimate the area from a level count

A spinless two-dimensional Dirichlet well has about 2000 levels up to $k = 50.0$ nm⁻¹. Use the leading counting law to estimate area. If an independent perimeter measurement gives 11.5 nm, apply the first boundary correction and discuss the precision of the inferred shape.

▼ **Show the asymptotic inversion**

Using $N \simeq Ak^2/(4\pi)$ gives $A \simeq 4\pi(2000)/2500 = 10.05$ nm². With $N \simeq Ak^2/(4\pi) - Pk/(4\pi)$ the inferred area increases by $P/k = 0.230$ nm² to about 10.28 nm². Low-order terms and fluctuations from integer level counts limit precision at finite k. Neither area nor perimeter fixes the full boundary, and even distinct shapes with identical spectra are possible. The estimate uses each spatial state once; including a separate spin degeneracy would double a counted list.

Recover a state from three spin measurements

Three independently prepared ensembles give probabilities for spin-up of 0.80 along x, 0.50 along y and 0.90 along z. Reconstruct the density matrix, decide whether the state is pure, and predict the result of an x measurement after a z measurement whose result is discarded. Repeat the purity check if the z probability was 0.95.

▼ **Worked solution**

For each Pauli measurement the mean is p(+) − p(−) = 2p(+) − 1. The Bloch vector is (0.6, 0, 0.8). In the z basis the diagonal entries are 0.9 and 0.1 and both off-diagonal entries equal 0.3. Its trace is one and its determinant vanishes. Thus its eigenvalues are one and zero. A state vector is (√0.9, √0.1) up to global phase.

An unread projective z measurement removes the off-diagonal terms. The z probabilities stay fixed, while the subsequent x-up probability becomes 0.50. The new purity is 0.9² + 0.1² = 0.82. Selecting the z-up subset instead would leave a pure z-up state, which also gives 0.50 in x. Testing z again distinguishes selection from mixing.

With z-up probability 0.95 the reconstructed Bloch length squared would become 0.6² + 0.9² = 1.17. The nominal matrix has a negative eigenvalue. Finite-sample fluctuations or calibration differences can produce such estimates. A physical reconstruction fits a positive matrix to the original counts. The three quoted probabilities alone do not specify their statistical uncertainty.

Explore these relations

- For dephasing, compare coherence loss and an echo under the stated static-detuning model. Open the live calculation.
- For the square-well doublet, change the rectangle’s aspect ratio and follow the degeneracy splitting. Open the live calculation.
- For the area estimate, compare discrete level counts with the area and boundary terms. Open the live calculation.

Which energy change really establishes a length change?

The numerical polygon gives a ground energy of 0.09670 eV for an electron with a 1 nm length unit. A second device yields 0.04835 eV and the same dimensionless ground eigenvalue. Infer its length if the mass is

unchanged. Construct a different explanation with unchanged geometry. Would measuring the second energy level distinguish the two explanations for a uniform effective mass?

▼ **Work through the solution**

The hard-wall problem gives

$$E_n = \hbar^2 \lambda_n / (2 m \ell^2).$$

A factor of two drop at fixed mass requires $\ell_2 / \ell_1 = \sqrt{2}$, or 1.414 nm. At the original scale, doubling the effective mass produces the same ground energy. Both mechanisms multiply every energy by one half under the stated assumptions. A second eigenvalue ratio therefore cannot resolve this ambiguity. An independent length or mass measurement can.

Use the OneSimulation laboratory to check your answer

Double mass or increase every length by the square root of two. Both changes halve every energy, while the stored dimensionless polygon eigenvalues stay fixed. Open the corresponding live laboratory reference.

Chapter 13 · When a free quantum packet initially gets narrower

A position distribution and a momentum distribution alone do not determine the early motion of a wave packet. Reconstruct free expansion from their covariance and compare two preparations with the same widths.

Ask for a width, rather than a single trajectory

Prepare many identical particles, release them in free space and measure one position from each preparation at a chosen delay. The spread in that collection of positions is the observable $\sigma_x(t)$. A particle detected early cannot also contribute an undisturbed late-time position to the same release. Quantum expectations below describe an ensemble of repeated preparations.

For a free particle in one dimension $H = p^2/(2m)$. In the Heisenberg picture the preparation stays fixed while the operators carry the time dependence. The identity $[AB, C] = A[B, C] + [A, C]B$ and $[x, p] = i\hbar$ give

$$\dot{p} = (i/\hbar)[H, p] = 0,$$

$$\dot{x} = (i/\hbar)[H, x] = p/m,$$

$$x(t) = x(0) + tp(0)/m.$$

Subtract the expectation value on both sides. Squaring this centred operator requires both multiplication orders for its cross term because x and p do not commute. Set $C = (\langle \Delta x \Delta p + \Delta p \Delta x \rangle)/2$, a real number with units of action. The exact result for any free preparation with these finite moments is

$$\sigma_x^2(t) = \sigma_x^2(0) + \frac{2t}{m} C + \frac{t^2}{m^2} \sigma_p^2(0).$$

Give a Gaussian a spatially varying phase

Let the initial wavefunction be proportional to $\exp[-x^2/(4s^2) + ibx^2/(4s^2)]$, with $s > 0$. Its density is a Gaussian of variance s^2 for every real b. Act with $-i\hbar\,\partial_x$ before taking moments. The derivative gives

$$p\psi = \hbar(b + i)x\psi/(2s^2).$$

Consequently $\langle p \rangle = 0$, $\sigma_p^2 = \|p\psi\|^2 = \hbar^2(1 + b^2)/(4s^2)$, and $C = \hbar b/2$, the real part of $\langle xp \rangle$. For negative b, probability on the right initially flows left and probability on the left flows right.

$$\frac{\sigma_x^2(t)}{s^2} = 1 + 2b\,u + (1 + b^2)u^2,$$

$$u = \frac{\hbar t}{2ms^2}.$$

Stationarity of the quadratic gives

$$u_* = -b/(1 + b^2).$$

A future focus occurs for negative b, and its width is $s/\sqrt{1+b^2}$. The preparation saturates $\sigma_x^2\sigma_p^2 - C^2 = \hbar^2/4$, so a position focus does not violate uncertainty. It removes the initial position–momentum correlation and leaves the same large momentum spread.

Predict two releases

Take $b = -2$ and $s = 4.0$ µm. At $u = 0.40$ the width falls to

$$4/\sqrt{5} = 1.79\,\mu\text{m}.$$

With $b = +2$, both initial marginal widths are unchanged, but at the same time the packet has width $4\sqrt{3.4} = 7.38$ µm. An initial image cannot distinguish the two releases. Measuring a second image at known positive delay resolves the sign.

For a third preparation use $b = 0$. This one starts with the smallest compatible momentum spread for the specified s and reaches $s\sqrt{2}$ at $u = 1$. In all three cases the mean stays at the origin. Translating the packet or giving it a common mean momentum would translate that mean trajectory without changing the variance formula.

Find the focus

Adjust the phase curvature. The plot compares the chosen preparation with an unchirped Gaussian of the same starting position width. Locate the focus with the derivative of the variance before reading its time from the result.

Open this experiment online

Try the calculation

Initial phase curvature b

-2

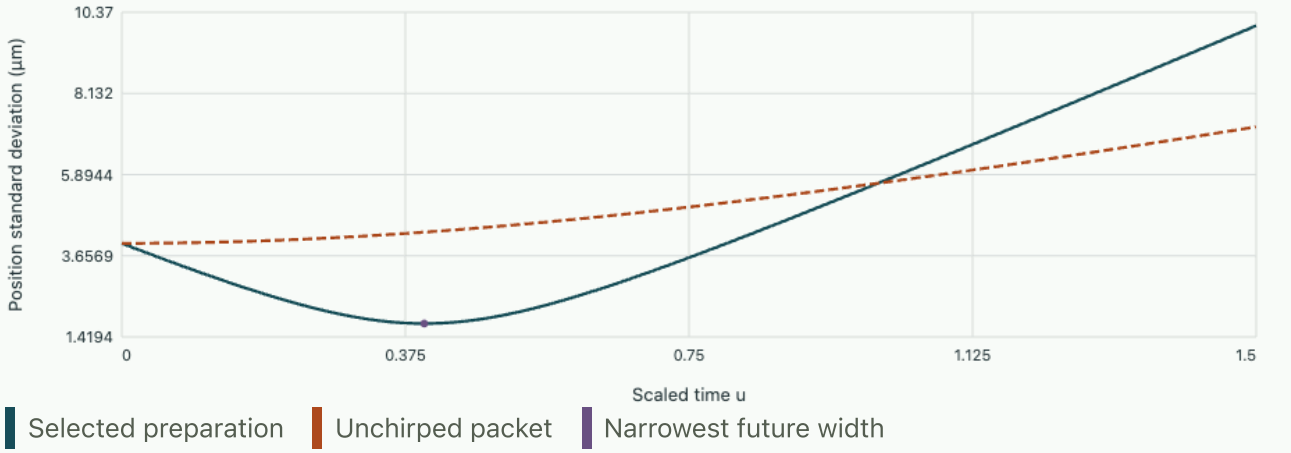

For u ≥ 0 the minimum width is 1.789 µm at u = 0.4.

Tutorial on a missing phase

A Gaussian initially 3.0 µm wide focuses to 1.5 µm after release. Find its initial chirp and dimensionless focusing time.

▼ **Work through the solution**

The ratio $s/s_{\min} = 2$ gives $1 + b^2 = 4$. A future focus selects $b = -\sqrt{3}$. Hence

$$u_* = \sqrt{3}/4 = 0.4330.$$

A packet with the opposite sign would have reached that same minimum before the stated release time under free evolution.

Two detectors measure packet widths at scaled delays $u = +0.20$ and $u = -0.20$ relative to a chosen reference preparation. Their squared widths differ by $-25.6\,\mu\mathrm{m}^2$, with $s = 4\,\mu\mathrm{m}$. Infer b from just this difference.

▼ **Work through the solution**

Subtracting the width-squared formula at opposite delays cancels the constant and quadratic terms, giving $4s^2 bu$. Hence

$$b = -25.6/(4 \times 16 \times 0.20) = -2.$$

The sign comes entirely from the position–momentum covariance. Choosing the reference preparation partway through free expansion makes measurements on both sides physically meaningful.

Can a non-Gaussian free state with finite second moments initially contract? What minimum variance follows from the general expression?

▼ **Work through the solution**

Differentiate the moment formula to get

$$t_* = -mC/\sigma_p^2.$$

A future minimum requires $C < 0$. Its variance is $\sigma_x^2(0) - C^2/\sigma_p^2$, bounded below by $\hbar^2/(4\sigma_p^2)$ using the Schrödinger–Robertson inequality. The Gaussian in this section attains the bound, but initial contraction requires only a negative covariance.

Further reading · Fließbach, Schrödingers Wellenmechanik · Tong, Quantum Mechanics, wavefunctions and time evolution.

Study the full derivation

What a momentum eigenstate means in a finite experiment · Quantum states, dynamics and measurement

A detector never counts a perfectly sharp momentum eigenket. Derive continuum normalisation from a finite box, change the spectral measure from momentum to energy and predict the full evolution of a normalisable Gaussian packet.

Chapter 14 · Boundary conditions that carry a current around a ring

An electron on an ideal narrow ring can sustain current in a stationary state. Work out how the topology of its configuration space and a magnetic flux through the hole select the allowed currents.

Recover the allowed state space from a boundary term

Wrap a coordinate s of length L around the ring. The points $s = 0$ and $s = L$ label the same physical place. For a uniform one-dimensional kinetic operator, integration by parts shows that the difference between $\langle \phi, H\psi \rangle$ and $\langle H\phi, \psi \rangle$ is proportional to

$$[\, \phi^* \partial_s \psi - (\partial_s \phi)^* \psi \,]_0^L .$$

Periodic values and periodic first derivatives make this expression vanish for any two admissible states. They also make the probability current continuous at the join. A plane wave $e^{iks}/\sqrt{L}$ then requires $kL = 2\pi n$ with integer n. The labels n and $-n$ have equal energy at zero flux but circulate oppositely. A real combination of the pair has zero net current. An energy measurement on its own does not tell us which combination was prepared.

Put flux through the hole

For a particle of charge q, the kinetic momentum along the ring is $-i\hbar\partial_s - qA_s$. Choose a constant tangential vector potential $A_s = \Phi/L$ giving the required circulation $\oint A_s ds = \Phi$. Define signed flux $\varphi = q\Phi/h$. Periodic plane waves remain eigenfunctions, while their kinetic momenta and energies now change:

$$p_{\mathrm{kin}} = \frac{h}{L}(n - \varphi),$$

$$E_n = E_L (n - \varphi)^2,$$

$$E_L = \frac{h^2}{2mL^2}.$$

A gauge transformation can eliminate the constant A_s locally. It multiplies a wavefunction by a phase whose value after one circuit differs from its initial value by $e^{-2\pi i\varphi}$ in this gauge. The transformed wavefunction must carry that twist at the join. Imposing both zero vector potential and unchanged periodicity would discard the physical flux dependence. Increasing φ by one reorders the spectrum by a shift of the integer labels.

Work out what a current measurement sees

One particle of velocity $v = p_{\mathrm{kin}}/m$ goes around the ring in a time $L/|v|$. The signed charge passing a fixed section per unit time is $I = qv/L$. Differentiating E_n at fixed n gives the same result:

$$I_n = \frac{qh}{mL^2}(n - \varphi) = -\frac{\partial E_n}{\partial \Phi}.$$

Take $\varphi = 0.30$. In units of E_L, the $n = 0$ energy is 0.09 and the $n = 1$ energy is 0.49. The currents in units of $qh/(mL^2)$ are −0.30 and +0.70 respectively. A probe that selects the lower-energy branch here sees a circulating state even though its probability density is uniform around the ring.

At $\varphi = 0.50$ the two energies agree at $0.25E_L$. Following a given eigenstate through this crossing and allowing a system to relax into its lowest energy at every flux describe different experiments. An ideal isolated perfectly uniform ring conserves its angular momentum label. Weak imperfections or a reservoir can change the populations and modify the behaviour near degeneracy.

Inspect the neighbouring branches

Change the signed flux continuously. Compare the slope of the branch containing $n = 0$ with its current. Predict where the lowest pair changes ordering and check that the complete set repeats with unit flux.

Open this experiment online

Try the calculation

Signed flux qΦ/h **0.3**

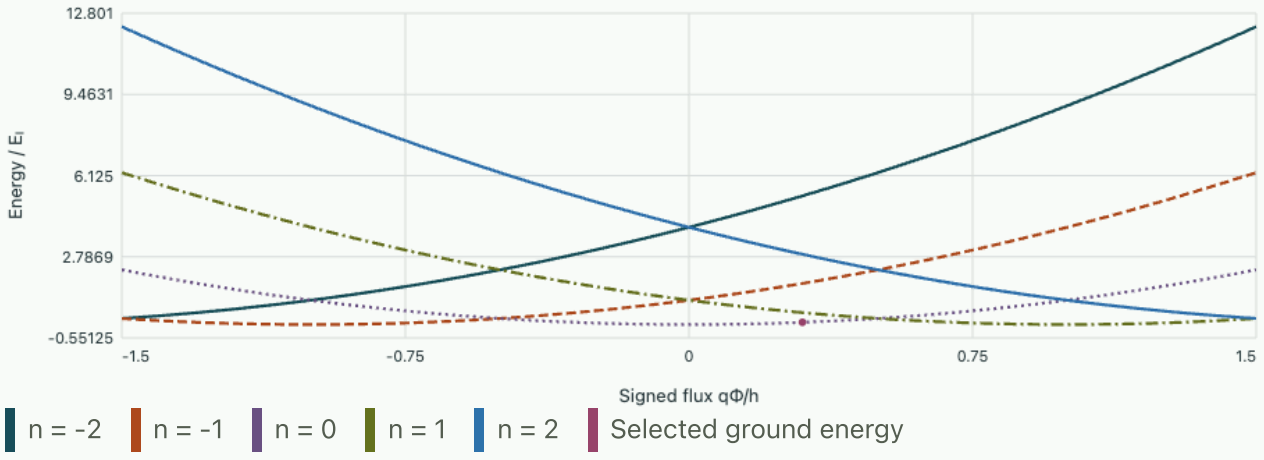

Lowest n: 0. Ground energy 0.09 E_I. The n = 0 current in units of qh/(mL²) is -0.3.

Tutorial on a ring of another size

Double the circumference at fixed mass and flux while preparing the same integer state. Predict the changes of its wave number, energy and current.

▼ **Work through the solution**

The canonical and kinetic wave numbers each halve because the integer label and signed flux are fixed. E_L scales as L^{-2}, and so does the current unit $qh/(mL^2)$. Energy and current both fall to one quarter. The flux period $h/|q|$ stays unchanged.

A ring with fixed flux is enlarged to twice its original circumference without changing the particle mass. How do energy gaps and branch currents change? Does the location of a crossing in flux move?

▼ **Work through the solution**

$E_L = h^2/(2mL^2)$ falls by a factor of four. The current scale $qh/(mL^2)$ falls by the same factor. The signed flux variable contains no circumference when total flux is held fixed, so crossings stay at half-integer flux. Holding the magnetic field fixed as ring area grows would instead change the enclosed flux.

At $\varphi = 0.50$, prepare an incoherent mixture with equal populations on $n = 0$ and $n = 1$. Calculate its current and energy. Can observing zero average current distinguish this mixture from zero flux?

▼ **Work through the solution**

The currents are opposite, each of magnitude half the current unit, and the mean vanishes. Both energies are $0.25E_L$. A zero-flux ground state also has zero current, but its energy is zero in the same calibration. Spectroscopy or a change of population separates the two cases.

Further reading · Tong, Particles in Magnetic Fields, gauge transformations and Aharonov–Bohm effect · Fließbach, Quantenmechanik, Eigenwerte und Eigenfunktionen.

Chapter 15 · Why the Coulomb radial series terminates

The hydrogen ground state admits a simple exponential guess. To understand whole spectral series, derive the radial recurrence and follow the two boundary conditions that force quantized energies.

Reduce the relative motion

A Coulomb attraction $V = -\kappa/r$ depends only on the separation of the electron and the positive centre. Separating the centre of mass leaves the reduced mass μ. Write the relative wavefunction as $u(r)Y_{\ell m}(\theta,\phi)/r$, with a normalized angular function. The r^2 in the volume element cancels the inverse square of the radial denominator, so normalization is

$$\int_0^\infty |u|^2 dr = 1.$$

The radial equation is

$$-\frac{\hbar^2}{2\mu}u'' + \left[\frac{\hbar^2\ell(\ell+1)}{2\mu r^2} - \frac{\kappa}{r}\right]u = Eu.$$

For a bound state define $k = \sqrt{-2\mu E}/\hbar$ and $\rho = 2kr$. At large radius the constant binding energy gives a decaying factor $e^{-\rho/2}$. At small radius the centrifugal term yields powers $\rho^{\ell+1}$ and $\rho^{-\ell}$. The regular Coulomb problem selects the first power. For $\ell = 0$ this enforces the regular three-dimensional boundary condition $u(0) = 0$. These two endpoint behaviours suggest factoring $u = \rho^{\ell+1}e^{-\rho/2}v(\rho)$.

Let normalizability select the energy

Differentiate the product twice and cancel the constant and inverse-square terms with the radial equation. The remaining polynomial equation involves only the dimensionless attraction $\lambda = \mu\kappa/(\hbar^2 k)$:

$$\rho v'' + (2\ell + 2 - \rho)v' + (\lambda - \ell - 1)v = 0.$$

Insert $v = \sum_{j=0}^{\infty} a_j \rho^j$ and equate equal powers. The derivative term proportional to a_{j+1} contains both the factor $j+1$ and the angular factor $j+2\ell+2$.

$$a_{j+1} = \frac{j + \ell + 1 - \lambda}{(j+1)(j+2\ell+2)}a_j.$$

A generic regular series has large-order coefficient ratio approaching $1/j$; its growing exponential component eventually overwhelms the front factor $e^{-\rho/2}$. For the regular Coulomb solution the coefficient of that growing component vanishes when the series terminates. A nonnegative integer n_r makes the numerator zero at $j = n_r$. Thus $\lambda = n_r + \ell + 1 = n$, and

$$E_n = -\frac{\mu\kappa^2}{2\hbar^2 n^2}.$$

Here n_r counts radial nodes in the open interval. Angular nodes and radial nodes enter the same energy label in this particular potential. Spin, fine structure, finite nuclear size and external fields are absent from this

calculation.

Distinguish two equal energies spatially

For $n = 2, \ell = 0$, the recurrence gives

$$a_1/a_0 = -1/2.$$

With $a = \hbar^2/(\mu\kappa)$, this state has $\rho = r/a$, so its radial probability vanishes at $r = 2a$. For $n = 2, \ell = 1$, the numerator vanishes already at $j = 0$. Here the radial probability is proportional to $r^4 e^{-r/a}$, with a maximum at $4a$ and no internal radial zero. The two preparations can yield the same spectroscopic energy but different radial histograms.

For an ideal hydrogenic ion with nuclear charge twice as large and the same reduced mass, κ doubles and a halves. Every radial feature just discussed moves to half its old radius. Each binding energy magnitude increases by four. For a Rydberg energy of 13.6 eV at charge one, the $n = 3$ to $n = 2$ photon has energy $13.6(1/4 - 1/9) = 1.89$ eV, rising to 7.56 eV for charge two.

Compare radial probabilities

Change the Coulomb strength while keeping the plotting unit equal to the original Bohr radius. The curves are separately normalized radial probabilities per unit of that original length. Locate the moving 2s node and the maximum of the 2p curve.

Open this experiment online

Try the calculation

Nuclear charge Z

1

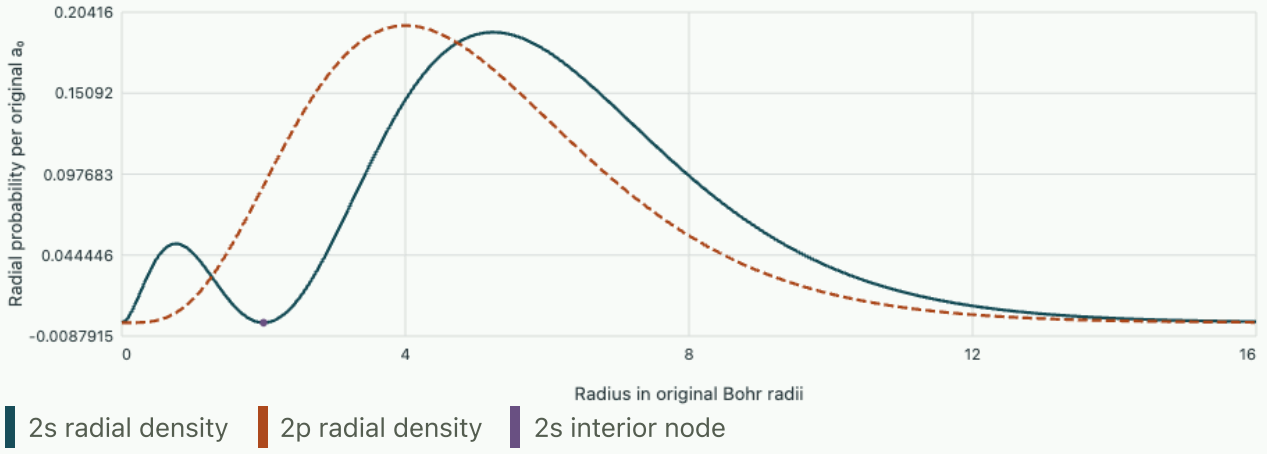

The 2s node is at 2 original Bohr radii. The most probable 2p radius is 4 original Bohr radii.

Tutorial on a third radial state

For $n = 3, \ell = 1$, obtain the first coefficient ratio, the radial node and its radius in units of a.

▼ **Work through the solution**

The recurrence at $j = 0$ gives

$$(2 - 3)/(2 + 2) = -1/4.$$

The next numerator vanishes, so $v = a_0(1 - \rho/4)$. The node is $\rho = 4$. Because $k = 1/(3a)$, the physical radius is

$$r = \rho/(2k) = 6a.$$

Compute the ratio of the most probable 2p radius to its expectation value. Why can those two characteristic radii differ for the same orbital?

▼ **Work through the solution**

With $x = Zr/a_0$ the radial density is $x^4 e^{-x}/24$. Its derivative vanishes at $x = 4$, whereas the mean is

$$\int_0^\infty x^5 e^{-x} dx/24 = 5.$$

Thus the ratio is $4/5$. The long tail shifts the mean above the mode. A local spatial probability density would be a different distribution again because it lacks the integrated spherical shell factor.

For an ideal Coulomb centre with $Z = 2$, predict the ratio of the $n = 3$ to $n = 2$ binding energies and the photon energy released from level three to level two in units of the hydrogen Rydberg. Neglect changes in reduced mass.

▼ **Work through the solution**

Binding magnitudes are $Z^2 \mathrm{Ry}/n^2$. Their ratio is $4/9$, and the positive emitted energy is

$$4\mathrm{Ry}(1/4 - 1/9) = 5\mathrm{Ry}/9.$$

The allowed orbital changes would also need to satisfy a radiative selection rule before an actual emission line appears.

Further reading · Tong, A Quantum Particle in Three Dimensions · Fließbach, Quantenmechanik, Dreidimensionale Probleme.

Study the full derivation

From a Coulomb orbit to the fine-structure multiplet · Quantum states, dynamics and measurement

Separate the hydrogenic Schrödinger equation, identify the origin of its large degeneracy and evaluate the three dominant relativistic corrections. The cancellation among them predicts an equality that spectroscopy can then test.

Chapter 16 · Measure an interaction by its low-energy scattering length

A cloud of cold, distinguishable atoms can scatter as if it had a size much greater than the range of its mutual force. Derive that length by matching a radial wave on the two sides of a simple potential well.

Replace a trajectory by a flux

For a localized central interaction an incident plane wave at large distance can be written as the sum of its incoming state and an outgoing spherical wave of amplitude $f(\theta)$. The spherical-wave intensity falls as $1/r^2$. Multiplying its radial probability current by the detector area $r^2 d\Omega$ and dividing by the incident current leaves $d\sigma/d\Omega = |f|^2$. This ratio is a geometrical area per unit solid angle even though the motion does not supply a definite classical impact parameter.

A partial wave with angular momentum ℓ feels the centrifugal term $\hbar^2 \ell(\ell+1)/(2\mu r^2)$. When $kR \ll 1$ for an interaction of range R, higher partial waves usually reach the interior only weakly. In the absence of higher-wave resonances keep the $\ell = 0$ wave. Far outside, its radial part is proportional to $\sin(kr + \delta_0)$. The phase shift measures how the interaction changes the oscillation relative to the free solution.

Background · Scattering length

$$k \cot \delta_0 = -1/a + O(k^2)$$

This sets the low-energy s-wave scattering scale for a short-range potential. Its sign carries phase information, and its size can greatly exceed the interaction range.

Match the zero-energy solutions

Choose $V = -V_0$ for $r < R$ and zero outside. At zero collision energy the exterior radial equation is $u'' = 0$, hence $u_{\text{out}} = A(r - a_s)$. The extrapolated intercept a_s defines the scattering length. Inside the well the regular solution is $B \sin(Kr)$,

$$K = \sqrt{2\mu V_0}/\hbar.$$

Continuity of u, u' at R allows their amplitude to cancel from the logarithmic derivative:

$$K \cot(KR) = \frac{1}{R - a_s},$$

$$\frac{a_s}{R} = 1 - \frac{\tan x}{x},$$

$$x = KR.$$

In the weak-well limit $\tan x = x + x^3/3 + \cdots$, so $a_s \simeq -Rx^2/3$. At $x = \pi/2$ a state reaches the bound-state threshold; the exterior zero-energy slope vanishes relative to its boundary value. The extrapolated intercept diverges. Its sign on either side of this crossing records the phase of the wave, rather than the sign of a physical radius.

Keep the finite-energy bound on the cross section

Expanding $\sin(kr + \delta_0)$ at small k gives $k \cot \delta_0 \simeq -1/a_s$, neglecting effective-range terms. The outgoing-wave amplitude and the integrated cross section of the isotropic component are

$$f_0 = \frac{1}{k \cot \delta_0 - ik},$$

$$\sigma_0 \simeq \frac{4\pi a_s^2}{1 + k^2 a_s^2}.$$

Far from resonance with $|ka_s| \ll 1$, $\sigma_0 \simeq 4\pi a_s^2$. At fixed nonzero k, increasing $|a_s|$ instead takes σ_0 towards $4\pi/k^2$. Keeping the imaginary term from outgoing flux is essential in this regime. A finite range adds terms of order k^2 to the real denominator. Identical atoms require exchange symmetrization and can have different allowed partial waves.

For $x = \pi/4$,

$$a_s/R = 1 - 4/\pi = -0.2732.$$

With $kR = 0.10$ the cross section is about $0.937R^2$. At $x = \pi$, $a_s = R$, and the same incoming wave gives $12.44R^2$. Thus increasing attraction does not produce a monotone core size. It changes a phase accumulated inside the well.

Change the depth of the well

The regularized cross-section curve stays finite across the first threshold. Vary kR to see how wavelength sets the height and breadth of the peak. Compare the selected depth with its two weak-scattering limits.

Open this experiment online

Try the calculation

Well phase KR **0.8**

Incoming wave number kR **0.1**

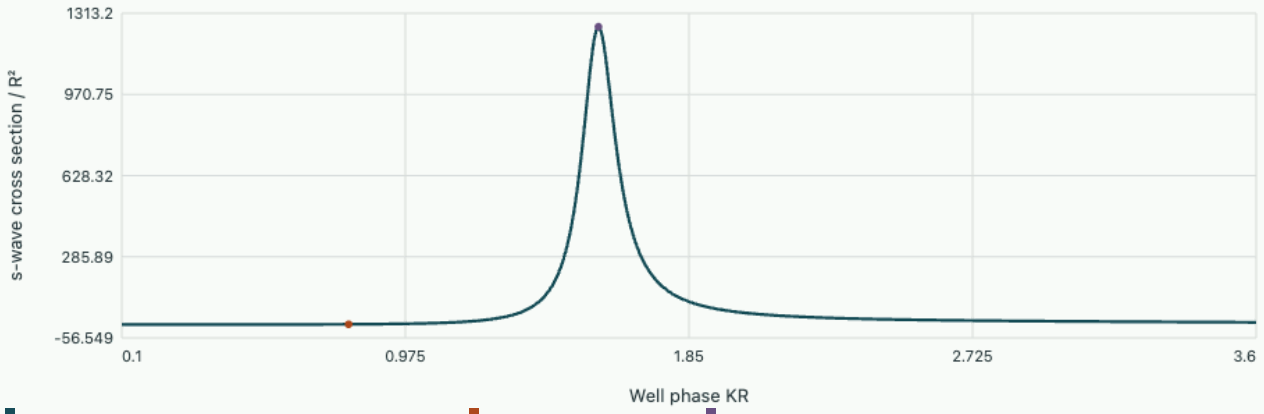

Leading low-energy cross section | Selected depth | First threshold

At KR = 0.8 the cross section is 1.035 R². The finite-k limit at resonance is 1257 R².

Tutorial on an ambiguous area

A low-energy measurement at $kR = 0.10$ finds $\sigma = 2\pi R^2$. Infer $|a_s|/R$ and identify what measurement is still missing.

▼ **Work through the solution**

Writing $b = |a_s|/R$, the equation $2\pi = 4\pi b^2/(1 + 0.01b^2)$ gives

$$b = 1/\sqrt{1.99} = 0.709.$$

The sign disappears in this cross section. A measurement sensitive to the scattering phase, or independent information about a near-threshold bound state, is needed to determine it.

A measured low-energy cross section equals half the unitarity value $4\pi/k^2$ at nonzero k. Infer the magnitude of the scattering length. Can this observation establish which side of a resonance the depth lies on?

▼ **Work through the solution**

Set $k^2a_s^2/(1 + k^2a_s^2) = 1/2$, giving $|a_s| = 1/k$. The total cross section depends on a_s^2 at this order. A phase-sensitive measurement of the scattering amplitude or independent knowledge of a bound state can resolve its sign. This inversion also presumes effective-range corrections are small.

Obtain the leading scattering length for a shallow well and predict how the zero-energy cross section changes when the well depth doubles.

▼ **Work through the solution**

Expand $\tan x = x + x^3/3 + \cdots$, giving $a_s/R \simeq -x^2/3$. Since $x^2 = 2mV_0R^2/\hbar^2$, the scattering length scales linearly and negatively with well depth. In the weak-scattering limit $\sigma \simeq 4\pi a_s^2$, so doubling a sufficiently shallow depth quadruples the cross section.

Further reading · Tong, Applications of Quantum Mechanics, Scattering Theory.

Study the full derivation

Inferring a short-range interaction from elastic scattering · Quantum states, dynamics and measurement

A beam of neutral particles encounters a weak, localised static potential. Derive differential cross section from current, extract phase shifts by radial matching and compare exact low-energy scattering with the Born approximation and its failure near a bound state.

Chapter 17 · A phase accumulated by slowly turning a magnetic field

A spin aligned with a slowly rotating field returns to its original measurement probabilities after one loop. Derive the additional phase that becomes observable when that path interferes with a reference path.

Separate following the state from acquiring a phase

Let a spin one-half have $H = -(\hbar\omega_0/2)\,\mathbf{n}(t)\cdot\boldsymbol{\sigma}$, with $\mathbf{n}$ a unit vector and constant $\omega_0 > 0$. The two instantaneous energies are separated by $\hbar\omega_0$. Prepare the lower-energy state and rotate $\mathbf{n}$ on a smooth path. If the rate of change of its eigenvector is small compared with the gap in angular-frequency units, transitions to the upper state are small. Following one eigenstate still allows a time-dependent complex phase.

Write $|\psi(t)\rangle = e^{i\alpha(t)}|n(t)\rangle$, substitute in the time-dependent Schrödinger equation and project onto $\langle n|$. Normalization makes $\langle n|\dot{n}\rangle$ imaginary. Consequently

$$\dot{\alpha} = -E_n/\hbar + i\langle n|\dot{n}\rangle.$$

The first contribution accumulates in proportion to elapsed time. The integral of the second depends on the path of the instantaneous eigenvector. The splitting into pieces changes with the arbitrary phase used to define $|n(t)\rangle$; after a closed loop the geometric phase in a single-valued gauge is invariant modulo 2π.

Evaluate a circle on the sphere of directions

Represent the aligned spin by the column $(\cos(\theta/2), e^{i\phi}\sin(\theta/2))^T$. Its mean spin points along the polar angles θ, ϕ. Hold θ fixed and increase ϕ by one revolution. Only the lower component contributes to the derivative. Its coefficient is its occupation weight:

$$i\langle n|\partial_\phi n\rangle = -\sin^2(\theta/2),$$

$$\gamma = -2\pi\sin^2(\theta/2) = -\pi(1-\cos\theta).$$

The solid angle enclosed by the latitude is $2\pi(1-\cos\theta)$, found by integrating $\sin\theta'\,d\theta'\,d\phi$ over the cap. The aligned spin accumulates minus half this angle. Reversing traversal changes its sign. As the cone closes to a point, the result vanishes continuously. Around the equator the result is $-\pi$.

Predict an interference fringe

Split a coherent preparation into two equal-amplitude paths and arrange the reference so the dynamical phases match. If only the signal path gains γ, an ideal balanced recombiner gives the bright-port probability $P = (1+\cos\gamma)/2$ for zero additional reference phase. The overall phase of one isolated spin would never produce this signal. The reference path supplies a phase comparison.

For $\theta = 60^\circ$, the solid angle is π, the geometric phase is $-\pi/2$, and $P = 1/2$. For a circuit of the equator the same port becomes dark. Reversing the 60° loop still gives one-half at zero reference phase. Introduce a known reference shift χ: the probability becomes $[1+\cos(\chi+\gamma)]/2$ in the chosen convention. Choosing $\chi = \pi/2$ changes the two directions of traversal into full and zero probability, revealing the sign.

Scan the reference phase

Set the cone angle, then vary the reference phase. The horizontal axis scans the cone angle and the marker value is returned in the readout. A source with finite path coherence would reduce the modulation depth, whereas unmatched dynamical phases would shift the fringe.

Open this experiment online

Try the calculation

Cone angle (degrees) **60**

Reference phase (degrees) **0**

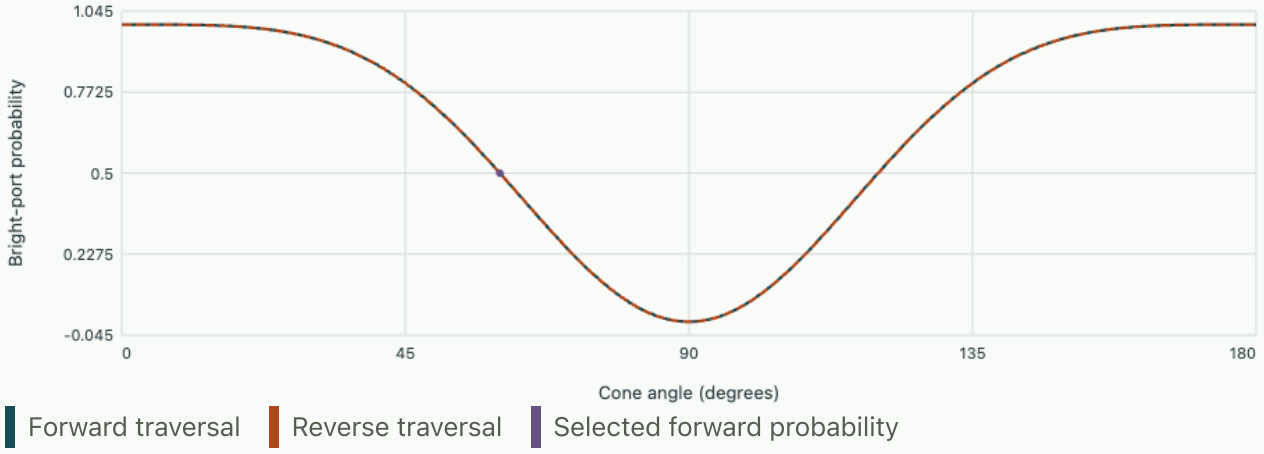

Forward geometric phase -0.5π radians. Forward port probability 0.5, reverse 0.5.

Tutorial on the following condition

A cone has angle 60°. At what rate scale does uniform azimuthal rotation stop being reliably adiabatic, and why can a change in loop duration alter a real experiment?

▼ **Work through the solution**

The off-diagonal eigenstate derivative has magnitude $|\dot{\phi}| \sin\theta/2$, so one requires $|\dot{\phi}| \sin\theta/(2\omega_0) \ll 1$. Increasing duration reduces that coupling but also changes the dynamical phase and can permit decoherence. An interferometer must control both even when the ideal geometric phase has no dependence on duration.

Keep the cone angle at 60° but make two complete positive circuits before recombination. What probability appears at the bright port with zero reference phase? Would performing one forward and one reverse circuit give the same result?

▼ **Work through the solution**

Each positive traversal contributes $-\pi/2$. Two give $-\pi$, hence zero bright-port probability. A positive and negative traversal instead add to zero, giving unit probability when dynamical phases are matched. These

outcomes test orientation and addition of the geometric phase over consecutive cycles.

Estimate the sensitivity to a cone-angle calibration error at 60°. Express the first-order phase error for one angular degree.

▼ **Work through the solution**

Differentiating in radians gives

$$d\gamma/d\theta = -\pi \sin\theta.$$

Multiplication by $\pi/180$ yields -0.0475 radians of phase error per added degree at 60°. At a port fringe extremum the intensity change from this small phase error begins quadratically. Near a half-bright fringe it appears linearly and is easier to detect.

Further reading · Tong, Topics in Quantum Mechanics, adiabatic theorem and Berry phase.

Study the full derivation

Adding angular momentum and predicting which lines can appear · Quantum states, dynamics and measurement

A three-component orbital wavefunction combines with a two-component spin. Derive the allowed total angular momenta, construct their amplitudes explicitly and predict both forbidden transitions and relative strengths in a polarisation-resolved spectrum.

Chapter 18 · A sudden expansion and the distribution of measured energy

A particle occupies the lowest state of an infinite square well. One wall is suddenly moved outward. Its position density initially stays where it was, although the list of energies an observer can measure changes immediately. An overlap integral predicts the complete distribution.

Preserve the state and change the basis

Take the old interval $0 < x < L$ and new interval $0 < x < bL$ with $b > 1$. The old state is $\sqrt{2/L}\sin(\pi x/L)$ inside its well, extended by zero across the newly accessible interval. The sudden approximation means evolution of the state during the control change can be neglected for the modes of interest. Immediately after the change, expand this same function in normalised new modes $\sqrt{2/(bL)}\sin(n\pi x/(bL))$. The measurement probability for mode n is the absolute square of its expansion coefficient.

With $u = x/L$ the coefficient is $2b^{-1/2}\int_0^1 \sin\pi u \sin(n\pi u/b)\,du$. Product-to-sum reduces it to the difference of two cosine integrals. At their upper endpoints $\sin[\pi(1 - n/b)] = \sin(\pi n/b)$ while $\sin[\pi(1 + n/b)] = -\sin(\pi n/b)$. Collecting the denominators gives

$$c_n = \frac{2}{\pi\sqrt{b}}\frac{\sin(\pi n/b)}{1-(n/b)^2}.$$

When $n/b = 1$ the expression has a removable singularity. The original integral then contains $\sin^2\pi u$, giving coefficient $1/\sqrt{b}$. When n/b is an integer other than one, orthogonality on the old interval makes the overlap zero. Completeness of the new sine basis requires $\sum_n |c_n|^2 = 1$, which gives a numerical truncation check.

Interpret an energy measurement

For $b = 2$, the first two coefficients are $8/(3\pi\sqrt{2})$ and $1/\sqrt{2}$. Their probabilities are 0.3603 and 0.5000. Mode three has coefficient $8/(5\pi\sqrt{2})$ and probability 0.1297. A measurement of the new ground energy therefore occurs only about 36 percent of the time. In contrast an adiabatic expansion that remains in the instantaneous nondegenerate ground state produces the new ground energy with probability essentially one.

The new level energies are $E_n = E_{\mathrm{old}}(n/b)^2$. The initial post-quench energy expectation is most safely evaluated as the kinetic quadratic form $(\hbar^2/2m)\int|\psi'|^2dx$. The wavefunction vanishes at the old boundary and its weak first derivative is square-integrable. The integral extends by zero and keeps its pre-expansion value E_{old}. For an instantaneous outward move the energy distribution broadens although this mean stays the same. The derivative jumps at the former wall, so applying the squared new Hamiltonian requires care. Generically $|c_n|^2 = O(n^{-4})$; the mean-energy sum converges, while the second energy moment diverges in this ideal step model.

See which observable converges first

Vary the expansion ratio and number of retained modes. The graph shows individual probabilities and accumulated probability. The readout separately sums the energy-weighted probabilities. A small omitted total probability can carry a larger fraction of the mean energy because higher modes have greater energy

per observation. Refine the cutoff until the quantity you intend to compare has converged. A real finite-speed wall or finite confining barrier changes the high-energy tail.

Open this experiment online

Try the calculation

Expansion ratio b 2

Retained eigenmodes 12

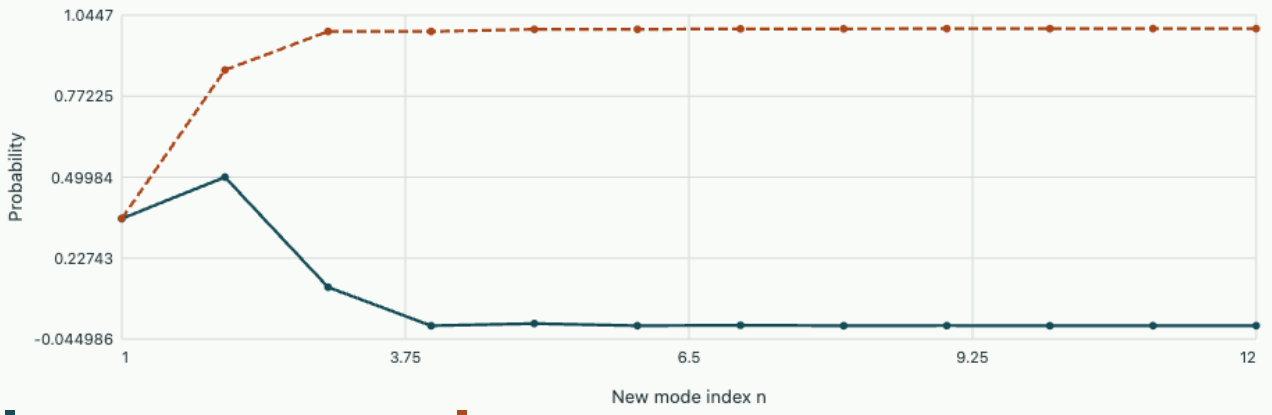

Probability retained 0.9997, omitted 0.0003189. Energy sum 0.9657 in units of the old ground energy. Exact infinite sum = 1.

Tutorial on expansion ratios and missing weight

Triple the box width. What is the probability that the measured post-quench energy equals the old ground energy? Which higher states have precisely zero weight?

▼ **Work through the solution**

The matching energy requires $n = 3$. Its coefficient is the removable-limit value $1/\sqrt{3}$, so probability is one third. States $n = 6, 9, 12, \ldots$ are orthogonal to the initial mode on the old interval and have zero overlap. Other modes generally have nonzero weight.

Is the original particle certain to be found somewhere in the old interval immediately after a sudden expansion? Is it certain to be there at a later time?

▼ **Work through the solution**

Initially the state has zero amplitude on the added interval, so the first probability is one. At later times the new energy components acquire unequal phases $e^{-iE_n t/\hbar}$. Their cancellations outside the old interval need no longer hold. Energy populations stay fixed under the new static Hamiltonian while position probabilities evolve.

Further reading · Cambridge quantum mechanics, sudden approximation and quantum quenches.

Study the full derivation

From a finite pulse to a continuum transition rate · Quantum states, dynamics and measurement

Switch a weak perturbation on for a measured duration. Its Fourier transform determines the amplitude of a discrete transition. Summing over a dense set of final states reveals where a time-independent transition rate comes from and where it fails.

Chapter 19 · Adiabatic following and a finite sweep through an avoided crossing

Two unperturbed energies cross as a control rises. A nonzero coupling opens a gap. Whether a prepared particle follows the lower energy through the crossing depends on how quickly the eigenvectors rotate compared with the dynamical phase they accumulate.

Find the gap and the rotation rate

Let $\hbar = 1$ temporarily, with Hamiltonian $H(t) = \begin{pmatrix} vt/2 & g \\ g & -vt/2 \end{pmatrix}$, $v > 0$, $g > 0$. The fixed basis is called diabatic because each basis vector keeps its identity as the detuning vt changes sign. Squaring the Hamiltonian gives $[(vt/2)^2 + g^2]I$; its energies are $\pm h(t)$ with $h = \sqrt{(vt/2)^2 + g^2}$. The minimum gap is $2g$. The corresponding Bloch axis points along $(g, 0, vt/2)$; write its polar angle as $\theta = \operatorname{atan2}(2g, vt)$.

Differentiating the angle gives

$$\dot{\theta} = -2gv/[(vt)^2 + 4g^2].$$

A spinor aligned with this axis rotates through half its geometric angle in Hilbert space. The magnitude of the off-diagonal derivative coupling is $|\dot{\theta}|/2$, largest at the crossing where it equals $v/(4g)$. Comparison with the minimum gap frequency $2g/\hbar$ gives a local adiabatic ratio $\hbar v/(8g^2)$. A slow sweep or large coupling suppresses transitions between instantaneous energy states. Following the lower adiabatic state exchanges the two fixed basis labels across the crossing.

Derive the fast-sweep scale independently

In the interaction picture of the diagonal part, the energy difference vt gives an accumulated relative phase $vt^2/(2\hbar)$. First-order perturbation theory gives a transition amplitude proportional to $(g/\hbar)\int_{-\infty}^{\infty} e^{ivt^2/(2\hbar)}dt$. The magnitude of this Fresnel integral is $\sqrt{2\pi\hbar/v}$, obtained by continuing the convergent Gaussian integral from a positive real damping parameter. The leading probability to change fixed basis states is therefore $2\pi g^2/(\hbar v)$, valid only when small. The exact infinite linear-sweep Landau–Zener solution resums this dependence: the probability of retaining the initial diabatic label is $\exp[-2\pi g^2/(\hbar v)]$. It refers to endpoints asymptotically far from the avoided crossing with constant coupling and no decoherence.

For $g = 0.5$ and $v = 1$ in units where $\hbar = 1$, the infinite-sweep retention probability is

$$e^{-\pi/2} = 0.2079.$$

Halving v squares it to 0.0432. The slow sweep therefore gives greater transfer between fixed basis states while preserving the lower-energy identity more faithfully. This distinction prevents reversing the interpretation of a measured probability.

Evolve a finite preparation numerically

Prepare the instantaneous ground state at detuning $vt = -10$ and finish at $+10$. Over each short step freeze H at its midpoint. Since $H^2 = h^2 I$, the constant-step propagator is exactly $\cos(h\Delta t)I - i\sin(h\Delta t)H/h$ in

the chosen units. The product is unitary up to roundoff; finite time resolution still changes accumulated phases. Compare the population of the original fixed basis state with its population in the instantaneous ground state. Their separation through the crossing directly shows where ideal adiabatic following fails. The calculated final population can differ from the infinite-sweep expression because of endpoint mixing and phases.

Open this experiment online

Try the calculation

Coupling g ($\hbar = 1$) **0.5**

Sweep rate v ($\hbar = 1$) **1**

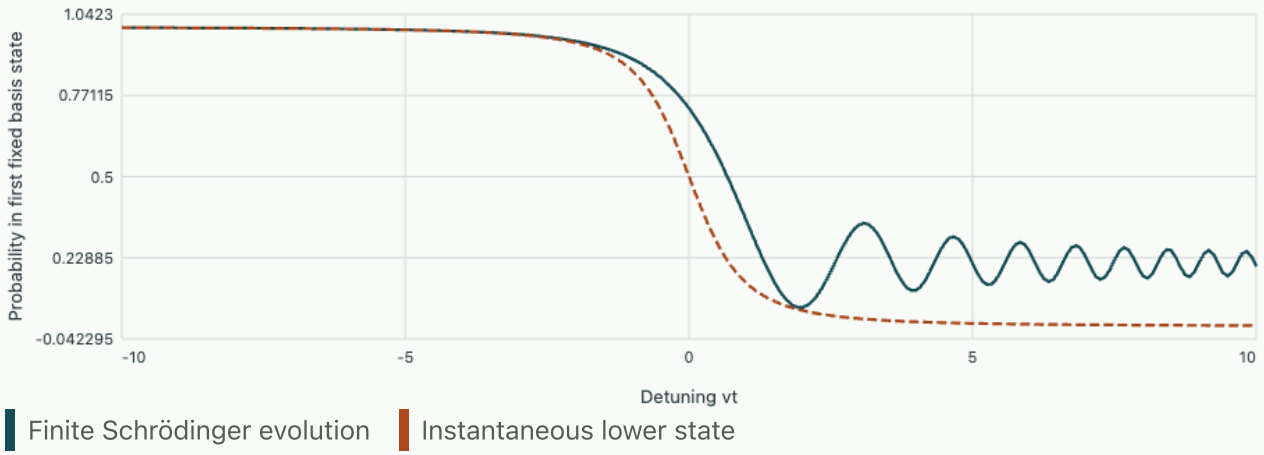

Final fixed-basis probability 0.2005, asymptotic Landau–Zener value 0.2079. Propagated state norm squared 1.

Tutorial on measurement choices

The coupling doubles while all else is fixed. How must the sweep rate change to leave the ideal infinite-sweep probability unchanged? Does a sweep over fixed detuning endpoints remain exactly equivalent?

▼ **Work through the solution**

Keep g^2/v fixed by multiplying v by four. The infinite-sweep exponent is unchanged. At fixed endpoints the ratios $|vt|/g$ at preparation and readout are now smaller, so the finite endpoints are less asymptotic. Their state mixing can change the observed result.

At zero detuning an energy measurement finds the lower eigenvalue with certainty. What would an immediate measurement in the fixed basis return?

▼ **Work through the solution**

The lower eigenstate there is $(1, -1)/\sqrt{2}$ in the chosen convention. Each fixed-basis outcome therefore occurs with probability one half. Certainty in the instantaneous energy is compatible with maximal uncertainty concerning the diabatic label.

Further reading · Tokmakoff, time-dependent quantum mechanics and Landau–Zener transitions · Rubbmark and colleagues, experimental passage through an avoided crossing.

Study the full derivation

When a magnetic field changes the quantum numbers of an atom · Quantum states, dynamics and measurement

Weak-field Zeeman slopes label total-angular-momentum multiplets. At stronger fields those slopes curve and intensities migrate. Solve an entire orbital-plus-spin manifold and infer coupling constants from measurable line positions and polarisation.

Chapter 20 · What a spin echo can recover

A collection of coherent spins can give nearly zero transverse signal because the phases of its members disagree. Reverse just one part of the evolution with a pulse and ask how much signal returns. The time dependence of the answer turns an apparent loss of coherence into a measurement of environmental memory.

Distinguish the state of a spin from the signal of a collection

Prepare each spin in the same equal superposition of the two z eigenstates. In a frame turning at the reference drive frequency, write its Hamiltonian during free evolution as $H(t) = \hbar\,\delta(t)\sigma_z/2$. The frequency error δ is angular frequency. At a constant detuning, the transverse spin vector simply rotates through angle δT. Its length does not change. An ensemble detector, however, adds the individual vectors. Two equally large transverse vectors pointing in opposite directions give no net signal.

It is useful to represent a transverse vector by a complex number and normalise its initially positive real value to one. The free-evolution contrast is the modulus of $W(T) = \langle e^{i\phi}\rangle$ with $\phi = \int_0^T \delta(t)\,dt$. The plus sign follows one choice of transverse quadrature and the modulus removes its arbitrary overall phase. If detunings are constant during each shot but Gaussian distributed between shots with standard deviation σ, the characteristic function gives

$$W = \exp(-\sigma^2 T^2/2).$$

Every spin can still be in a pure state even when this average is very small.

Now apply an ideal π rotation about the x axis at $T/2$. The Pauli identity $\sigma_x\sigma_z\sigma_x = -\sigma_z$ reverses the sign of the detuning Hamiltonian in the frame that follows the pulse. This algebra gives more than the common picture of fast and slow runners. It specifies exactly which operator the reversal changes. A Hamiltonian proportional to σ_x would retain its sign, and irreversible transitions into other degrees of freedom generally would not be undone. Treat the pulse as instantaneous and perfect for the derivation, and take preparation and readout axes to be phase locked.

Encode a pulse sequence by its signed accumulation time

Introduce the switching function $y(t)$. It is one throughout a free experiment, and it is one up to $T/2$ and minus one thereafter for a single echo. Observable dephasing is then controlled by $\phi_y = \int y(t)\delta(t)\,dt$, up to a global sign. A constant δ accumulates equal and opposite signed times and cancels exactly at the echo. For a linear drift $\delta(t) = a + bt$, the constant term cancels while the residual phase is $-bT^2/4$. The echo rejects a static error even when the distribution of that error is not Gaussian. The cancellation says nothing yet about time-dependent fields.

Assume next a zero-mean stationary Gaussian process with covariance $C(t-s) = \sigma^2 e^{-|t-s|/\tau_c}$. Its integrated phase is again Gaussian because integration is a linear operation. The logarithm of the coherence is minus one half of the phase variance. Substituting the two integrals gives

$$W_y = e^{-\chi_y},$$

$$\chi_y = \frac{1}{2}\int_0^T dt \int_0^T ds\, y(t)y(s)C(t-s).$$

The double integral counts correlated pieces of phase accumulation. In free evolution every product $y(t)y(s)$ is positive. For the echo, a time before the pulse and a time after it carry opposite signs. Long-lived correlations make substantial negative contributions that compensate the two positive within-half blocks. Noise whose memory is short cannot maintain a useful correlation across much of the pulse interval. Its effects are correspondingly harder to refocus.

Evaluate the covariance integral and inspect its limits

For free evolution the integrand is symmetric about $s = t$. Twice the area below that diagonal cancels the explicit one half. At fixed t the inner integral of the exponential covariance is $\sigma^2\tau_c(1 - e^{-t/\tau_c})$. Integrate once more and introduce $x = T/\tau_c$. For the echo it is efficient to partition the square into two identical within-half blocks and two cross blocks. Across one cross block the exponential separates into one integral over the first half and one over the second; its integral is $\sigma^2\tau_c^2(1 - e^{-x/2})^2$. Combining these terms gives

$$\chi_{\mathrm{free}} = \sigma^2\tau_c^2(x - 1 + e^{-x}),$$
$$\chi_{\mathrm{echo}} = \sigma^2\tau_c^2(x - 3 + 4e^{-x/2} - e^{-x}).$$

A direct subtraction is revealing:

$$\chi_{\mathrm{free}} - \chi_{\mathrm{echo}} = 2\sigma^2\tau_c^2(1 - e^{-x/2})^2.$$

It is nonnegative for this covariance at all T. This inequality depends on the chosen positive exponential correlation. A coherent narrow-band field can accumulate more phase under the echo than under free motion at certain frequencies, as the frequency filter will show.

For $T \ll \tau_c$, expand the exponentials before inserting numbers. The leading free exponent is $\sigma^2T^2/2$; the echo instead starts at $\sigma^2T^3/(12\tau_c)$. A cubic leading exponent carries information that simply fitting a single decay time would throw away. For $T \gg \tau_c$, both exponents grow with leading slope $\sigma^2\tau_c$ but differ by a constant. The long-time experiments can therefore have equal logarithmic slopes and visibly unequal amplitudes. A result quoted as one fitted coherence time needs its sequence, fitting interval and underlying noise assumption.

Open editable notebook Download .ipynb

Translate the pulse into a frequency filter

Test the same sequence using a single complex Fourier component of the detuning. Its contribution is proportional to $Y(\omega) = \int_0^T y(t)e^{i\omega t}dt$. Integrating the two constant pieces, and taking moduli to eliminate irrelevant origin phases, gives

$$|Y_{\mathrm{free}}|^2 = \frac{4\sin^2(\omega T/2)}{\omega^2},$$
$$|Y_{\mathrm{echo}}|^2 = \frac{16\sin^4(\omega T/4)}{\omega^2}.$$

Each quantity has units of time squared. At the origin take limits before evaluating: free accumulation gives T^2 and the echo gives zero. At one complete noise cycle during T, the free integral vanishes. The echo can instead turn the two signed halves into constructive accumulation, depending on the phase of the driving signal. Thus the relative response contains information about which frequencies contribute to the observed loss. Measuring a single echo is insufficient to reconstruct an arbitrary noise spectrum; many spectra have the same weighted integral.

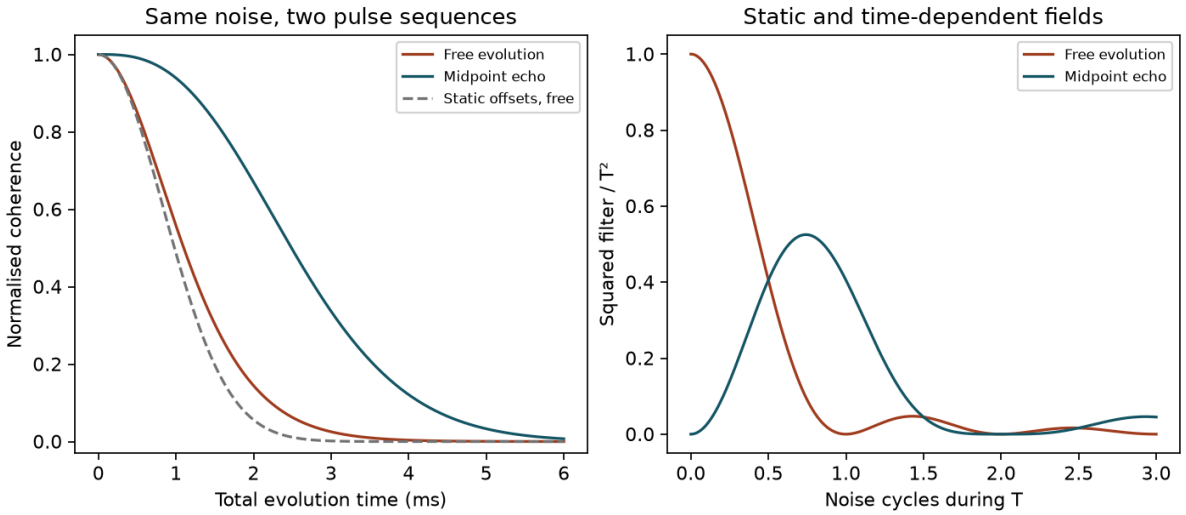

The midpoint pulse removes sensitivity to a constant offset. The right panel also shows frequencies at which it makes the experiment more sensitive than free evolution over the same total time.

Try it · Echo

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Noise correlation time (ms) 1

Rms detuning (rad/ms) 1

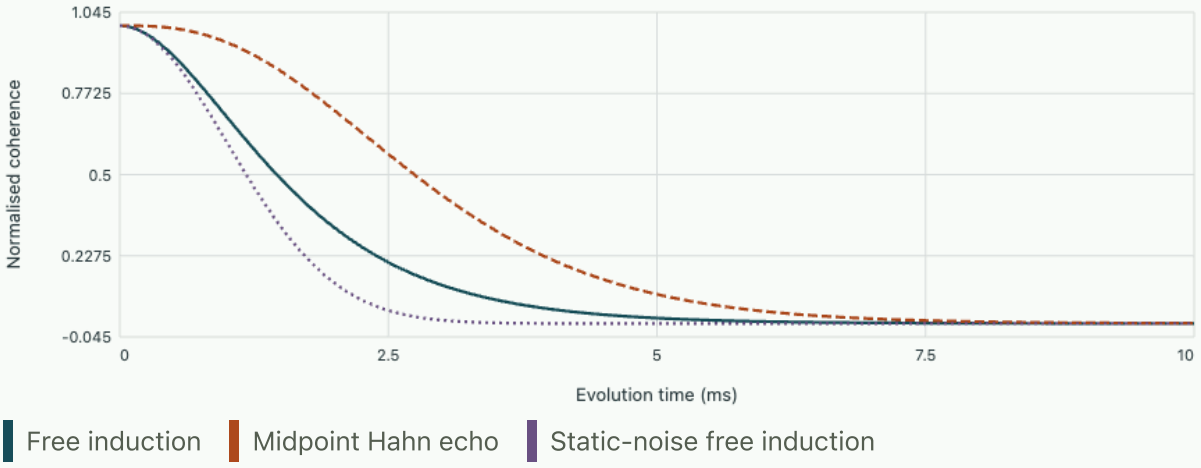

Stationary Gaussian noise with exponential covariance. The echo suppresses slowly varying noise; faster noise cannot be completely refocused.

Download the corresponding editable notebook. The browser recalculates the reduced equations directly; the notebook exposes the full calculation and its checks.

In the notebook take $\sigma = 1.2$ rad/ms, $\tau_c = 1.5$ ms and $T = 2.0$ ms. The free contrast is about 0.145 and the echo gives about 0.670. First vary only τ_c, keeping the instantaneous width σ fixed. Then compare the analytic result with a double sum over a time grid. Make the pulse lie on a cell boundary; an odd number of equal cells otherwise assigns unequal signed areas unless the central cell is split. A separate optional OneSimulation eigensolve constructs the propagator of one spin with a static detuning. Its unit return probability and the finite ensemble echo describe different preparations and different levels of modelling.

What can an imperfect apparatus identify?

A uniform offset in pulse timing puts the nominal π pulse at $T/2 + \varepsilon$. A static detuning now accumulates signed time 2ε, so a Gaussian ensemble is attenuated by $\exp[-2\sigma^2\varepsilon^2]$. A small flip-angle error also leaves part of the coherence in the wrong component. Neither effect is environmental loss of memory. One

diagnostic is to vary the programmed pulse position slightly in both directions while holding the final observation time fixed. Static inhomogeneity predicts a peak centred at balanced signed times.

Longitudinal relaxation, non-Gaussian telegraph noise, finite readout contrast and pulse bandwidth can all alter a measured sequence. An independently measured readout visibility can normalise the signal, while a separate population-relaxation measurement tests whether the transverse amplitude must decay even with perfect static refocusing. The Gaussian characteristic-function derivation uses the full probability law, not just its second moment. Specifying the same covariance for a non-Gaussian process need not reproduce the same contrast at strong coupling. The simpler second-moment calculation survives as a leading weak-phase approximation and can be challenged by repeating the experiment at different noise amplitudes.

Tutorial on timing, units and environmental memory

A static ensemble has an rms ordinary-frequency spread of 100 Hz. What free contrast and ideal echo contrast are expected after 2.0 ms?

▼ **Work through the solution**

Convert cycles per second to angular frequency first. The dimensionless spread of phase is $2\pi(100)(0.002) = 1.257$, giving $\exp[-1.257^2/2] \simeq 0.454$ for the free contrast. The ideal echo cancels every static phase and returns unit contrast when there are no other losses. Inserting 100 directly as an angular frequency would substantially overestimate the free signal.

The midpoint pulse in that experiment is late by 0.20 ms. Compute the echo attenuation due only to this error.

▼ **Work through the solution**

The remaining signed time is twice the timing error, or 0.40 ms. The rms residual phase is about 0.251 rad. Its characteristic function is about 0.969. A timing mistake can therefore look like a few per cent intrinsic signal loss even when every individual spin remains perfectly coherent.

A single-cycle sinusoidal disturbance causes little free dephasing over its exact period. Must the echo also remove it? Compare the normalised filter powers.

▼ **Work through the solution**

For $\omega T = 2\pi$, the free numerator vanishes. The echo numerator contains $\sin^4(\pi/2) = 1$, giving $|Y|^2/T^2 = 4/\pi^2 \simeq 0.405$. The phase of a deterministic sinusoid still matters for an individual shot. Averaging uniformly random starting phase turns the squared filter into the relevant sensitivity factor.

Further reading · E. L. Hahn · Spin Echoes, original experiment

Problem sheet 7 · Measuring two phases with one interferometer

An ideal coherent interferometer has two paths for a charged spin-half particle. One loop encloses a magnetic flux. A controlled field direction can also rotate the spin adiabatically in one arm. Use independent calibrations of the flux and spin phase to predict what the recombiner measures. Assume identical path intensities, perfect coherence and matched dynamical phases.

Calibrate the magnetic flux from a ring spectrum

A separate narrow ring made for the same particle has $E_n/E_L = (n-\varphi)^2$ with signed flux $\varphi = q\Phi/h$. Near zero flux, the $n=0$ energy is $0.0625E_L$. An independent current measurement on that branch gives

$$I/(qh/mL^2) = -0.25.$$

Find the flux and the nearest excited energy. Why are both observations useful?

▼ **Work through the solution**

The energy gives $|\varphi| = 0.25$. The signed current is $n-\varphi = -\varphi$ on this branch, so $\varphi = +0.25$. The nearest excited branch is $n=1$ at $0.5625E_L$. The excitation gap is $0.50E_L$. Energy alone leaves the sign undetermined.

Read off the energy branches and the signed ring current

Calibrate the spin phase by reversing a loop

Prepare the aligned spin state and carry the field direction once around a positive latitude at cone angle 60°. Find the geometric phase for each traversal direction. With no magnetic flux, predict the bright-port probabilities at reference shifts $\chi = 0$ and $\chi = +\pi/2$ for each direction.

▼ **Work through the solution**

The solid angle is $2\pi(1-\cos 60^\circ) = \pi$, hence the forward phase is $-\pi/2$ and the reverse phase $+\pi/2$. In $P = [1+\cos(\chi+\gamma)]/2$, zero reference phase gives one half for either traversal. A reference shift $+\pi/2$ gives one for the forward loop and zero for its reverse. The second setting reveals the sign.

Reverse the latitude and read the two bright-port probabilities

Predict the combined measurement

In the interferometer the paths enclose the flux once with an orientation giving magnetic phase $q\Phi/\hbar = 2\pi\varphi$. At zero reference shift use the preceding calibrations to predict the bright-port probability with the forward spin loop and with its reversal. Then reverse the magnetic flux without changing the spin protocol.

▼ **Work through the solution**

The magnetic phase is $+\pi/2$, so the forward total is zero and the port is bright with probability one. The reverse spin loop gives total phase π and zero probability. Reversing the flux swaps those outcomes because the total phases become $-\pi$ and zero. The two individual phases have different physical origins but their amplitudes multiply on the same path.

Check an imperfect observation

At the combined bright and dark settings an experiment instead records probabilities 0.91 and 0.09. Fit a visibility V in $P = [1 + V\cos\phi]/2$. Can these two points alone determine whether the loss was caused by dephasing or unequal path powers? Propose one independent observation.

▼ **Work through the solution**

Their difference is $V = 0.82$ and their sum equals one, consistent with the balanced offset in the model. An incoherent admixture and unequal path amplitudes can both reduce the observed visibility. Measuring each arm intensity with the other arm blocked tests imbalance. Scanning the reference phase through a full period additionally checks whether a phase offset displaced the nominal extrema.

Problem sheet 8 · How a pair of neutral oscillators acquires an attractive interaction

Two neutral atoms in their ground states have zero permanent electric dipoles. A molecule nevertheless feels a long-range attraction to another neutral molecule. Treat one coordinate of each atom as a quantised oscillator. Solve the coupled pair twice and decide which parts of the result would survive in three dimensions.

First make a prediction at fixed separation

Two equal masses m have displacements x_1, x_2 about fixed positive centres a distance R apart. In isolation both have angular frequency ω_0. The two induced electric moments are proportional to the displacements. The leading electrostatic interaction has energy $g x_1 x_2 / R^3$, with orientation and charges absorbed into g. Explain why a classical point resting at each minimum gains no energy, then ask what part of the Hamiltonian a quantum ground state would sample.

▼ **Work through the solution**

Write $k = m\omega_0^2$ and $\lambda = g/(kR^3)$. Along $x_2 = 0$ the potential rises as $k x_1^2/2$. Along the diagonals $x_1 = \pm x_2$ it rises at different rates. For $|\lambda| < 1$ both slopes of restoring force are positive and the minimum remains at the origin with classical potential zero. The quantum oscillators have finite widths at zero temperature. Their ground state depends on both diagonal stiffnesses. A shift in that energy can occur even if the average position of each mass stays zero.

Diagonalise the complete Hamiltonian

Starting from kinetic energy $(p_1^2 + p_2^2)/(2m)$, introduce normal coordinates and their conjugate momenta. Find both frequencies and the ground-state shift. Estimate it at $\lambda = 0.40$, in units of $\hbar\omega_0$, sufficiently accurately to distinguish the exact result from its leading approximation.

▼ **Work through the solution**

Use $x_\pm = (x_1 \pm x_2)/\sqrt{2}$ and the identical orthogonal transformation on momenta. Orthogonality preserves the kinetic energy and the canonical commutators. The product in the interaction is $(x_+^2 - x_-^2)/2$. The two uncoupled Hamiltonians have spring constants $k(1 \pm \lambda)$. Each costs half a quantum in its own ground state. Subtract the energy of the two separated oscillators before comparing separations.

$$\omega_\pm = \omega_0 \sqrt{1 \pm \lambda},$$

$$\frac{\Delta E}{\hbar\omega_0} = \frac{\sqrt{1+\lambda} + \sqrt{1-\lambda}}{2} - 1.$$

At 0.40 the two frequency ratios are 1.18322 and 0.77460. Their mean lies 0.02109 below one. The frequency that decreases loses more than the other gains because a square root is concave. Changing the sign of the coupling exchanges the normal coordinates and leaves the energy unchanged. This symmetry already excludes odd powers of the interaction.

Recover the same force from virtual transitions

Calculate the first two perturbative corrections in the basis of the isolated pair. Identify the intermediate state and its energy denominator. Obtain a power law for the force with the sign convention that increasing R is positive.

▼ **Work through the solution**

Let $x_0 = \sqrt{\hbar/(2m\omega_0)}$. Acting on an isolated ground state, x creates one quantum with amplitude x_0. The expectation of the interaction in $|00\rangle$ vanishes because each coordinate has zero mean. One action of its product leaves only $|11\rangle$. The matrix element is $k\lambda x_0^2 = \lambda\hbar\omega_0/2$, and the excited state lies two quanta above the ground state. The denominator $E_{00} - E_{11}$ is therefore negative.

$$\Delta E^{(2)} = -\frac{\hbar\omega_0}{8}\lambda^2 = -\frac{C_6}{R^6},$$

$$C_6 = \frac{\hbar\omega_0 g^2}{8k^2},$$

$$F_R = -\frac{6C_6}{R^7}.$$

At the specified coupling the quadratic approximation gives −0.02000. Expanding each square root and adding yields the next correction $-5\lambda^4/128$. At 0.40 that term is −0.00100. The formal expansion parameter is a stiffness ratio. Calling the interaction weak requires a separation large enough to make this ratio small, even if each microscopic charge is individually large.

What would distinguish the power law in data?

The downloadable OneSimulation model assembles the two-by-two stiffness matrix and computes its modes with a contour eigensolver. Taking square roots and summing their zero-point energies reproduces the plotted exact curve. The two separate frequencies in the CSV provide a stronger check than their sum.

Use separation units R_0 and set the coupling there to 0.40. Before opening the graph, predict the factor by which the leading energy shift changes at $2R_0$. Is the corresponding factor the same for force? Then change the coupling at the reference separation. Compare the exact and perturbative curves where the soft frequency approaches zero.

Open this experiment online

Try the calculation

Dipole coupling at the reference separation **0.4**

Separation to inspect (R_0) **1.4**

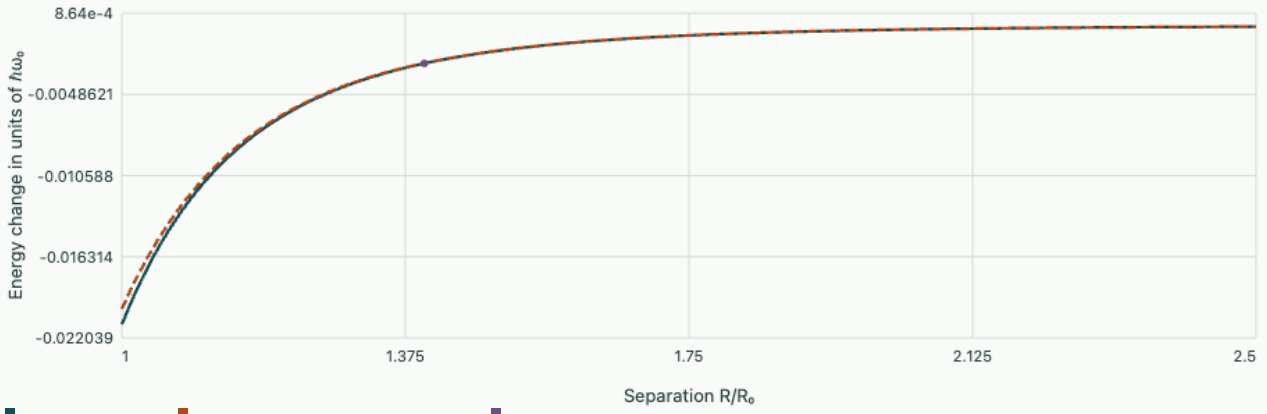

Exact pair | Second-order energy | Selected measurement

At the selected separation the energy shift is -0.002674 $\hbar\omega_0$. Second order gives -0.0026562 $\hbar\omega_0$.

Download the OneSimulation model

In your OneSimulation environment, run the model with these inputs.

```
python spectra-and-orbits.py pair --coupling 0.4 --separation 1.4
```

▼ Work through the solution

Doubling the separation divides the coupling by eight. The leading energy falls in magnitude by $8^2 = 64$, and the force by 128. For the reference coupling the exact shifts at R_0 and $2R_0$ have magnitudes about 0.02109 and 0.000313 in units of one quantum. The ratio of exact energies is about 67.4. A fit near strong coupling therefore gives a separation exponent steeper than six. Measuring it does not establish a new long-distance law. Repeat farther apart and check whether the estimate approaches six.

A mode becoming soft signals the breakdown of the quadratic model. At small separation real atoms overlap and exchange matters. At distances comparable with a transition wavelength, electromagnetic propagation time matters. This two-coordinate electrostatic model predicts neither limit. Extending it to three dimensions also requires carrying the dipole orientation tensor through the intermediate states. The coefficient found here is specific to the stated geometry.

Follow the independent development of stationary perturbation theory in the MIT Quantum Physics III course material.

Notebook practice · What a spin echo can recover

Quantum states, dynamics and measurement

Build the connection between preparation, time evolution and measurement. Develop the algebra, approximations and scattering theory needed to predict a spectrum or a set of count rates.

1. ## From two-path interference to the geometry of a quantum state

   Two preparations give the same counts at one analyser. Place a second analyser in a rotated basis and the counts separate. From this experiment develop amplitudes, phase, projective measurement, mixed states and the geometry of distinguishability.

2. ## What a momentum eigenstate means in a finite experiment

   A detector never counts a perfectly sharp momentum eigenket. Derive continuum normalisation from a finite box, change the spectral measure from momentum to energy and predict the full evolution of a normalisable Gaussian packet.

3. ## Why an oscillator can move without its packet spreading

   Start with the measured mass and curvature of a potential. Factor its Hamiltonian, prove the energy ladder, construct a freely moving Gaussian from number states and find when squeezing or anharmonicity changes that prediction.

4. ## Adding angular momentum and predicting which lines can appear

   A three-component orbital wavefunction combines with a two-component spin. Derive the allowed total angular momenta, construct their amplitudes explicitly and predict both forbidden transitions and relative strengths in a polarisation-resolved spectrum.

5. ## From a Coulomb orbit to the fine-structure multiplet

   Separate the hydrogenic Schrödinger equation, identify the origin of its large degeneracy and evaluate the three dominant relativistic corrections. The cancellation among them predicts an equality that spectroscopy can then test.

6. ## When a magnetic field changes the quantum numbers of an atom

   Weak-field Zeeman slopes label total-angular-momentum multiplets. At stronger fields those slopes curve and intensities migrate. Solve an entire orbital-plus-spin manifold and infer coupling constants from measurable line positions and polarisation.

7. ## Two electrons, exchange splitting and the cost of localisation

   Build a two-electron state that respects exchange, derive its charge and spin correlations, and then solve a two-site model exactly. The two examples reveal when a triplet is favoured and when an apparently weak hopping lowers the singlet instead.

8. ## From a finite pulse to a continuum transition rate

   Switch a weak perturbation on for a measured duration. Its Fourier transform determines the amplitude of a discrete transition. Summing over a dense set of final states reveals where a time-independent transition rate comes from and where it fails.

9. Inferring a short-range interaction from elastic scattering

A beam of neutral particles encounters a weak, localised static potential. Derive differential cross section from current, extract phase shifts by radial matching and compare exact low-energy scattering with the Born approximation and its failure near a bound state.

10. Distinguishing entanglement from ordinary shared randomness

Two distant analysers sometimes agree even when individual outcomes look random. Derive which correlations shared classical information can produce, compute the quantum prediction and identify what a finite table of coincidence counts can establish.

11. How an unobserved environment changes a quantum state

Model a qubit and a two-state environment by a joint unitary. Derive channels, the continuous-time master equation and the relation between longitudinal and transverse relaxation. A recorded event and an unread detector produce different conditional predictions.

12. How a turning point supplies the missing semiclassical phase

Far from a turning point the wavelength follows the classical momentum. At the turning point that approximation becomes singular. Derive the Airy connection, quantise a smooth well, distinguish a hard wall and estimate a tunnelling exponent with a defensible error check.

13. Measure geometric phase from three spin overlaps

A numerical eigensolver is free to choose a new complex phase for every eigenvector. A product of three measured overlaps removes these arbitrary choices. Relate its argument to a spherical triangle, then see why a vanishing overlap destroys phase information.

14. From electron exchange to atomic shells and molecular binding

An orbital model can explain why a closed atomic shell is special and why a shared electron can bind two protons. Both predictions need an energy as well as a rule for counting states. We construct the energies and check them against their separated-particle limits.

15. The scattering integral equation and successive Born terms

A weak potential often produces an angular pattern resembling its Fourier transform. What controls that approximation? A particle can propagate between two interactions with the same potential. Accounting for that propagation builds an integral equation and supplies a calculable second term.

16. Spontaneous radiation, stimulated transitions and thermal detailed balance

A coherent laser can drive Rabi oscillations in two atomic levels. A large cavity filled with thermal radiation presents many modes with random phases. In this limit one can infer the ratios of absorption, stimulated emission and spontaneous emission coefficients by requiring the observed equilibrium to be independent of the atoms used to probe it.

17. Which properties of an unknown spin can an instrument copy?

A device calibrated with spin up and down can write two perfectly distinguishable records. Asking what happens to a superposition reveals a limit on information extraction and explains why state tomography needs an ensemble.

Quantum states, dynamics and measurement · 1. From two-path interference to the geometry of a quantum state

Two preparations give the same counts at one analyser. Place a second analyser in a rotated basis and the counts separate. From this experiment develop amplitudes, phase, projective measurement, mixed states and the geometry of distinguishability.

Begin with an experiment that can distinguish two preparations

A single photon enters a lossless interferometer and can take either an upper or a lower path. Think of the paths as two distinguishable spatial modes labelled $0, 1$, with equal travel times unless a phase shifter is inserted. Record one binary result from each freshly prepared photon. Our model predicts outcome frequencies over many repetitions. A preparation that always follows path 0 is represented by $|0\rangle$, one that always follows path 1 by $|1\rangle$. Orthonormality, $\langle i|j\rangle = \delta_{ij}$, encodes perfect distinguishability by a path measurement. A general pure preparation has two complex amplitudes,

$$|\psi\rangle = a|0\rangle + b|1\rangle,$$

with $|a|^2 + |b|^2 = 1$. The probability rule $P(0) = |a|^2$ and its partner for path 1 are physical postulates. Linear algebra lets us derive testable consequences of them; it does not derive probability from arithmetic alone.

Take preparation A with $a = b = 1/\sqrt{2}$. For preparation B choose either path with ordinary probabilities one half using a classical random switch. A detector placed directly in either arm records equal count rates for A and B. One shared probability distribution is therefore insufficient to identify the state. To discriminate the preparations, bring the paths together at a lossless recombiner. Choose its phase references so an incoming $|0\rangle$ emerges as $(|0\rangle + |1\rangle)/\sqrt{2}$, while $|1\rangle$ emerges as $(|0\rangle - |1\rangle)/\sqrt{2}$. The matrix $H = 2^{-1/2}\left(\begin{smallmatrix} 1 & 1 \\ 1 & -1 \end{smallmatrix}\right)$ is unitary because its columns are orthonormal. Other beam-splitter conventions put factors of i on reflection; a consistent set of propagation phases yields the same observed fringes.

Insert a relative phase ϕ on the lower path in preparation A. After recombination the amplitudes are $(1 + e^{i\phi})/2$ and $(1 - e^{i\phi})/2$. Multiplying each by its complex conjugate gives

$$P(0) = \frac{1 + \cos\phi}{2},$$

$$P(1) = \frac{1 - \cos\phi}{2}.$$

For preparation B each possible input produces equal exit probabilities. Averaging those probabilities leaves a constant one half at both ports for every phase. The contrast term results from multiplying the sum of amplitudes by its conjugate. Summing probabilities over alternatives is justified when the alternatives correspond to distinct preparations, or when a physical record distinguishes the two paths and that record is ignored. Confusing these operations predicts a fringe in the random-switch experiment. Counting both outputs checks normalisation independently of visibility, provided detector efficiencies have been calibrated.

At $\phi = \pi/3$, 400 independently prepared photons from A have expected output-0 count 300 and binomial standard deviation $\sqrt{400(3/4)(1/4)} = 8.66$. Preparation B predicts 200 with standard deviation 10. Repeating the measurement at $-\pi/3$ gives exactly the same probabilities: an analyser sensitive only to the cosine cannot determine the sign of the phase. Add a known $\pi/2$ offset before one further run. It converts a

sine into a cosine and removes that ambiguity. This simple experiment is a model of tomography, the recovery of state parameters from several distinct measurement settings.

Remove the phase that no ideal analyser can observe

Let a downstream analyser transmit normalised state $|v\rangle$. Born's rule gives transmission probability $|\langle v|\psi\rangle|^2$. Replacing the entire preparation by $e^{i\gamma}|\psi\rangle$ multiplies the overlap by a phase but leaves its squared modulus invariant for every v. Physical pure states are therefore rays. Relative phases survive because changing the phase of only one term is not multiplication of the whole vector by the same number. A basis vector itself may be assigned an arbitrary phase if its coordinates and every operator matrix are transformed consistently. Such a convention change cannot affect the phase measured relative to an actually inserted optical delay.

A unit vector in two complex dimensions has four real components. Normalisation removes one degree of freedom and the common phase removes another. Away from the coordinate poles a convenient parametrisation is

$$|\psi(\theta,\varphi)\rangle = \cos(\theta/2)|0\rangle + e^{i\varphi}\sin(\theta/2)|1\rangle, \qquad 0 \le \theta \le \pi.$$

Projecting onto $|0\rangle$ gives $(1+\cos\theta)/2$. The corresponding point on the unit sphere has coordinates $\mathbf{r} = (\sin\theta\cos\varphi, \sin\theta\sin\varphi, \cos\theta)$. Orthogonal states sit at opposite ends of a diameter. Their resulting angle in ordinary three-space is π, while their angle defined from the Hilbert-space overlap is $\pi/2$. This factor of two is why the spinor parametrisation contains half angles. At the poles φ drops out of all probabilities. Attempting to estimate an azimuth of an exactly path-localised preparation is a singular parametrisation, rather than a failure of the experiment.

Represent the analyser for the pure direction $\mathbf{n}$ by the projector $P_{\mathbf{n}} = |v\rangle\langle v|$. Projectors satisfy $P^\dagger = P$ and $P^2 = P$. Let $\sigma_x, \sigma_y, \sigma_z$ be the Pauli matrices with products $\sigma_i\sigma_j = \delta_{ij}I + i\epsilon_{ijk}\sigma_k$. Then $P_{\mathbf{n}} = (I + \mathbf{n}\cdot\boldsymbol{\sigma})/2$. Since each Pauli matrix has zero trace and $\operatorname{tr}\sigma_i\sigma_j = 2\delta_{ij}$,

$$\operatorname{tr}(P_{\mathbf{r}}P_{\mathbf{n}}) = \frac{1+\mathbf{r}\cdot\mathbf{n}}{2} = \cos^2(\gamma/2),$$

where γ is the spatial angle between the two directions. Thus transmission through a rotated analyser determines an inner product between real geometric directions, even though the original amplitude was complex. The vector for a path interferometer describes a two-dimensional mode space; rotating it need not mean physically rotating the entire apparatus in the laboratory.

Differentiate a probability from the state after a measurement

For an ideal projective measurement $\{P_j\}$, mutually orthogonal projectors sum to the identity. The probability of result j in a pure state is $\langle\psi|P_j|\psi\rangle$. Conditional on actually observing that result, the state for subsequent predictions in an ideal repeatable experiment is $P_j|\psi\rangle/\sqrt{p_j}$, provided $p_j > 0$. Applying the same measurement immediately gives the same result with probability one, because $P_j^2 = P_j$. Measuring a different observable in between can change the answer. The formula for an ideal repeatable measurement is an assumption about the measuring device as well as about the observable. An absorbing photon counter

ends the photon's trajectory, so a sequential model applies to a nondestructive analyser or to a filtered ensemble prepared in the transmitted state.

Prepare $|0\rangle$ and then sequentially measure a direction tilted 60° from $+z$, followed by z. If one keeps only the positive tilted outcome, three quarters of the original ensemble survive the selection. Among those survivors the fraction with final $+z$ is again three quarters. The joint probability of both positives is $9/16$; reporting $3/4$ for this joint event would discard the probability cost of selection. If all outcomes of the intermediate measurement are retained but unread, half is no longer the right default either. One must add over both paths of recorded outcomes. The final $+z$ probability is

$$(3/4)^2 + (1/4)^2 = 5/8.$$

Without the intermediate device it would be one. Removing data after a measurement and removing the measurement itself are experimentally different operations.

For a broad outcome corresponding to several basis states, a projector can preserve coherence inside the accepted subspace. Suppose a three-mode preparation is $(|0\rangle + |1\rangle + |2\rangle)/\sqrt{3}$. A device that only asks whether the state is in the span of $0, 1$ accepts with probability $2/3$ and ideally transmits $(|0\rangle + |1\rangle)/\sqrt{2}$. A device measuring the individual paths and later combining their labels into “accepted” yields the same acceptance probability but no surviving phase relation between 0 and 1 unless its record is coherently erased. A recombiner acting on the accepted ensemble distinguishes the implementations. An observable's probability table by itself does not specify every disturbance produced by a physical instrument.

Recover a state matrix from three families of counts

For a classical mixture of pure preparations $|\psi_j\rangle$ with weights w_j, average their individual probabilities. The result is $\mathrm{tr}(\rho P)$ with $\rho = \sum_j w_j |\psi_j\rangle\langle\psi_j|$. This matrix has trace one, is Hermitian and has nonnegative expectation on every vector. Changing the ensemble decomposition without changing ρ makes no difference to any measurement performed on the system alone. Preparation B is $I/2$. A 50:50 mixture of the two opposite x-analysed states is the same matrix. Trying to infer which randomisation mechanism was used requires access to its preparation record.

Every Hermitian two-state density matrix of trace one has the form $\rho = (I + \mathbf{r}\cdot\boldsymbol{\sigma})/2$. Diagonalise along $\mathbf{r}$. The eigenvalues are $(1 \pm |\mathbf{r}|)/2$. Positivity is therefore equivalent to $|\mathbf{r}| \le 1$. Interior points correspond to mixtures and surface points to pure states. Measuring σ_i gives $p_i^+ = (1 + r_i)/2$, so three noncoplanar analyser directions suffice to reconstruct an otherwise unknown two-level state given enough identically prepared copies. Noncommutation places no obstacle to using independent copies for these settings.

For 1000 trials per direction take $n_x^+ = 740,\ n_y^+ = 320,\ n_z^+ = 750$. The linear estimates are

$$r_x = 0.48,\ r_y = -0.36,\ r_z = 0.50.$$

Their length is

$$\sqrt{0.61} = 0.7810.$$

The eigenvalues are 0.8905 and 0.1095. The probability for $\mathbf{n} = (3/5, 4/5, 0)$ is $1/2$ even though the state has strong coherence in the path basis. The setting is accidentally perpendicular to its transverse projection.

A fourth such measurement supplies an out-of-sample check of state preparation and analyser calibration. If finite data instead give $|\mathbf{r}| > 1$, a negative estimated eigenvalue signals statistical inconsistency with exact physical probabilities. Fitting all binomial data under $|\mathbf{r}| \leq 1$ retains a physical model. Clipping each Cartesian expectation separately fails to enforce the spherical constraint.

Recover a phase from a closed loop of projections

There is phase information that cannot be associated with any one point on the sphere. For three nonorthogonal preparations A, B, C, the complex product $G = \langle A|B\rangle\langle B|C\rangle\langle C|A\rangle$ is unchanged by independent phase changes of the three state vectors. The phase acquired in each overlap cancels from the complete product. A loop of conditional projections can therefore have an observable phase relative to a reference arm even if every intermediate selection probability depends only on a modulus. The probability of successfully traversing the filtered path also matters for the attainable interference contrast.

Take $A = |0\rangle,\ B = (|0\rangle + |1\rangle)/\sqrt{2},\ C = (|0\rangle + i|1\rangle)/\sqrt{2}$. The overlaps are $1/\sqrt{2}, (1+i)/2, 1/\sqrt{2}$. Their product is $(1+i)/4$, giving phase $+\pi/4$. Traversing the reversed overlap product gives its complex conjugate and changes the sign. The directed path in a laboratory must be matched carefully to the order of bra and ket in the adopted product. If the analysers actually project a state in the order $A \to B \to C \to A$, the amplitude in the reference convention is $\langle A|C\rangle\langle C|B\rangle\langle B|A\rangle = G^*$. A path with a zero overlap has zero surviving interference amplitude and an undefined phase. This condition prevents interpreting a phase jump at a dark output as a measured singular rotation of a physical state.

Design the best analyser for discriminating two known signals

Imagine a source chosen at random from two calibrated preparations ρ_0, ρ_1 with known prior probabilities η_0, η_1. One copy arrives; the task is to guess which preparation was used and minimise the probability of a wrong guess. This differs from estimating all unknown entries of a state matrix from many repetitions. For example, with highly unbalanced priors the best choice may be to guess the more common preparation after many different measured outcomes. A measurement along the direction connecting the two Bloch vectors is optimal for equal priors under the criterion derived below, but generally need not be a universal optimal analyser for other decisions.

Let a binary measurement have two positive effects $E_0, E_1 = I - E_0$, where the eigenvalues of E_0 lie between zero and one. These positive-operator-valued effects give probabilities $\mathrm{tr}(\rho E_j)$, extending projectors to include randomisation or coarse responses. Adopt the decision rule of reporting label j after measurement result j. The average probability of success is then

$$P_{\mathrm{good}} = \eta_0 \,\mathrm{tr}(E_0\rho_0) + \eta_1 \,\mathrm{tr}[(I - E_0)\rho_1] = \eta_1 + \mathrm{tr}(E_0\Delta),$$

$$\Delta = \eta_0\rho_0 - \eta_1\rho_1.$$

Diagonalise the Hermitian Δ. The trace depends only on diagonal elements of E_0 in this basis. Each lies in $[0, 1]$. To maximise the sum, take one for positive eigenvalues and zero for negative ones. The resulting projector is an allowed measurement, proving optimality rather than only providing an upper bound. Zero eigenvalues permit arbitrary treatment without changing average success. Since the sum of positive eigenvalues equals $[\|\Delta\|_1 + \mathrm{tr}\,\Delta]/2$, substitution gives

$$P_{\mathrm{good}}^{\max} = (1 + \|\Delta\|_1)/2.$$

The trace norm here is the sum of singular values; for this Hermitian matrix the singular values are absolute eigenvalues. This operational meaning explains why trace distance between preparations appears repeatedly when comparing quantum processes.

For two equally likely pure states of overlap magnitude $c = |\langle\psi_0|\psi_1\rangle|$, use the two-dimensional span of the states. The eigenvalues of $\Delta = (\rho_0 - \rho_1)/2$ are $\pm\sqrt{1-c^2}/2$. Consequently $P_{\text{good}}^{\max} = [1 + \sqrt{1-c^2}]/2$. Orthogonal states can be perfectly distinguished. Identical preparations attain success one half, exactly the prior random-guess limit. For a concrete optical example let preparation 0 be $|0\rangle$ and preparation 1 be $\cos 30^\circ|0\rangle + \sin 30^\circ|1\rangle$. Their Hilbert-space overlap is $\sqrt{3}/2$, while the Bloch vectors have a spatial separation angle of 60°. The best success probability is $3/4$. Measuring the path basis instead and reporting 0 whenever it transmits yields success $[1 + \sin^2 30^\circ]/2 = 0.625$. The best analyser deliberately allows some errors on both sources to reduce the average error over the pair.

The physical realisation is simple in this example. The difference of Bloch vectors is $(-\sqrt{3}/2, 0, 1/2)$, already of unit length. Use that direction for the positive projector labelled 0. Its dot product with the preparation-0 vector is $+1/2$, and with the preparation-1 vector is $-1/2$, giving correct guess probability $3/4$ for either source independently. The sign of this vector matters because it assigns the outcome label. Reversing the analyser axis and interchanging the decision labels leaves the same performance. An experimental confusion matrix recording each conditional outcome rate checks considerably more than an aggregate fraction of successful guesses, particularly if the priors in the actual source differ from those in the calibration trial.

If two independent copies of the same unknown signal can be stored and measured jointly, the pure-state overlap becomes c^2. The optimum success rises to $[1 + \sqrt{1-c^4}]/2$. For the previous optical example this is about 0.831. The calculation assumes both copies were identically prepared from one shared unknown label. It would be invalid if each copy came from a newly and independently drawn source label. Such a distinction arises naturally when comparing a two-photon coincidence experiment with two unrelated records in a database. Increasing the number of independent preparations for tomography and increasing copies available per individual discrimination decision are different resources.

Calibrate an imperfect path detector with its classical record intact

An ideal path measurement perfectly distinguishes the two modes and destroys their mutual coherence in an unread outcome. A tunable weaker measurement provides an intermediate experiment linking information and disturbance quantitatively. Suppose an instrument reports $+$ or $-$ with operators $M_+ = \mathrm{diag}(\sqrt{(1+d)/2}, \sqrt{(1-d)/2})$ and $M_- = \mathrm{diag}(\sqrt{(1-d)/2}, \sqrt{(1+d)/2})$, in the path basis, with $0 \le d \le 1$. Squaring and summing gives identity, so probabilities normalise for every input. Applied separately to the known 0 and 1 calibration sources, the $+$ result occurs with probabilities $(1+d)/2$ and $(1-d)/2$; their difference measures d. At zero this detector returns an independent coin flip; at one it reads the path without misclassification.

For an input with Bloch component z, the result probability is

$$p_+ = (1 + dz)/2.$$

Conditional on that result, the path-0 and path-1 weights must be renormalised by p_+. The new longitudinal component and transverse coherence satisfy

$$z' = \frac{z + d}{1 + dz}, \qquad (\rho_{01})' = \frac{\sqrt{1 - d^2}}{1 + dz}\,\rho_{01}.$$

If the outcome is unread, the map is the sum of both $M_j \rho M_j^\dagger$. The populations are unchanged and the off-diagonal entry becomes $\sqrt{1 - d^2}\rho_{01}$. A balanced pure input consequently gives a recombiner fringe visibility $\sqrt{1 - d^2}$. This equality links this calibrated detector's ability to distinguish path preparations with the contrast left in a complementary measurement. It depends on the chosen instrument being a minimal diagonal measurement. A device with the same two outcome probabilities can also apply an uncontrolled relative phase, yielding a smaller visibility. The existence of such instruments is why probability effects alone do not determine post-measurement states.

Choose $d = 0.60$ and a balanced pure $+x$ input. Either outcome has probability one half. Given $+$, the new state has $z' = 0.60$ and $x' = 0.80$, so it is pure. The other branch has $z' = -0.60, x' = 0.80$. If the outcomes are ignored their average has $\mathbf{r} = (0.80, 0, 0)$, which is mixed. Feeding each branch through a separate calibrated rotation can align both with the original $+x$ direction, restoring full visibility when the input is known. This feed-forward preparation does not invert arbitrary unknown states through a closed deterministic quantum map; it makes use of the recorded result and knowledge of the preparation. A single fixed rotation cannot increase the purity of the unread mixture because unitaries preserve the eigenvalues of ρ.

There is also a practically important calibration ambiguity. Suppose a perfect projective path detector is followed by a classical reporting error which flips the stored label with probability $(1 - d)/2$. Its outcome rates on every input agree exactly with those of the weak instrument above. But its unrevealed physical interaction has already produced orthogonal environmental records of the true path; unread coherence is then zero for every value of the classical label reliability. Test the distinction by preparing a known balanced superposition, measuring with the candidate instrument and recombining its exits while retaining both count channels. Matching path-classification accuracy cannot substitute for this separate disturbance test. Conversely observing reduced visibility alone does not reveal whether useful path information is accessible in the recorded labels.

Count statistics when the analyser itself has a bias

Suppose the final binary detector has efficiency ϵ_0 at port 0 and ϵ_1 at port 1. If lost events are discarded and only a ratio of recorded clicks retained, the observed fraction is $p_{\mathrm{obs}} = \epsilon_0 p / [\epsilon_0 p + \epsilon_1 (1 - p)]$. At a balanced state unequal efficiencies can imitate a longitudinal preparation bias. For $p = 1/2, \epsilon_0 = 0.90, \epsilon_1 = 0.75$, after selection the recorded port-0 fraction is

$$0.90/1.65 = 0.5455.$$

Treating it as an ideal probability would infer $z = 0.091$. Collecting arbitrarily many photons reduces the scatter about that biased result without reducing the bias itself.

A measured ratio from an independently characterised balanced reference can determine relative efficiency $r = \epsilon_0/\epsilon_1$. In odds form $p_{\mathrm{obs}}/(1 - p_{\mathrm{obs}}) = r\, p/(1 - p)$. If the reference is unavailable but a known unitary exchanges the measurement ports, record a second measurement on the same preparation with that exchange. Its observed odds are $r(1 - p)/p$. Their product is r^2; their ratio is $[p/(1 - p)]^2$. Thus the two

runs can separate relative gain and true probability, provided the gain and preparation remain stable across them. For observed odds 1.8 and 0.8 one obtains $r = 1.2$ and true odds 1.5, hence $p = 0.60$. Stability is testable by alternating the two settings across the acquisition instead of making one long run after the other. This class of self-calibrating measurement is useful far beyond a two-path interferometer; the proposed physical symmetry is built into both the inversion formula and its independent checks.

Reconstruct a state and check positivity

Change the recorded counts along two spin axes. Determine whether the inferred state can predict valid probabilities for every fourth analyser setting.

Open this experiment online

Try the calculation

Up counts along y (out of 25) **19**

Up counts along z (out of 25) **19**

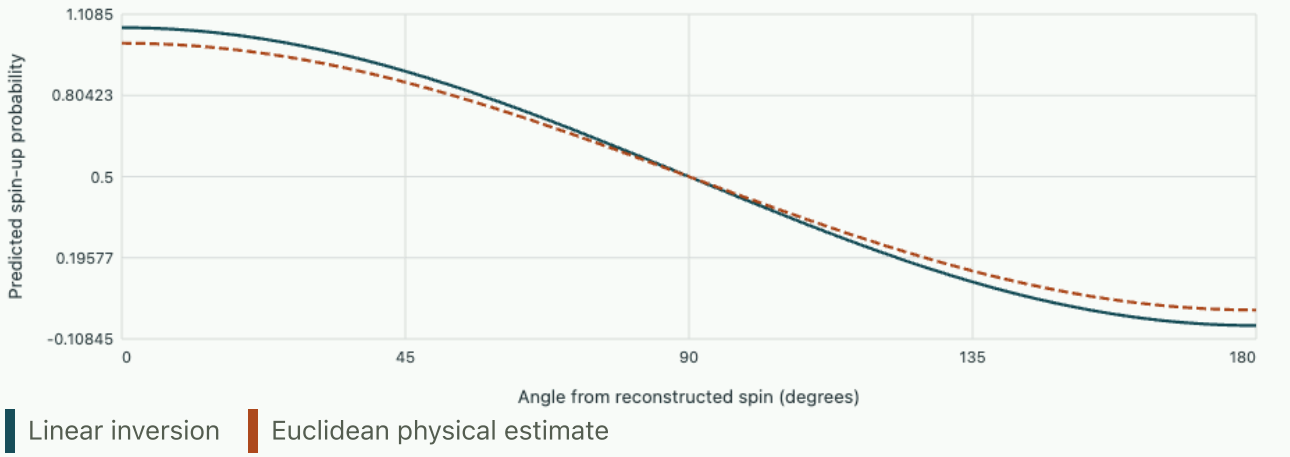

Reconstructed Bloch radius 1.1164. Eigenvalues 1.0582 and -0.058211.

Separate component uncertainty from experimental readout resolution

A variance bound can be proved directly from Hilbert-space geometry without assuming a model of a particular measurement apparatus. For a normalised pure state set $|u\rangle = (A - \langle A\rangle)|\psi\rangle$, $|v\rangle = (B - \langle B\rangle)|\psi\rangle$, for observables and states for which the required products exist. The squared norms are $(\Delta A)^2, (\Delta B)^2$. Cauchy–Schwarz bounds their product below by $|\langle u|v\rangle|^2$. Separating the real and imaginary parts of the inner product gives

$$(\Delta A)^2(\Delta B)^2 \geq \left|\frac{\langle [A,B]\rangle}{2i}\right|^2 + \left|\frac{1}{2}\langle \{\Delta A_{\mathrm{op}}, \Delta B_{\mathrm{op}}\}\rangle\right|^2,$$

where $\Delta A_{\mathrm{op}} = A - \langle A \rangle I$ denotes a centred operator rather than its scalar standard deviation. The real term is the squared symmetrised covariance. The commonly quoted commutator-only inequality drops this nonnegative contribution. The bound describes spreads across identical preparations; the electronics' error in reading a single outcome is a different quantity and enters an experimental analysis through the detection model.

For Pauli measurements $A = \sigma_x, B = \sigma_y$, the elementary lower bound on the product of variances is r_z^2, since $[\sigma_x, \sigma_y] = 2i\sigma_z$. The full covariance is $-r_x r_y$ because their anticommutator vanishes before centring. The actual product is $(1 - r_x^2)(1 - r_y^2)$. Subtracting the full lower bound leaves $1 - r_x^2 - r_y^2 - r_z^2$, directly testing the Hilbert-space positivity constraint on the Bloch vector. Every pure qubit saturates the full inequality for this pair, while a mixed interior state can give strict excess. A zero commutator expectation in a particular state does not prove that the two observables commute or can be jointly assigned precise values there. For $|+x\rangle$ one variance is zero, while the other is maximal.

Decide whether an apparent uncertainty violation is an instrument problem

A device reconstruction reports $r_x = 0.80,\ r_y = 0.60,\ r_z = 0.20$. Test the commutator-only bound, the full covariance inequality and the eigenvalue positivity condition. Why is testing the weaker bound alone insufficient?

▼ **Derive and check the result**

The measured components imply variances 0.36 and 0.64, with product 0.2304. The commutator-only bound is 0.04 and passes. The covariance-squared contribution is $0.8^2\,0.6^2 = 0.2304$, so the full bound is 0.2704 and fails by 0.0400. Independently $|\mathbf{r}|^2 = 1.04$ predicts a slightly negative eigenvalue. Finite count noise or inappropriate calibration can make separate empirical component estimates leave the physical Bloch ball even though every individual estimate lies in its allowed range. Check uncertainty regions jointly and, when fitting a physical state, retain measurement likelihoods and positivity. Satisfying one weaker inequality cannot certify that all estimated outcomes come from a common quantum preparation.

Tutorial · use interference, sequential selection and tomography

Infer visibility and a path-coherence phase

An equal-population path state gives normalised output-0 probabilities 0.80 with zero phase offset and 0.70 after applying $+\pi/2$ to path 1. With the recombiner convention above, find the complex density matrix entry ρ_{01}, its allowed magnitude, and the highest achievable output probability.

▼ **Work through the solution**

The phase plate sends ρ_{01} to $e^{-i\phi}\rho_{01}$, so $p_0 = 1/2 + \mathrm{Re}(e^{-i\phi}\rho_{01})$. The observations give real part 0.30 and imaginary part 0.20. Its magnitude is 0.3606, below the positivity bound $\sqrt{\rho_{00}\rho_{11}} = 0.5$. The maximum exit probability is 0.8606, obtained when the phase offset cancels the argument of ρ_{01}; here $\phi = \arctan(2/3) = 33.7^\circ$. The minimum is 0.1394. The usual fringe visibility $(p_{\max} - p_{\min})/(p_{\max} + p_{\min})$ is 0.721. The shortfall from unity in ideal equal-population data measures reduced path coherence. Unequal

detection efficiencies or an unbalanced recombiner would need an independent calibration before drawing that conclusion from raw counts.

Is a succession of nearly parallel filters an efficient rotator?

Start in $+z$. Place N ideal positive-channel filters whose directions rotate in equal steps from $+z$ to $-z$ along a great circle. Find the transmitted fraction and evaluate it for $N = 2, 10$. Compare with applying a single final filter.

▼ **Work through the solution**

Adjacent directions differ by π/N, so each conditional transmission probability is $\cos^2(\pi/2N)$. All accepted states end at $-z$, with joint transmission $[\cos^2(\pi/2N)]^N$. For one filter overlap is zero. Two filters give $1/4$, whereas ten give about 0.7805. Taking logarithms gives $2N \ln\cos(\pi/2N) = -\pi^2/(4N) + O(N^{-3})$, tending to zero, so the ideal transmission tends to one. Each retained state has been prepared along the new axis by the filtering. Finite loss per filter eventually reverses the trend: multiplying by an efficiency 0.99^N penalises arbitrarily large N. Tracking both conditional and joint rates distinguishes this sequence from a unitary spin rotation.

Set an error bar on a prediction along a new direction

Use the three tomography runs above to predict p_+ for $\mathbf{n} = (3/5, 4/5, 0)$ with an approximate standard error. Test whether observing 535 positives in 1000 new trials is surprising at the level of a few standard deviations.

▼ **Work through the solution**

For a component estimate $r_i = 2n_i^+/N_i - 1$, the plug-in variance is $(1 - r_i^2)/N_i$ for independent binomial trials. Independent settings give $\mathrm{var}(\hat{p}) = \frac{1}{4}[0.36(1 - 0.48^2) + 0.64(1 - 0.36^2)]/1000 = 0.0002085$, or standard error 0.0144. The new fraction has sampling error about 0.0158. Comparing two independent estimates requires adding these variances, giving combined standard error about 0.0214. The discrepancy 0.035 is about 1.6 standard errors. Treating the reconstructed state as exactly known would overstate the evidence of miscalibration. For operational work near the Bloch-sphere boundary, a likelihood analysis handles correlations and positivity more accurately than this linear interior approximation.

Further reading

- Fließbach · Abstrakte Formulierung, Hilbert-space states and observables
- Cambridge · quantum foundations, density matrices and generalised measurements
- Course investigation · spin measurement and joint preparations

Quantum states, dynamics and measurement · 2. What a momentum eigenstate means in a finite experiment

A detector never counts a perfectly sharp momentum eigenket. Derive continuum normalisation from a finite box, change the spectral measure from momentum to energy and predict the full evolution of a normalisable Gaussian packet.

Follow a normalised state as its allowed box grows

Consider a particle moving in one spatial dimension. A periodic box of length L is a mathematical model for an observation made far from any reflecting wall. The inner product is the integral of $\psi^*(x)\chi(x)$ over the box. The momentum operator is $-i\hbar\, d/dx$ with periodic boundary conditions. Integration by parts cancels its boundary form for periodic differentiable functions. Its normalised modes $L^{-1/2}e^{ik_n x}$ have $k_n = 2\pi n/L$ and momenta $p_n = \hbar k_n$. Their overlap is one if the integers coincide and zero otherwise. Increasing L reduces the spacing $\Delta p = 2\pi\hbar/L$ but also reduces the amplitude of each individual box mode at a fixed position. If one simply replaces the Kronecker delta by a Dirac delta while carrying over $L^{-1/2}$, the units and probabilities will fail.

Let c_n be the unit-normalised amplitudes with $\sum_n |c_n|^2 = 1$. For a smooth normalisable packet the probability in an interval dp has a limit $|\phi(p)|^2 dp$. Thus $c_n \simeq \sqrt{\Delta p}\,\phi(p_n)$. Substituting in the synthesis sum and using $L\Delta p = 2\pi\hbar$ produces

$$\psi(x) = \frac{1}{\sqrt{2\pi\hbar}} \int_{-\infty}^{\infty} e^{ipx/\hbar}\phi(p)\, dp,$$

$$\phi(p) = \frac{1}{\sqrt{2\pi\hbar}} \int_{-\infty}^{\infty} e^{-ipx/\hbar}\psi(x)\, dx.$$

The distributions $\langle x|p\rangle = (2\pi\hbar)^{-1/2}e^{ipx/\hbar}$ have $\langle p|p'\rangle = \delta(p-p')$. The delta acts under integration; $\delta(0)$ has no finite interpretation as the probability for a detector. A physical probability is $\int_{p_1}^{p_2} |\phi(p)|^2 dp$. Changing to wave number gives

$$|\phi_k(k)|^2 = \hbar|\phi(\hbar k)|^2.$$

A graph labelled density in k with a vertical axis copied directly from the momentum plot loses the Jacobian and in general ceases to normalise to one.

Change from momentum bins to energy bins

For a free particle $E = p^2/(2m)$. Each positive energy has two distinct momenta $p_\pm = \pm\sqrt{2mE}$. A detector measuring energy but not direction sums over these mutually orthogonal alternatives. Apply the ordinary change-of-variables rule on each monotone half-line:

$$w(E) = \frac{m}{\sqrt{2mE}}\Big[|\phi(+\sqrt{2mE})|^2 + |\phi(-\sqrt{2mE})|^2\Big], \qquad E > 0.$$

The factor $m/|p|$ is $1/|dE/dp|$, with units of momentum per energy. Measuring a finite energy interval around zero can have a density diverging as $E^{-1/2}$ if $\phi(0) \neq 0$, yet the integrated probability is finite because $\int_0^\delta E^{-1/2} dE = 2\sqrt{\delta}$. The same effect appears in a box as relatively dense low-energy modes. Standing-wave box boundary conditions change the counting of travelling directions; include degeneracy once and use the appropriate mode spacing, rather than doubling both.

On an infinite line it is useful to introduce $|E, +\rangle = \sqrt{m/|p|}\,|p\rangle$ at $p > 0$, and a corresponding negative branch. In the overlap,

$$\delta(p(E) - p(E')) = \delta(E - E')/|dp/dE|;$$

the two square-root factors give $\delta(E - E')$. The resolution of identity contains $\sum_{\pm} \int_0^\infty dE\, |E, \pm\rangle\langle E, \pm|$. Measuring energy while coherently manipulating the two degenerate directions requires keeping both amplitudes. Adding their probabilities is sufficient only for an instrument whose outcome identifies energy without mixing the directions before detection.

Integrate the free evolution of a Gaussian

Prepare a minimum-uncertainty wave packet centred at x_0 with average momentum p_0 and initial position standard deviation s:

$$\psi(x, 0) = (2\pi s^2)^{-1/4} \exp[-(x - x_0)^2/(4s^2) + i p_0 (x - x_0)/\hbar].$$

Its squared modulus is a Gaussian with variance s^2, which explains the factor four in the amplitude exponent. Complete the square in the Fourier integral to find a momentum probability Gaussian centred at p_0 of variance $\hbar^2/(4s^2)$. The initial state has zero symmetrised position–momentum covariance. Multiplying by a chirp factor $\exp[i\kappa(x - x_0)^2]$ leaves its initial position density unchanged but changes momentum variance and covariance. A snapshot of one density does not determine its later free spread.

Each momentum component evolves by $e^{-ip^2 t/(2m\hbar)}$. Split $p = p_0 + q$. The constant term in p^2 produces a common phase; the linear term shifts the envelope by $p_0 t/m$; the quadratic term changes the width. Carrying out the remaining Gaussian integral yields

$$|\psi(x, t)|^2 = \frac{\exp[-(x - x_0 - p_0 t/m)^2/(2 s_t^2)]}{\sqrt{2\pi}\, s_t},$$

$$s_t^2 = s^2 \left[1 + \left(\frac{\hbar t}{2 m s^2}\right)^2\right].$$

An independent route uses the Heisenberg operator $x(t) = x(0) + p(0) t/m$. Its variance is $s^2 + (t/m)\langle\{\Delta x, \Delta p\}\rangle + (t/m)^2 (\Delta p)^2$. Inserting zero covariance and the Gaussian momentum width reproduces s_t^2. For a chirped packet the covariance term can be negative; momentary narrowing during free evolution is consistent with uncertainty because momentum width and the phase correlation must be included. The density integral stays one, so spreading lowers the height of the peak without removing particles.

Take a model atom of mass 1.44×10^{-25} kg prepared with $s = 1.00\,\mu\text{m}$. Its spread time $t_s = 2ms^2/\hbar = 2.73$ ms. At $t = t_s$ the standard deviation is 1.41 µm. With mean velocity 2.00 mm/s its mean

has advanced 5.46 µm. At long flight times $s_t \simeq t\Delta p/m$. Measuring position after a known long flight therefore maps the momentum distribution for this free packet. At $0.5t_s$ the neglected initial position variance is still dominant, so the same inference would substantially overestimate Δp.

A detector window and a Fourier grid do different jobs

Suppose the position-sensitive detector has a Gaussian point-spread function of width d. The measured density is a convolution of the true density with that response if detections are independent and gain is uniform. Gaussian variances then add:

$$s_{\mathrm{obs}}^2 = s_t^2 + d^2.$$

This broadening belongs to the observation model; it does not increase the particle's momentum spread before the measurement. For $d = 1.20\,\mu\mathrm{m}$ at $t = t_s$, the observed standard deviation becomes $1.85\,\mu\mathrm{m}$. A fit to a single image with unknown d cannot disentangle the two contributions. A sequence of flight times constrains the coefficient of t^2 independently of a fixed instrument width.

A discrete Fourier propagation represents a periodic interval sampled at spacing Δx. Its wave-number spacing is $2\pi/L$, and the highest resolved wave number is order $\pi/\Delta x$. Enlarging L at fixed Δx samples the momentum profile more finely and delays periodic wraparound. Reducing Δx at fixed L raises the bandwidth without moving the periodic boundary. Either operation alone can leave the other error intact. A norm-preserving FFT propagator can therefore return a perfectly normalised wave packet that has wrapped through an unphysical boundary or whose carrier momentum has aliased. Compare the phase, group velocity and width against independent formulas, then refine the two scales separately.

Why boundary conditions are part of a momentum observable

The expression $p = -i\hbar\, d/dx$ does not by itself specify a complete operator. On a finite interval $[0, L]$, integrate two sufficiently smooth square-integrable functions by parts:

$$\langle f, pg\rangle - \langle pf, g\rangle = -i\hbar[f^*(x)g(x)]_0^L.$$

A symmetric domain makes this boundary form vanish for every pair in the domain. A self-adjoint operator additionally has the same domain as its adjoint. This distinction matters for a physical generator of unitary translations and for assigning sharp spectral outcomes. Imposing a real or imaginary matrix representation in a finite basis can conceal boundary conditions that become essential as the basis is enlarged.

For a ring allow quasi-periodicity $g(L) = e^{i\theta}g(0)$ with fixed θ common to all functions in the domain and square-integrable first derivative. Then the boundary form vanishes. Conversely, for an adjoint candidate f the requirement that it vanish for all possible $g(0)$ forces $f(L) = e^{i\theta}f(0)$, so the same condition reappears. Solving $pg = pg$ yields the momenta $p_n = \hbar(2\pi n + \theta)/L$. Thus boxes with precisely the same local differential expression have momentum spectra shifted by an externally fixed phase. In a physically realised ring such a phase can encode a magnetic flux through a gauge-transformed boundary condition; a gauge consistent with that interpretation must also be used in the Hamiltonian and physical momentum.

Compare with forcing $g(0) = g(L) = 0$ as in a hard-wall well. The boundary form vanishes for all f, g satisfying both zero conditions, so the derivative is symmetric there. But the adjoint boundary form vanishes

for arbitrary endpoint values of f when g vanishes at both ends. Consequently this two-zero derivative domain is smaller than its adjoint domain. There is no ordinary rigid translation that carries every state of a hard-wall interval back into that same interval without assigning what happens at its boundary. Meanwhile the kinetic energy $-\hbar^2 d^2/(2m\,dx^2)$ with suitable Dirichlet domain is self-adjoint and has a well-defined discrete spectrum. The symbol $p^2/(2m)$ in a Dirichlet well does not turn its first derivative with the same two endpoint conditions into a self-adjoint sharp momentum observable.

A release-and-detect experiment makes this distinction concrete. Prepare the first Dirichlet state $\sqrt{2/L}\sin(\pi x/L)$ supported in $0<x<L$, then remove the walls suddenly and measure momentum on the full line. Its momentum wavefunction is the Fourier transform of the spatial function extended by zero outside. Multiplying the exponential and sine and integrating gives a continuous amplitude proportional to $[1+e^{-ipL/\hbar}]/[(\pi/L)^2-(p/\hbar)^2]$. The apparent singularities at $p/\hbar=\pm\pi/L$ are removable because the numerator also vanishes. It is a mistake to predict two delta peaks by replacing the standing wave with the sum of two infinite plane waves. The finite support multiplies that sum by a spatial window and broadens both momentum components.

The mean momentum for this real function is zero. Its second moment after release is $\hbar^2\int|\psi'|^2dx=\pi^2\hbar^2/L^2$, consistent with the original kinetic energy. Thus a continuous released momentum distribution can have an rms momentum whose square exactly recovers a discrete energy of the original bounded system. This is a useful quantitative example where the equality of moments across two measurement procedures does not imply equality of their full spectral probability distributions. Adding a short laboratory impulse to every particle after release should translate the whole distribution by the corresponding momentum kick, preserving its width. That response provides an independent check on a momentum detector's sign and scale.

Derive the free propagator with a convergent limiting prescription

Partitioning the free unitary into momentum eigenstates produces its position kernel

$$K(x,t;x',0)=\frac{1}{2\pi\hbar}\int_{-\infty}^{\infty}\exp\left[\frac{i}{\hbar}\left(p(x-x')-\frac{p^2t}{2m}\right)\right]dp.$$

It is oscillatory rather than absolutely integrable. Define it by first damping $e^{-\varepsilon p^2}$, evaluating a convergent Gaussian and then taking $\varepsilon\downarrow 0$ on smooth normalisable test packets. Completing the square in the exponent gives, for $t>0$,

$$K=\sqrt{\frac{m}{2\pi i\hbar t}}\exp\left[\frac{im(x-x')^2}{2\hbar t}\right].$$

The phase convention of the square root for positive time follows continuously from the regularised Gaussian. The units are inverse length so $\psi(x,t)=\int K(x,t;x',0)\psi(x',0)\,dx'$ has the same units as the initial wavefunction. As $t\to 0^+$, the kernel does not approach a finite spike pointwise. Its increasingly fast oscillations across x' cancel except near $x'=x$, recovering the initial state in the appropriate strong or distributional sense. For a constant initial amplitude one may recover unity formally, but such an amplitude on the full line is itself not a normalised test state.

Two checks involve no wave packet fitting. Inserting two kernels for t_1, t_2 and integrating over an intermediate position gives the kernel for $t_1 + t_2$, either by another Gaussian integration or, more directly, by composing the momentum phase factors. Complex conjugation and reversing the time arguments yield the inverse evolution. Neither statement holds for an arbitrary choice of magnitude prefactor assigned merely to obtain the correct phase. In Fourier numerics the corresponding checks are unit norm after repeated evolution and recovery of the original sampled complex state after a forward followed by backward propagation. These are necessary algorithm checks, even though a periodic or aliased computational model can still satisfy them exactly.

This kernel also clarifies the long-time momentum mapping. Factor $e^{imx^2/(2\hbar t)}$ from K and identify the Fourier transform of $e^{imx'^2/(2\hbar t)}\psi(x', 0)$ evaluated at $p = mx/t$. If the occupied initial region is narrow enough that $mx'^2/(2\hbar t)$ varies little across it about the selected origin, the extra phase can be neglected and

$$|\psi(x,t)|^2 \simeq \frac{m}{t}\left|\phi\Big(\frac{mx}{t}\Big)\right|^2.$$

The factor m/t is the Jacobian dp/dx. When the initial packet is displaced, use a suitable shifted origin or account for the induced linear phase explicitly. This derivation applies well beyond Gaussians and gives a scale for testing the approximation before interpreting a flight image as a momentum histogram. The quadratic residual phase provides the matter-wave analogue of the near-field correction in diffraction.

Distinguish a detector crossing from a snapshot in space

A detector counts particles passing through an aperture over time. Its signal is related to flux, whereas a camera image of identical preparations at one time probes a position distribution. For a one-dimensional free wavefunction the continuity equation is $\partial_t|\psi|^2 + \partial_x j = 0$ with $j = (\hbar/m)\,\mathrm{Im}(\psi^*\partial_x\psi)$. Integrating over the region $x < X$ gives the change in probability to the left of a fixed plane. If there is no flux at negative infinity, the net integrated current through X equals the reduction in that probability. A positive-reading counter with absorption is a more detailed instrument and cannot in general be identified with a signed current in the presence of reflected waves or interference.

For a minimum-uncertainty Gaussian initially at x_0, the exact phase in the freely expanding envelope has a quadratic term. Differentiating the phase gives the local velocity entering current,

$$\frac{j(x,t)}{|\psi(x,t)|^2} = \frac{p_0}{m} + \frac{t}{t_s^2 + t^2}\left(x - x_0 - \frac{p_0 t}{m}\right).$$

At the centre this equals the group velocity p_0/m, but ahead of the centre it is larger and behind it smaller. A single constant velocity multiplying density is inadequate for a significantly spreading packet. For the earlier atom at $t = t_s$, one current-standard-deviation distance ahead corresponds to $1.41\ \mu\mathrm{m}$ and the additional speed is about 0.259 mm/s. The trailing side gives the opposite correction. Direct differentiation of its Gaussian density verifies that this current satisfies continuity. A code reproducing s_t but evolving the wrong complex phase may fail this current test; judging an entire propagation by the changing width alone would miss the error.

A useful inverse question is whether two time-of-flight images suffice to recover an unknown complex wavefunction. In the long-time limit the image yields $|\phi(p)|^2$, which contains no direct momentum-dependent

phase. Translation of the initial wave packet changes $\phi(p)$ by a linear phase, leaving the momentum density unchanged. Complex conjugate or chirped preparations can also share particular intensity records. Additional calibrated evolution, known reference interference or structural assumptions can provide phase information. Even an exact free propagator only defines the forward map; its invertibility on complex amplitudes does not imply intensity-only data identify every possible input. The experiment must specify which class of preparations and which measurement settings render the desired parameters estimable.

Track a continuous superposition through focus

Change a position-momentum correlation and compare the full packet width with an initially unchirped packet. A narrow width at one time does not determine its subsequent spread.

Open this experiment online

Try the calculation

Initial phase curvature b -2

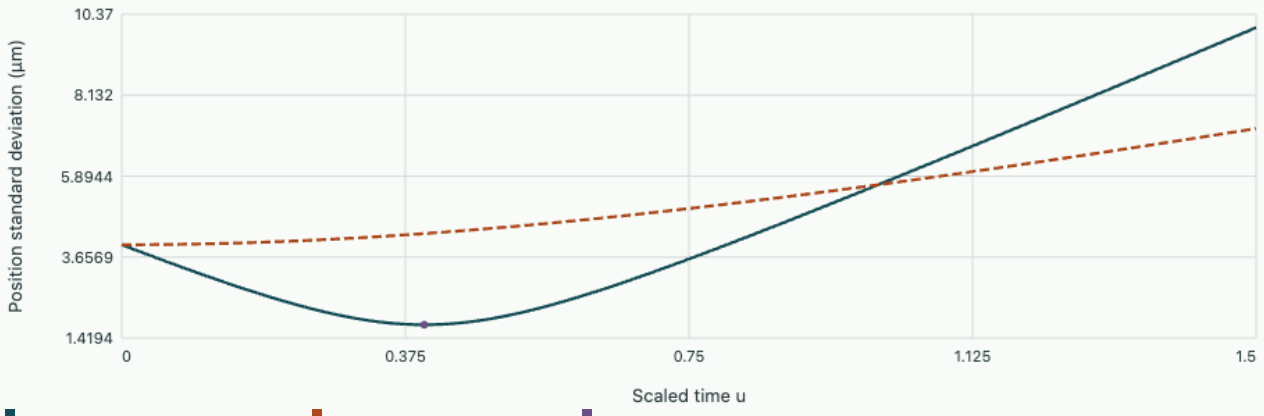

For u ≥ 0 the minimum width is 1.789 µm at u = 0.4.

Plan an FFT comparison using physical position and momentum resolution

A grid propagator can track the same packet analytically solved earlier, then carry over its verified conventions to less elementary potentials. A common trap is to regard the number of points N as a single accuracy parameter. Define a periodic interval of physical length $L = N\Delta x$. Discrete Fourier modes have $k_j = 2\pi j/L$, with negative wave numbers represented by negative indices after reordering. Their spacing is $\Delta k = 2\pi/L$; the largest distinct magnitude is close to $\pi/\Delta x$. A narrow structure in momentum requires large L. A high carrier momentum requires small Δx. These conditions can fail independently even when the norm stays exactly one.

If $\psi_l = \psi(x_l)$ is a sampled wavefunction in inverse square-root metres and the forward transform sums without $1/N$, the approximated continuum k-amplitude is

$$\phi_k(k_j) \simeq \Delta x \sum_l \psi_l e^{-ik_j x_l}/\sqrt{2\pi}.$$

The phase includes the physical coordinate origin adopted for the grid. Parseval then gives

$$\Delta k \sum_j |\phi_k(k_j)|^2 = \Delta x \sum_l |\psi_l|^2.$$

A library returning a unitary discrete transform instead assigns a factor $1/\sqrt{N}$ to both directions; multiplying its output by $\sqrt{\Delta x/\Delta k}$ converts sample units. Plotting a library output directly against physical momentum without these scale factors is especially misleading when comparing two meshes. Total discrete ℓ^2 norm and continuum probability density are different normalisations of the same samples.

For the atom of mass 1.44×10^{-25} kg, $s = 1.00$ µm and $v_0 = 2.00$ mm/s,

$$k_0 = mv_0/\hbar = 2.73\,\mu\mathrm{m}^{-1},$$

$$\Delta k_{\mathrm{rms}} = 1/(2s) = 0.50\,\mu\mathrm{m}^{-1}.$$

A box 80 µm wide gives spacing $0.0785\,\mu\mathrm{m}^{-1}$, about 6.4 modes per rms momentum width. With $N = 256$, grid spacing is 0.3125 µm and Nyquist wavenumber is $10.05\,\mu\mathrm{m}^{-1}$, exceeding $k_0 + 5\Delta k_{\mathrm{rms}} = 5.23\,\mu\mathrm{m}^{-1}$. This comfortably samples the initial carrier distribution. It is still necessary to check that the envelope remains inside the periodic box at the desired final time.

At $t = 10$ ms the packet centre has moved 20 µm and $s_t \simeq 3.80$ µm. Starting at the centre of an interval spanning -40 to 40 µm, its five-standard-deviation front approaches 39 µm and begins testing the boundary margin. At 20 ms the centre alone reaches the periodic boundary. A perfectly unitary calculation will show it reappearing on the opposite side. This is correct for a ring and incorrect for a free-space flight image. Enlarging N while keeping the same L improves the Nyquist limit but does nothing to postpone this crossing. An open-boundary absorber can delay reflected artefacts but attenuates norm; a lost norm in that model should be compared with counted escape rather than naively enforced to zero.

Infer displacement and speed from a wrapped packet without mistaking periodicity for reflection

A code on a periodic box of length 80 µm reports a sharp packet peak at $+35$ µm in one frame and -36 µm in the next frame 4.50 ms later. Find a consistent uniform drift and identify the ambiguity if no speed prior is available.

▼ Derive and check the result

The smallest forward displacement is $+9$ µm after adding one full box length to the second position. It corresponds to 2.00 mm/s and agrees with the earlier preparation. Direct subtraction without unwrapping would give -71 µm, or -15.8 mm/s, an apparent reversal. All displacements $9 + 80n$ µm for integer n yield the same wrapped pair, so the frames alone cannot identify the number of turns. A momentum measurement or a bound on speed relative to frame spacing resolves the branch. For an extended packet crossing a seam, use a periodic circular moment such as $\langle e^{2\pi i x/L} \rangle$ so an ordinary spatial mean does not spuriously jump

toward the box centre. If that circular moment is nearly zero, even its angular centre is ill-defined and a full density comparison is preferable.

Tutorial · design a momentum measurement

Distinguish an unresolved line from a finite-width packet

A periodic calculation of length $L = 100\,\mu\mathrm{m}$ contains mode n. A detector bins momentum in intervals $1.0 \times 10^{-29}\,\mathrm{kg\,m/s}$. How many box-mode spacings fit into a bin? What changes if the box length doubles at fixed physical packet?

▼ Work through the solution

Here $\Delta p = h/L = 6.626 \times 10^{-30}\,\mathrm{kg\,m/s}$, so the detector bin spans about 1.51 spacings. The precise number of grid points falling within a bin depends on its endpoints. Doubling L gives roughly 3.02 spacings per bin. Each corresponding coefficient squared decreases approximately in proportion to Δp, while the number included increases. The total probability should approach a stable integral on a progressively finer grid. If a graph compares $|c_n|^2$ across the boxes, divide by the corresponding Δp before calling the ordinate a continuum density.

Calculate a threshold energy interval without evaluating a delta at zero

A packet's momentum density is nearly constant w_0 over $-p_c < p < p_c$. Obtain the probability for $0 < E < E_c = p_c^2/(2m)$ by two different integrations. What changes when only positive-momentum particles reach the counter?

▼ Work through the solution

Direct integration in momentum gives $2p_c w_0$. In energy the two branches contribute $2mw_0/\sqrt{2mE}$. Its integral gives $4mw_0\sqrt{E_c}/\sqrt{2m} = 2p_c w_0$ as required. A positive-direction selector halves this particular probability when applied to the original ensemble. If the entire subsequent experiment is renormalised to the selected particles, divide again by the total positive-momentum probability. Those two reported probabilities correspond to different conditioning and should be labelled explicitly.

Infer the packet width from two images

The model atom has an initial measured width 1.56 µm and a width 1.85 µm after 2.73 ms. Assume a time-independent Gaussian instrument response and a freely evolving unchirped minimum-uncertainty state. Infer s and d, and identify an additional check.

▼ Work through the solution

Subtract squared widths to remove the detector:

$$s_{\mathrm{obs}}(t)^2 - s_{\mathrm{obs}}(0)^2 = \hbar^2 t^2/(4m^2 s^2).$$

The observed difference is about $0.99\,\mu\mathrm{m}^2$, giving $s \simeq 1.0\,\mu\mathrm{m}$ and $d = \sqrt{s_{\mathrm{obs}}(0)^2 - s^2} \simeq 1.2\,\mu\mathrm{m}$, within rounding. A third flight time tests the predicted quadratic in t. A term linear in time would point to a nonzero initial correlation or to measurement drift. Independent calibration of the point-spread function distinguishes those interpretations from a nonminimum momentum width.

Further reading

- Fließbach · Eigenwerte und Eigenfunktionen
- TUM · Theoretische Physik 3, representations and momentum states
- Course investigation · probability density and current

Quantum states, dynamics and measurement · 3. Why an oscillator can move without its packet spreading

Start with the measured mass and curvature of a potential. Factor its Hamiltonian, prove the energy ladder, construct a freely moving Gaussian from number states and find when squeezing or anharmonicity changes that prediction.

Use the equilibrium length and momentum scales

Near a stable minimum a smooth one-dimensional potential has expansion $V(x) = V(0) + m\omega^2 x^2/2 + O(x^3)$ after the origin is shifted to the minimum. Choose the arbitrary zero $V(0) = 0$. The fitted curvature is $m\omega^2$, not ω itself; measuring both the oscillation frequency and restoring slope separately identifies the effective mass. The classical energy ellipse becomes a circle if position and momentum are rescaled consistently. Define $X = \sqrt{m\omega/\hbar}\, x$ and $P = p/\sqrt{m\hbar\omega}$, so $[X, P] = i$ and $H = \hbar\omega(X^2 + P^2)/2$. Since position and momentum fail to commute, multiplying two linear factors must retain their order.

Set $a = (X + iP)/\sqrt{2}$ and $a^\dagger = (X - iP)/\sqrt{2}$. Then

$$a^\dagger a = (X^2 + P^2 + i[X, P])/2 = (X^2 + P^2 - 1)/2.$$

Thus $H = \hbar\omega(N + 1/2)$, $N = a^\dagger a$ and $[a, a^\dagger] = 1$. For every normalised state in the relevant domain $\langle N\rangle = \|a|\psi\rangle\|^2 \geq 0$. The additive half quantum is an energy contribution that survives in a state with vanishing classical means of x, p; discarding the commutator would lose it. Transition frequencies within an ideal ladder measure energy differences, so those frequencies alone cannot determine an absolute zero of energy.

Prove discreteness and construct the ground state

Using the product rule for commutators gives

$$[N, a] = -a,$$
$$[N, a^\dagger] = a^\dagger.$$

If a normalisable energy eigenstate has N eigenvalue ν, application of a either annihilates it or yields an eigenstate of eigenvalue $\nu - 1$ and norm squared ν. Repeated lowering cannot yield a state of negative norm squared. For noninteger nonnegative ν the chain would reach $0 < \nu' < 1$ with nonzero image at a negative eigenvalue. Therefore the chain must terminate at an integer step in an annihilated ground state. This argument uses the bound oscillator's spectrum and states in the ladder operators' domains. Existence and completeness can also be checked by solving the one-dimensional differential equation with decaying boundary conditions.

In coordinate representation $P = -i\, d/dX$, so $a\psi_0 = 0$ reduces to

$$d\psi_0/dX = -X\psi_0.$$

Its square-integrable solution is

$$\psi_0(X) = \pi^{-1/4} e^{-X^2/2}.$$

Normalised excited states follow by $|n\rangle = (a^\dagger)^n|0\rangle/\sqrt{n!}$. Orthogonality follows from the self-adjoint number operator with distinct eigenvalues. The coordinates in metres have ground variance $\sigma_x^2 = \hbar/(2m\omega)$, momentum variance $\sigma_p^2 = m\hbar\omega/2$, and zero symmetrised covariance. The geometric ground-state width is thus fixed by mass and trap stiffness. A semiclassical estimate minimising $\hbar^2/(8m\sigma_x^2) + m\omega^2\sigma_x^2/2$ gives $E_{\min} = \hbar\omega/2$, because the Gaussian saturates the uncertainty inequality used in that estimate.

In a number state both $\langle x\rangle$ and $\langle p\rangle$ vanish. Their fluctuations do not. Substitute $x = \sigma_x(a + a^\dagger)$. Terms $a^2, (a^\dagger)^2$ change the occupation by two and have zero diagonal element, whereas $aa^\dagger$ and $a^\dagger a$ contribute $n + 1$ and n. Hence $\langle x^2\rangle_n = (2n + 1)\sigma_x^2$; similarly $\langle p^2\rangle_n = (2n + 1)\sigma_p^2$. The mean kinetic and potential energies each equal half the total. Measuring a nonzero rms displacement does not imply that the oscillator centre is undergoing a coherent motion.

Find a state whose centre follows the classical trajectory

Search for a normalised eigenstate of a with complex eigenvalue α. It is permissible for this non-Hermitian operator to have complex eigenvalues; they label states rather than directly observable outcomes. Expand $|\alpha\rangle = \sum_n c_n|n\rangle$. Matching every number-state coefficient yields

$$c_{n+1} = \alpha c_n/\sqrt{n + 1}.$$

Using the exponential series to normalise gives

$$|\alpha\rangle = e^{-|\alpha|^2/2}\sum_{n=0}^{\infty}\frac{\alpha^n}{\sqrt{n!}}|n\rangle,$$

$$P(n) = e^{-|\alpha|^2}\frac{|\alpha|^{2n}}{n!}.$$

The mean and variance of occupation are both $|\alpha|^2$. Different coherent states are nonorthogonal: summing the overlap yields

$$|\langle\beta|\alpha\rangle|^2 = e^{-|\alpha-\beta|^2}.$$

At displacement difference one in this dimensionless complex plane the overlap probability is still $e^{-1} = 0.368$. A measurement that distinguishes all such nearby coherent preparations perfectly from one copy would contradict this overlap.

Each coefficient acquires $e^{-i(n+1/2)\omega t}$. Take the common zero-point phase outside the sum, leaving exactly the same expansion with $\alpha(t) = \alpha e^{-i\omega t}$. Hence $\langle x(t)\rangle = 2\sigma_x \operatorname{Re}\alpha(t)$ and $\langle p(t)\rangle = 2\sigma_p \operatorname{Im}\alpha(t)$. The centred operator $a - \alpha$ annihilates the state, so its centred quadrature variances match those of the ground state. The packet moves on the classical phase-space ellipse without spreading. In a free particle the quadratic momentum phase causes spreading; in the bound oscillator equally spaced energy levels preserve this particular relative-phase structure.

For a vibrational mode with $m = 2.00 \times 10^{-25}$ kg and $\omega = 2\pi \times 100$ kHz, $\sigma_x = 20.5$ nm. Prepare $\alpha = 1.5i$. Initially the average position is zero and the momentum is positive; a quarter period later the centre reaches $3\sigma_x = 61.5$ nm. The mean occupation is 2.25 and the ground probability is 0.1054. This ensemble has a nonzero ground-state probability even though its position oscillation is clearly resolved by repeated

measurements. A pure number state of average energy near this value could give similar rms positions but would show no time-dependent centre.

Background · Which measurement distinguishes two oscillators of equal mean energy?

Compare the first excited number state with a coherent state of amplitude $\alpha = 1$. Both have mean occupation one and hence the same mean energy $3\hbar\omega/2$. For a number state the expectation of the lowering operator vanishes because it connects orthogonal levels. For a coherent state its expectation equals α and evolves with the oscillator phase.

$$\langle x(t)\rangle_{\mathrm{coh}} = \sqrt{\frac{2\hbar}{m\omega}}\,\mathrm{Re}(\alpha e^{-i\omega t}),$$

$$\langle x(t)\rangle_{n=1} = 0.$$

Synchronous position measurements on freshly prepared states distinguish the moving centroid from the stationary one. A repeated energy measurement also separates the preparations. The number state always gives one quantum, whereas the coherent state gives a Poisson distribution with mean and variance both equal to one. A mean energy by itself omits both distinctions.

Rotate an uncertainty ellipse and test a nonideal trap

A sudden frequency change provides an experimentally meaningful way to generate a state with unequal quadrature fluctuations. Prepare the ground state at frequency ω_i, then instantaneously change to ω_f. The state in the position representation is continuous through an ideal sudden change. Relative to the final oscillator,

$$\langle X_f^2\rangle = \omega_f/(2\omega_i),$$

$$\langle P_f^2\rangle = \omega_i/(2\omega_f);$$

their product remains $1/4$. Their sum exceeds one unless the two frequencies coincide. Thus the state is a squeezed vacuum for the new trap, with mean occupation $\frac{1}{4}(\omega_f/\omega_i + \omega_i/\omega_f - 2)$. Unlike a coherent displacement its centre remains zero. Measuring a variance as the ellipse rotates distinguishes these preparations even if their average energies match.

For $\omega_f = 2\omega_i$ the mean occupation is 0.125. Since $X_f(t) = X_f(0)\cos\omega_f t + P_f(0)\sin\omega_f t$, variance oscillates between 1 and 0.25 in these units with angular frequency $2\omega_f$. The ground variance for the final trap would be 0.5, so part of the cycle lies below it. Momentum variance rises correspondingly, and at intermediate angles a nonzero covariance is needed in the uncertainty determinant. Abrupt changes also inject work via the time-dependent Hamiltonian. Energy measured before and after the sudden change need not agree even though the state vector at the switching instant is unchanged.

Real traps eventually depart from exact harmonic confinement. Write a model diagonal spectrum $E_n/\hbar = \omega(n + 1/2) + \chi n(n-1)/2$ for weak amplitude-dependent spacings over the occupied range. Coherent-state sums give

$$\langle a(t)\rangle = \alpha e^{-i\omega t}\exp[|\alpha|^2(e^{-i\chi t} - 1)].$$

Its magnitude is $|\alpha|\exp[-|\alpha|^2(1-\cos\chi t)]$. The collapse of the centre can occur under strictly unitary evolution while the number probabilities stay fixed. Their phases reassemble at $2\pi/|\chi|$ in this ideal spectrum. Irreversible noise and uncontrolled fluctuations can suppress a measured revival; loss of the mean alone does not diagnose dissipation. A position trace and a number-resolved distribution probe complementary parts of the dynamics.

Factor a measured displacement into force history and oscillator response

A linear force couples to position,

$$H(t) = \hbar\omega(N + 1/2) - F(t)x.$$

Even if the force time series is irregular, this driven oscillator is exactly soluble within the harmonic approximation. In the Heisenberg picture use $\dot{a} = (i/\hbar)[H, a]$. Since $[x, a] = -\sigma_x$, the equation is

$$\dot{a} = -i\omega a + i\sigma_x F(t)/\hbar.$$

Multiply by $e^{i\omega t}$ and integrate:

$$a(t) = e^{-i\omega t}\left[a(0) + \frac{i\sigma_x}{\hbar}\int_0^t F(t')e^{i\omega t'}dt'\right].$$

A force therefore displaces the rotating oscillator according to one complex Fourier component of its history. Its phase determines which final quadrature changes. The centred operator still evolves as in free oscillation. Thus a known linear drive displaces a ground state into a coherent state without squeezing it. The same formula shifts a thermal state without removing its thermal fluctuations. The statement about pure coherence requires knowing the initial state as well as measuring the final displacement.

Take a constant force F_0 switched on for time T then off. For an initially stationary ground state the final amplitude, measured at switch-off in the original oscillator basis, is $\alpha_F = F_0\sigma_x[1 - e^{-i\omega T}]/(\hbar\omega)$. The corresponding mean position is $F_0(1-\cos\omega T)/(m\omega^2)$, agreeing with the classical step response of a mass on a spring. Its excitation number is $4(F_0\sigma_x/\hbar\omega)^2\sin^2(\omega T/2)$. A constant force applied for exactly a full period creates zero residual excitation at switch-off even though the oscillator moved substantially during the pulse. Motion of the trap centre and persistent excitation after the trap returns must be experimentally distinguished.

For $m = 2.00 \times 10^{-25}$ kg and 100 kHz as above, let $F_0 = 2.00 \times 10^{-21}$ N. Its static equilibrium displacement is

$$F_0/(m\omega^2) = 25.3\,\mathrm{nm}.$$

A half-period pulse leaves the mean at 50.7 nm with zero mean momentum and $|\alpha| = 50.7/(2 \times 20.5) \simeq 1.24$. Its mean added occupation is roughly 1.53. The elastic energy $m\omega^2\langle x\rangle^2/2$ at switch-off matches $\hbar\omega|\alpha|^2$, independently of any counting argument. This check subtracts the unchanged zero-point energy. Using the total energy including zero point in the comparison would incorrectly attribute a half quantum to the drive even when the applied force vanishes.

The sudden impulse limit provides another independent test. Let a short pulse deliver total impulse $J = \int F dt$ over a time much shorter than $1/\omega$. The amplitude increment is approximately $i\sigma_x J/\hbar$ in the appropriate rotating phase. Since $2\sigma_p\sigma_x = \hbar$, the mean momentum changes by J. For a smooth pulse spread around one period the same total impulse can produce a much smaller residual excitation because the force acts at different phases of the oscillator. A force sensor based on final occupation measures the squared magnitude of a windowed resonant integral. Reconstructing a complete time-dependent force from that single energy measurement is generally impossible. Varying the drive start time or measuring both final quadratures gives additional information.

Use operator ordering to predict a phase with zero final displacement

Define the displacement operator $D(\alpha) = \exp(\alpha a^\dagger - \alpha^* a)$. Its generator is anti-Hermitian, so it is unitary. The commutator of $\alpha a^\dagger - \alpha^* a$ with a is a constant. Differentiating $D^\dagger(s\alpha) a D(s\alpha)$ with respect to s and integrating gives

$$D^\dagger(\alpha) a D(\alpha) = a + \alpha.$$

Consequently $D(\alpha)|0\rangle$ is an annihilation eigenstate of eigenvalue α, giving an independent physical construction of the number expansion. The normal-order factorisation $D = e^{-|\alpha|^2/2} e^{\alpha a^\dagger} e^{-\alpha^* a}$ follows from the Baker–Campbell–Hausdorff formula because the relevant commutator is a scalar. Acting on the ground state recovers the Poisson coefficients directly. Omitting the scalar prefactor gives a state with the wrong norm.

Two displacements do not commute as operators despite their effects on the mean phase-space position adding in either order. Their scalar commutator gives

$$D(\alpha) D(\beta) = e^{(\alpha\beta^* - \alpha^*\beta)/2} D(\alpha + \beta).$$

The product around a closed displacement loop is

$$D(\alpha) D(\beta) D(-\alpha) D(-\beta) = e^{\alpha\beta^* - \alpha^*\beta} I.$$

For $\alpha = 1,\ \beta = i/2$ the phase is -1 rad in this written operator order. Reversing the physical order conjugates the phase. The oscillator returns to its initial phase-space position and, apart from this phase, to its initial state. In a single isolated preparation the common phase is invisible. To observe it, place the oscillator interaction under coherent control of one branch of a separate two-level reference, recombine the reference and compare the branch that went around the loop with the branch that did not. Surviving coherence of the reference is essential. The same result can be described as an oriented area in the $\mathrm{Re}\,\alpha, \mathrm{Im}\,\alpha$ plane and ultimately follows from $[x, p] = i\hbar$.

It is useful to trace the actual motion before interpreting a claimed geometric phase. Operators act from right to left on kets, so plotting a particular rectangular loop in time can correspond to the conjugate of a product written left to right. A lab comparison should specify the four pulse phases in chronological order and the ordinary free rotation occurring between them. In the rotating frame the displacement vectors are simple complex increments; in the laboratory frame the oscillator itself continually advances in phase. An uncompensated delay can leave a residual displacement rather than a closed loop. The overlap of that

residual displaced ground state with the reference branch is $e^{-|\Delta\alpha|^2/2}$ in amplitude, reducing observed fringe contrast as well as shifting its phase.

Measure displaced parity to distinguish a number state from a mixture

The number parity operator $\Pi = (-1)^N$ gives $+1$ for even and -1 for odd occupation. It implements the coordinate reflection $x \to -x$ for the harmonic oscillator. A position-space Wigner distribution may be sampled by displacing the state and measuring its parity. In the complex phase-space convention used here define $W(\beta) = (2/\pi)\,\mathrm{tr}[D(-\beta)\rho D^\dagger(-\beta)\Pi]$; it integrates to one in $d^2\beta$. At $\beta = 0$ it is just twice the even-minus-odd probability divided by π. Its sign is determined by observable population statistics, but its value at a point is not the probability that a particle occupies a sharp simultaneous position and momentum.

For a coherent input $|\alpha\rangle$, displacement produces $|\alpha - \beta\rangle$ up to a phase. Summing the even-minus-odd Poisson occupations gives $e^{-2|\alpha-\beta|^2}$. Hence $W(\beta) = 2e^{-2|\alpha-\beta|^2}/\pi$, a positive Gaussian. A first excited number state instead has $W(0) = -2/\pi$. Mix that state with ground-state population,

$$\rho = p|1\rangle\langle 1| + (1-p)|0\rangle\langle 0|.$$

Then $W(0) = (2/\pi)(1-2p)$, negative only when $p > 1/2$. An observed lack of negativity at one displacement does not imply a Gaussian state or a globally classical description. It may simply mean that the probe is examining a positive region or that losses have mixed the odd preparation with the ground state.

Suppose an experimental detector distinguishes only “zero excitations” from “one or more”. A measured nonzero fraction of 0.80 does not determine parity without further information: all those events might contain one excitation or a substantial even population might be present. Number resolution or a separate calibrated parity mapping is needed. If only ground and first excited population are independently established, an observed $p_1 = 0.70 \pm 0.02$ gives $W(0) = -0.255 \pm 0.025$ in these dimensionless units. Propagating that uncertainty requires the assumption about higher occupations. Combining a displacement scan with number-resolved or parity-sensitive data therefore tests both the preparation and the adequacy of a finite oscillator basis.

Recover a thermal occupation from two independent observables

For an oscillator in thermal equilibrium at inverse temperature β_T, the occupation probabilities are $P_n = (1-q)q^n$ with $q = e^{-\beta_T\hbar\omega}$. The geometric series gives $\bar{n} = q/(1-q)$ and variance $\bar{n}(\bar{n}+1)$. Unlike a coherent state with the same mean occupation, the number distribution is broader and the free mean displacement is zero. The symmetrised position variance is $(2\bar{n}+1)\sigma_x^2$. A variance measurement alone cannot distinguish a thermal distribution from an incoherent nonthermal mixture with the same mean energy. Resolving more than the variance probes the full probability law.

For example, at $\bar{n} = 2$ thermal equilibrium predicts ground occupation $1/3$, number variance 6 and parity $1/(2\bar{n}+1) = 1/5$. A coherent state with $|\alpha|^2 = 2$ gives ground probability 0.1353, variance 2 and parity $e^{-4} = 0.0183$. A uniform randomisation of the coherent phase removes the mean displacement but leaves these Poisson number statistics unchanged. It can therefore resemble the thermal state on a quadrature average while differing sharply in ground population and parity. A displaced thermal preparation adds a

classical centre motion on top of the thermal fluctuations; fitting only its energy would confuse displacement with temperature. Centre, centred variance and at least one independent number statistic make a coherent experimental set of observables for these physically different preparations.

Follow a coherent oscillator state in phase space

Change the magnitude and phase of a coherent amplitude independently. Predict which control changes the mean occupation and which shifts the timing of the position measurement.

Open this experiment online

Try the calculation

Coherent amplitude magnitude **1.3**

Preparation phase (degrees) **30**

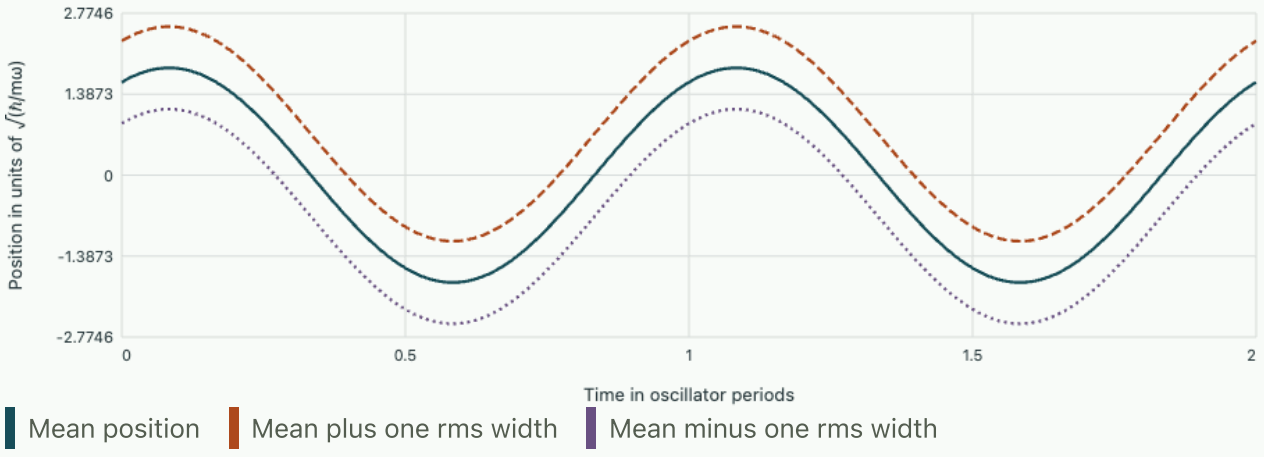

Mean occupation 1.69 quanta, standard deviation 1.3. Ground-state probability 0.18452. Position variance in the plotted units is 1/2.

Recover individual transition probabilities after a sudden trap change

The sudden-frequency experiment predicted mean occupation from the second moments. Number-resolved probabilities provide a stronger test. The initial ground-state wavefunction is even, and changing the trap curvature about the same centre preserves parity. The initial state therefore overlaps only even eigenstates of the final trap. A measured odd population beyond readout uncertainty diagnoses a displaced minimum, parity-breaking drive or imperfect initial preparation. The vanishing of odd occupation is stronger than knowing only the mean and variance of a quadrature.

Write the old annihilator in terms of the new operators,

$$a_i = \cosh r\, a_f - \sinh r\, a_f^\dagger,$$

with $r = \frac{1}{2}\ln(\omega_f/\omega_i)$ for this choice of sign. This follows directly by expressing both in the same x, p and comparing coefficients. The relative sign can change under a different convention for defining the squeeze parameter but all probabilities below are invariant. The old vacuum obeys $(a_f - \tanh r\, a_f^\dagger)|0_i\rangle = 0$. In the final number basis $c_n = \langle n_f|0_i\rangle$, the coefficient of $|n\rangle$ gives

$$\sqrt{n+1}c_{n+1} = \tanh r\sqrt{n}\, c_{n-1}.$$

The $n = 0$ equation forces $c_1 = 0$, and iteration gives

$$c_{2j} = c_0 \frac{\sqrt{(2j)!}}{2^j j!}(\tanh r)^j,$$

$$c_{2j+1} = 0.$$

To normalise, use the central-binomial series $\sum_j \binom{2j}{j} q^j/4^j = (1-q)^{-1/2}$ at $q = \tanh^2 r$. Thus

$$|c_0|^2 = 1/\cosh r.$$

A separate coordinate-space check follows by integrating the overlap of two Gaussian ground states, giving $P_0 = 2\sqrt{\omega_i\omega_f}/(\omega_i + \omega_f)$. This agrees with $1/\cosh r$, independently of the operator recursion. Differentiating the normalisation-generating series gives $\bar{n} = \sinh^2 r$, agreeing with the energy inferred from the sudden change of the Hamiltonian.

For a doubling of frequency, $\tanh r = 1/3,\ \cosh r = 3/(2\sqrt{2})$, so

$$P_0 = 0.9428,\ P_2 = P_0/18 = 0.05238,$$

$$P_4 = P_0(3/8)(1/81) = 0.00436.$$

Higher even levels contain the remaining probability. At a mean excitation of 0.125, a coherent state would instead have $P_1 = 0.1103$. Even imperfect number resolution discriminating odd versus even can strongly distinguish a pure curvature quench from a coherent displacement with equal final energy. A finite switching time can reduce excitation if the oscillator tracks the moving ground state adiabatically. A shift of trap centre during a nominal curvature change adds a displacement and can populate odd levels. Measure the mean and parity along with the energy estimate before inferring a squeeze strength solely from a single occupation number.

Separate temperature from a parity-preserving quench

Before a sudden centred frequency doubling, a calibrated trap has thermal mean occupation 0.10. Predict the fraction of odd-number results after the quench and its final mean occupation. How might an experiment mistakenly interpret the odd results?

▼ **Derive and check the result**

In the initial thermal ensemble parity expectation is

$$1/(2\bar{n}+1) = 1/1.2 = 0.8333.$$

Parity commutes with both centred trap Hamiltonians, and measuring even or odd final occupation measures that same parity operator. The odds fraction is therefore $(1 - 0.8333)/2 = 0.0833$, unchanged by the quench. The final energy fluctuations scale with the initial quadrature variances multiplied by $2\bar{n} + 1$, giving $\bar{n}_f + 1/2 = (\bar{n}_i + 1/2)(\omega_f/\omega_i + \omega_i/\omega_f)/2 = 0.75$. Thus $\bar{n}_f = 0.25$. Assuming a perfectly cold initial state might turn the observed odd fraction into a false diagnosis of trap displacement. Measure initial parity and post-quench parity independently to isolate processes that actually break the symmetry. The increase of mean occupation alongside unchanged parity is the predicted signature of a centred quench.

Tutorial · analyse a state from amplitude and fluctuation data

Test two preparations with the same mean energy

Take a coherent state with $|\alpha|^2 = 3$ and number state $|3\rangle$. Compare energies, number variance, ground probability and the largest possible mean displacement.

▼ **Work through the solution**

Both have $\langle H \rangle = 3.5\hbar\omega$. The coherent number variance is 3 while the number-state variance is zero. Ground probabilities are $e^{-3} = 0.0498$ and zero. By selecting a suitable oscillator phase, the coherent state reaches $2\sqrt{3}\sigma_x$; the number state has zero centre at every time. An energy average alone leaves these preparations indistinguishable. The event of measuring exactly three quanta has probability $e^{-3}3^3/3! = 0.224$ in the coherent case and unity in the number case. Repeated shots allow this comparison without assuming a single shot reveals the input state.

Estimate numerical truncation before trusting a simulated coherent state

Expand a coherent state with $\alpha = 2$ in number states $0, \ldots, 8$. Is the missing probability negligible at a desired accuracy of one part in a thousand?

▼ **Work through the solution**

Generate probabilities successively from $p_0 = e^{-4}$, $p_{n+1} = 4p_n/(n+1)$; summing through 8 gives 0.97864, leaving 0.02136 above the cutoff. Renormalising the retained vector hides this discarded probability from a norm check and distorts quantities weighted by n. Through $n = 11$ the tail is about 0.00092, just below the specified probability tolerance. Energy moments require a separate tail weighted by n; a probability bound alone can be inadequate for them. Vary the cutoff while monitoring both number statistics and the centre motion to separate truncation artefacts from real dephasing.

Infer the coherent displacement and a squeezing test

At a known phase, measurements give

$$\langle x \rangle = 1.2\sigma_x, \ \langle p \rangle = -1.6\sigma_p.$$

Infer the coherent amplitude consistent with those means, then say what must still be measured to confirm a coherent preparation.

▼ **Work through the solution**

The means imply $\mathrm{Re}\,\alpha = 0.60,\ \mathrm{Im}\,\alpha = -0.80$, so $|\alpha| = 1$ and the putative coherent mean occupation is one. They do not fix the variances. A displaced thermal or squeezed state can produce the same means with different fluctuations. Measure the variance along independent quadratures, including a rotated axis to infer covariance. In an ideal normalised state the expectation of the positive operator $(a-\alpha)^\dagger(a-\alpha)$ equals $\langle N\rangle - |\alpha|^2$. If it vanishes exactly, the support is in the kernel of $a-\alpha$, giving the coherent state. With finite uncertainty a near-zero value instead bounds extra excitation in the displaced frame.

Further reading

- Fließbach · Operatorenmethode
- Course experiment · operators and oscillator spectra
- TUM · quantum mechanics, harmonic oscillator and operator formulation

Quantum states, dynamics and measurement · 4. Adding angular momentum and predicting which lines can appear

A three-component orbital wavefunction combines with a two-component spin. Derive the allowed total angular momenta, construct their amplitudes explicitly and predict both forbidden transitions and relative strengths in a polarisation-resolved spectrum.

Read the angular momentum algebra as a statement about rotations

Rotate a prepared quantum state through a small angle $\delta\theta$ about a unit laboratory direction $\mathbf{n}$. A continuous norm-preserving transformation has form $U = I - i\delta\theta\,\mathbf{n}\cdot\mathbf{J}/\hbar + O(\delta\theta^2)$, with Hermitian generators J_i. Rotations about different axes need not commute. Comparing the commutator of two infinitesimal rotations with the rotation around the third axis gives

$$[J_i, J_j] = i\hbar\epsilon_{ijk}J_k.$$

The sum $J^2 = J_x^2 + J_y^2 + J_z^2$ commutes with each component. One can therefore label eigenstates of J^2 and one chosen component J_z simultaneously, while generally leaving the other components uncertain.

Define $J_\pm = J_x \pm iJ_y$. Their commutators with J_z are $\pm\hbar J_\pm$; with J^2 they vanish. If $J_z|m\rangle = \hbar m|m\rangle$, the two ladder operations change m by one without changing total angular momentum. To determine their coefficients, use $J_-J_+ = J^2 - J_z^2 - \hbar J_z$ and $J_+J_- = J^2 - J_z^2 + \hbar J_z$. Nonnegative squared norms bound the ladder above and below. Let j be its maximum magnetic quantum number. Vanishing raising norm there fixes the J^2 eigenvalue at $\hbar^2 j(j+1)$; vanishing lowering norm at the bottom gives

$$m_{\min} = -j.$$

The number of unit steps between the two endpoints is $2j$, so j is an integer or half-integer. With a consistent phase convention

$$J_\pm|j,m\rangle = \hbar\sqrt{j(j+1) - m(m\pm1)}\,|j,m\pm1\rangle.$$

Spatial orbital eigenfunctions on an ordinary single-valued scalar wavefunction carry integer orbital l. Spin-$\frac{1}{2}$ has two states that cannot be modelled by adding another spatial coordinate angle to a scalar orbital. The operator algebra describes both. In $|j,j\rangle$, J_z is definite but J_x, J_y fluctuate. Their variance sum is $\hbar^2 j(j+1) - \hbar^2 j^2 = \hbar^2 j$; rotational symmetry around z splits it evenly. Drawing a classical vector with precisely fixed components perpendicular to the axis would miss these fluctuations.

Construct a coupled basis from the stretched state

For independent angular momenta $\mathbf{L}, \mathbf{S}$, the total generator is $\mathbf{J} = \mathbf{L} + \mathbf{S}$. Operators on different factors commute. In the product basis $|m_l, m_s\rangle$, J_z is diagonal with eigenvalue $\hbar(m_l + m_s)$. The Hamiltonian may instead depend on J^2, for example through $\mathbf{L}\cdot\mathbf{S}$. Then states with the same $m_l + m_s$ can mix. Resolving that degeneracy requires a unitary change of basis, whose entries are Clebsch–Gordan coefficients.

Take $l = 1, s = 1/2$. The unique state $m = 3/2$ is $|1,\uparrow\rangle$, so it must be $|3/2, 3/2\rangle$. Lower once with $J_- = L_- + S_-$. The left side is $\hbar\sqrt{3}|3/2, 1/2\rangle$; the right side is $\hbar(\sqrt{2}|0,\uparrow\rangle + |1,\downarrow\rangle)$. Normalising and constructing an orthogonal state in the same m plane gives

$$|3/2,1/2\rangle = \sqrt{2/3}|0,\uparrow\rangle + \sqrt{1/3}|1,\downarrow\rangle,$$

$$|1/2,1/2\rangle = \sqrt{1/3}|0,\uparrow\rangle - \sqrt{2/3}|1,\downarrow\rangle.$$

The sign of an entire coupled vector is conventional. The relative sign of its two components matters whenever a later measurement recombines amplitudes. Applying J_+ to the second combination gives zero, which verifies the top of the $j = 1/2$ multiplet independently of orthogonality. Lowering produces four states for $j = 3/2$ and two for $j = 1/2$, exhausting the six-dimensional product space. Merely listing both possible j's does not supply the amplitude with which a prepared product state belongs to each, which is what spectroscopy needs.

There is a counting derivation for arbitrary j_1, j_2. Start with one stretched state at $m = j_1 + j_2$ and remove its entire ladder. The next highest unused m supplies the next multiplet and so on down to $|j_1 - j_2|$. The identity $\sum_{j=|j_1-j_2|}^{j_1+j_2}(2j+1) = (2j_1+1)(2j_2+1)$ checks completeness. Distinct coupling schemes for three angular momenta are different bases of a common space. An intermediate angular momentum may be a reliable state label when a corresponding interaction dominates, yet cease to be conserved if an external field favours another axis or coupling.

Derive magnetic and parity selection rules

Consider a weak optical field coupled by the electric dipole interaction $-\mathbf{d}\cdot\mathbf{E}$, where $\mathbf{d}$ transforms as a vector and has odd parity. Its spherical components relative to z are

$$d_0 = d_z,\; d_{\pm1} = \mp(d_x \pm i d_y)/\sqrt{2}.$$

For a definite tensor component $[J_z, d_q] = q\hbar d_q$. Sandwiching the commutator between initial and final J_z eigenstates yields

$$(m_f - m_i - q)\langle f|d_q|i\rangle = 0.$$

Thus nonzero amplitudes require

$$m_f = m_i + q.$$

Coupling j_i with a rank-one tensor also requires $|j_i - 1| \le j_f \le j_i + 1$, and $0 \to 0$ is excluded. Parity gives a separate test: under inversion the matrix element gains factor $-\pi_f\pi_i$, so it vanishes for equal initial and final parity in a parity-symmetric Hamiltonian. For purely orbital motion with parity $(-1)^l$, the triangle condition plus odd parity gives

$$\Delta l = \pm 1.$$

A direction label on a photon depends on propagation and polarisation conventions. Stating the driven operator component q is less ambiguous than calling a beam “right circular” without defining the viewing direction.

For a spin-independent dipole transition the spin factor is $\langle m_s'|m_s\rangle$. Begin in $l_i = 0, m_l = 0, m_s = +1/2$ and drive $q = 0$ to $l_f = 1, m_l = 0$. In the product basis there is one equally normalised orbital amplitude D. Its projections on the two coupled final levels have squared amplitudes $2|D|^2/3$ and $|D|^2/3$, giving a 2:1 line-strength ratio under equal field intensity, comparable detection and negligible population saturation. For $q = +1$ the product state is $|1, \uparrow\rangle$. The transition to final $j = 1/2$ vanishes because no $m = 3/2$ belongs to that multiplet. These statements predict relative angular factors. Actual count rates can also depend on frequencies, radial matrix elements, initial populations and how fluorescence is collected.

As an explicit orbital check,

$$Y_{00} = 1/\sqrt{4\pi},$$

$$Y_{10} = \sqrt{3/(4\pi)}\cos\theta.$$

The angular factor of $d_0 \propto r\cos\theta$ is $\int Y_{10}^* \cos\theta Y_{00}\, d\Omega = 1/\sqrt{3}$, while for $Y_{1,\pm1}$ the azimuthal integral vanishes. A radial integral can still vanish accidentally even when all angular tests permit the transition. Symmetry rules predict zeros robust to changes preserving that symmetry; a permitted channel has an amplitude that must still be calculated.

Separate atomic structure from analyser orientation with a reduced matrix element

The angular selection rule from the J_z commutator identifies which transitions are forced to zero but does not alone give their relative strengths. A rotationally covariant way to connect the allowed amplitudes is to view an irreducible tensor component $T_q^{(k)}$ as carrying angular momentum rank k. Under rotations its $2k+1$ components transform among themselves with the same representation matrices as the magnetic substates of an angular momentum k. Coupling $T_q^{(k)}$ to an initial $|jm\rangle$ therefore furnishes the same product representation as coupling $|jm\rangle$ to $|kq\rangle$. The component projecting onto a final j' must transform as j'. When this representation occurs once in the product, its relative magnetic amplitudes are fixed by Clebsch–Gordan coefficients up to one common factor independent of m, q, m'.

In one standard convention the Wigner–Eckart theorem writes

$$\langle \alpha' j' m'|T_q^{(k)}|\alpha j m\rangle = \frac{\langle jm; kq|j'm'\rangle}{\sqrt{2j'+1}}\,\langle \alpha' j'\|T^{(k)}\|\alpha j\rangle.$$

Additional labels α, α' specify the radial, spin or configuration information absent from the angular labels. Conventions placing a different square-root factor in the reduced element are equally common, so a table of reduced strengths should be checked before combining it with an external coefficient table. The theorem predicts ratios within a fixed j, j' multiplet. It does not remove differences in radial integrals between unrelated configurations, nor does it assert the reduced element is nonzero. This factorisation explains why angular data alone may fix a branching ratio while offering no absolute radiative lifetime.

For a vector operator connecting j to the same j, there is an especially transparent example. Within a given irreducible j space, any vector operator has matrix elements proportional to the corresponding J_q matrix elements, provided additional degeneracies are separated. For the conventional spherical components $J_0 = J_z$, $J_{+1} = -J_+/\sqrt{2}$, ask for the ratio of line strengths out of $j = 2, m = 1$ into $m' = 2$ via $q = +1$

and into $m' = 1$ via $q = 0$, with equal field amplitudes in the specified components. The squared matrix elements of J_q are $\hbar^2[6 - 1(2)]/2 = 2\hbar^2$ and $\hbar^2$, hence the ratio is two. The same numerical angular ratio can apply to any appropriate nonzero reduced vector coupling between these specified manifolds. For an electric dipole transition parity must still allow the initial and final configurations; using J_q here computes angular ratios and does not turn angular momentum itself into an electric dipole operator.

Summing over a complete set of final magnetic substates and all tensor components supplies a useful normalisation audit. For a fixed initial m, rotational invariance of the scalar sum $\sum_q T_q^\dagger P_{j'} T_q$ makes its diagonal expectation independent of m. A detector with a preferred polarisation or direction samples only a subset and may exhibit strong dependence on m. This difference is essential in interpreting "unpolarised" line strengths: an unpolarised pump that populates all m equally, an unanalysed detector that sums final polarisations, and a source with a definite m viewed along one axis are three different experimental ensembles.

Compare successive Stern–Gerlach measurements with addition of spin

Two spin-$1/2$ particles provide a minimal addition experiment with no spatial angular wavefunction needed. The coupled triplet basis is $|1,1\rangle = |\uparrow\uparrow\rangle$,

$$|1,0\rangle = (|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle)/\sqrt{2},$$
$$|1,-1\rangle = |\downarrow\downarrow\rangle;$$

the singlet is $(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)/\sqrt{2}$. An apparatus measuring the total z component alone cannot distinguish the two $m = 0$ states. Measuring J^2 along with J_z can. Their local one-spin reduced matrices are both $I/2$, so even complete tomography of either spin alone fails to reveal the difference. The distinction lies in a joint correlation.

The singlet projector can be obtained as a polynomial in $\mathbf{S}_1 \cdot \mathbf{S}_2$. That scalar has eigenvalues $-3\hbar^2/4, +\hbar^2/4$, so $P_s = I/4 - \mathbf{S}_1 \cdot \mathbf{S}_2/\hbar^2$. Expressing $\mathbf{S}_i = (\hbar/2)\boldsymbol{\sigma}_i$ gives

$$P_s = [I - \sum_i \sigma_i \otimes \sigma_i]/4.$$

If three independently prepared ensembles yield measured same-axis correlations $(-0.80, -0.78, -0.82)$, their inferred singlet fraction is

$$(1 + 2.40)/4 = 0.85.$$

This result uses all three axes. Perfect z-anticorrelation alone constrains the population to a two-dimensional $m = 0$ space but says nothing about the relative phase within it. If each measurement has uncertainty 0.03 and the estimates are independent, the resulting statistical uncertainty in singlet fraction is $\sqrt{3}(0.03)/4 = 0.013$, before calibration uncertainties.

For a coherent initial product $|\uparrow\downarrow\rangle$, a Hamiltonian $H = J_{\mathrm{ex}} \mathbf{S}_1 \cdot \mathbf{S}_2/\hbar^2$ evolves the singlet and triplet amplitudes with distinct phases. Writing the initial state as $(|1,0\rangle + |0,0\rangle)/\sqrt{2}$, the probability to find spins reversed later is $\sin^2(J_{\mathrm{ex}} t/(2\hbar))$. The total m remains exactly zero at all times. If $J_{\mathrm{ex}}/h = 5.00$ MHz, the

first full reversal occurs at $t = 1/(2 \times 5\,\mathrm{MHz}) = 100$ ns. An incoherent mixture with the same singlet and triplet populations has a constant reversal probability one half. This calculation makes the signs in the coupled basis observable and distinguishes a true exchange evolution from random spin flips inferred solely from individual populations.

Rotate the quantisation axis without changing the atomic Hamiltonian

A transverse magnetic or polarisation analyser can interrogate a state prepared relative to z along a different laboratory axis. For spin $1/2$, $e^{-i\theta S_y/\hbar} = \cos(\theta/2)I - i\sin(\theta/2)\sigma_y$, from $\sigma_y^2 = I$. To obtain the rotation probabilities for a spin-one state without memorising a rotation matrix, represent its stretched state as the symmetric product of two up spinors. Under a common rotation each spinor transforms as $\cos(\theta/2)|\uparrow\rangle + \sin(\theta/2)|\downarrow\rangle$. Expanding the two-spin product gives the triplet amplitudes $\cos^2(\theta/2), \sqrt{2}\cos(\theta/2)\sin(\theta/2), \sin^2(\theta/2)$ for the three m values. This construction exploits the representation and applies to any angular momentum one, regardless of its microscopic origin.

At $\theta = \pi/2$, a stretched $m = 1$ state projected on the rotated axis gives probabilities $(1/4, 1/2, 1/4)$. The expectation of its rotated component vanishes, but outcomes $+\hbar$ and $-\hbar$ still occur and the variance is $\hbar^2/2$. A source prepared in $m = 0$ along z gives a different test. Rotate its symmetrised up–down form on both spins and project on $m = 0$; the amplitude is $\cos\theta$ and the probability $\cos^2\theta$. At a right angle the “zero” outcome is absent, although the mean along either axis is zero. A single mean value cannot distinguish alignment from an unpolarised mixture. Measuring the full resolved substate distribution, or its quadrupole moment, can.

The general stretched state $|j, j\rangle$ can be represented by $2j$ symmetrised spinors all up. Its rotated m outcome distribution is binomial:

$$P(m) = \binom{2j}{j-m} \cos^{2(j+m)}(\theta/2)\sin^{2(j-m)}(\theta/2).$$

Its mean angular momentum component is $\hbar j\cos\theta$, with variance $\hbar^2(j/2)\sin^2\theta$. At fixed nonzero $\cos\theta$, the relative fluctuation falls like $j^{-1/2}$ as j grows. This supplies a precise sense in which a large stretched state can approximate a classical oriented angular momentum while retaining quantum projection noise. A state with the same large j but $m = 0$ need not resemble an oriented classical arrow. The representation dimension alone does not determine the suitability of a classical approximation.

Distinguish an angularly forbidden line from coherent cancellation

Suppose two degenerate lower magnetic states $|g_1\rangle, |g_2\rangle$ both couple to one excited state $|e\rangle$ under controlled coherent drives. In a rotating-frame model the raising part is $|e\rangle(\Omega_1\langle g_1| + \Omega_2\langle g_2|)\hbar/2$, with the lowering term its adjoint. A normalised superposition proportional to $\Omega_2|g_1\rangle - \Omega_1|g_2\rangle$ has zero excitation amplitude because the two terms cancel. Each angular channel can be individually allowed and strong, yet this particular coherent combination is dark. Its orthogonal bright state is proportional to $\Omega_1^*|g_1\rangle + \Omega_2^*|g_2\rangle$ and has effective Rabi coupling $\sqrt{|\Omega_1|^2 + |\Omega_2|^2}$.

For equal real drives and an initially balanced antisymmetric superposition, the transition ideally vanishes. A randomly prepared mixture of the same two basis states populates the bright and dark subspaces equally. If a resonant pulse transfers the bright component completely, that mixture produces 50% excitation. Changing the relative optical phase changes which superposition is dark but does not change the selection rule for either individual transition. Preparing a calibrated ground-state phase and scanning the drive phase can therefore distinguish coherent cancellation from a forbidden angular matrix element or a disconnected drive channel. Two-photon detuning, magnetic splitting of the lower states and dephasing can spoil the destructive interference even if every angular coefficient is exactly known.

Track a spin while rotating the preparation axis

Change the path of a spin direction and compare a change in its geometric phase with an ordinary dynamical phase.

Open this experiment online

Try the calculation

Cone angle (degrees) **60**

Reference phase (degrees) **0**

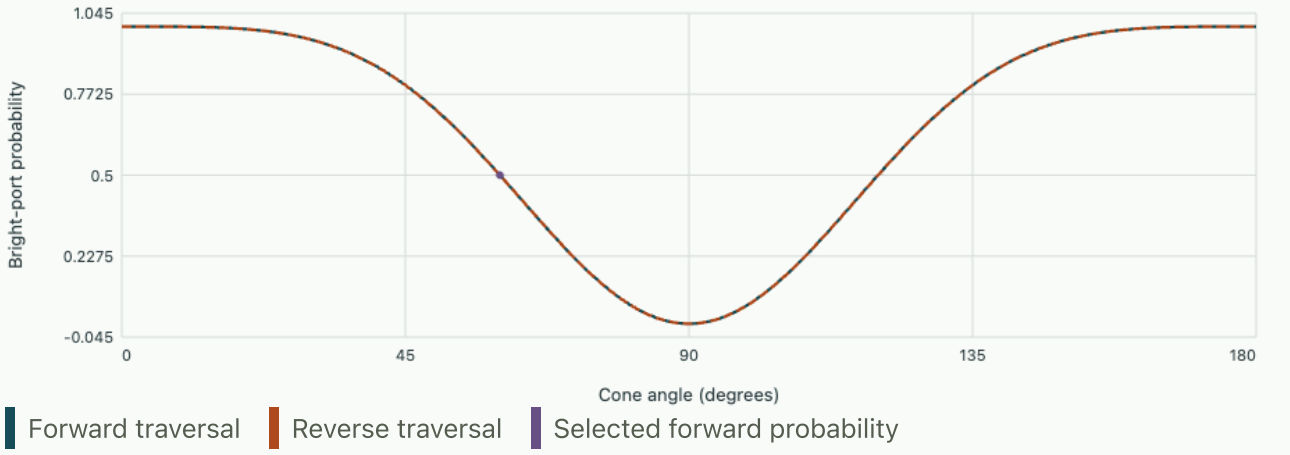

Forward geometric phase -0.5π radians. Forward port probability 0.5, reverse 0.5.

Predict which perturbation can make a dark transition visible

Electric dipole selection rules use the approximation that the driving field is essentially uniform across the atom. Expand a monochromatic plane wave $e^{i\mathbf{k}\cdot\mathbf{r}} = 1 + i\mathbf{k}\cdot\mathbf{r} + \cdots$. For a source of size a the next spatial term is suppressed in amplitude by a factor of order ka when $ka \ll 1$. Consistently expanding the light–matter Hamiltonian yields higher electric multipoles and magnetic couplings with their own symmetry. The electric quadrupole is a rank-two even-parity tensor. Thus it can connect same-parity states with $|j_i - 2| \le j_f \le j_i + 2$. A magnetic dipole is a rank-one operator of even spatial parity, permitting some same-parity transitions denied to an electric dipole. A line forbidden to one specified multipole need not vanish in all experimental conditions.

Take a model atom with a lower $j = 0$ even state and an upper $j = 2$ even state. Electric dipole fails both parity and angular-momentum tests. Magnetic dipole has the right parity but still cannot supply total angular rank two. Electric quadrupole has the correct parity and satisfies the triangle condition; its actual reduced matrix element and the applied field gradient still determine strength. In a $j = 0 \to j' = 2$ transition with $m_i = 0$, each spherical quadrupole component q excites only $m_f = q$. A simple claim that changing circular to linear polarisation opens an electric dipole line in this field-free atom would be inconsistent with both global selection rules. A static field that mixes opposite parity or different angular momentum can change the states themselves and thereby modify that conclusion.

For a 600 nm wavelength and characteristic radius $a = 0.10$ nm,

$$ka = 2\pi a/\lambda = 1.05 \times 10^{-3}.$$

The squared geometric factor is about 1.1×10^{-6} relative to a typical allowed electric dipole strength in comparable units, before distinct radial and angular factors. This is an order estimate, not a quantitative branching ratio for two arbitrary lines. In particular, dividing a finite quadrupole amplitude by a strictly zero dipole amplitude of the same ideal transition is meaningless. Compare with a separate reference line whose dipole matrix element and detection efficiencies are known if a rate ratio is needed.

An independent way to brighten an otherwise forbidden line is mixing through a perturbation W. Suppose an upper state $|u\rangle$ has zero electric dipole amplitude from $|g\rangle$, while a nearby state $|v\rangle$ allows amplitude D. If $|\langle v|W|u\rangle| \ll |E_u - E_v|$, the perturbed upper state contains coefficient $\epsilon = \langle v|W|u\rangle/(E_u - E_v)$ of $|v\rangle$, and its borrowed dipole intensity scales as $|\epsilon|^2|D|^2$. Where W is proportional to a reversible static field, borrowed intensity typically begins quadratically in that field for a previously dark isolated line. A native higher-multipole contribution may already be present at zero field and can interfere if it feeds the same final channel with coherent angular dependence.

Use polarisation and bias sweeps to identify borrowed intensity

In a two-state model the bright level is 2.00 GHz away and a perturbation mixes it with the dark state at off-diagonal $W/h = 20$ MHz. Predict relative borrowed intensity, assess the leading approximation, and identify an independent spectroscopic consequence.

▼ **Derive and check the result**

The small parameter is

$$20/2000 = 0.010.$$

The leading borrowed intensity is 10^{-4} of the corresponding bright reference, with corrections of relative order 10^{-4} in the ideal two-level mixing expansion. The dark state's energy shifts by approximately $W^2/(E_u - E_v)$, whose magnitude expressed as ordinary frequency is

$$20^2/2000 = 0.20\,\mathrm{MHz}.$$

Its sign is away from the neighbouring bright level. Measuring this level repulsion together with the same polarisation pattern as the borrowed bright transition tests the mixing explanation. Changing the perturbing field by a factor of two should increase both the leading borrowed intensity and the magnitude of the second-

order level shift by a factor of four while the coupling remains weak. A field-independent residual or a different polarisation dependence motivates examining other multipoles or additional states.

Tutorial · infer the state that a spectrum actually selects

Measure the orbital label after total angular momentum preparation

Prepare $|j = 1/2, m = 1/2\rangle$ in the coupled p manifold. Predict m_l outcomes, the spin expectation and the sum $\langle L_z + S_z \rangle$. Is either product component a state with definite J^2?

▼ **Work through the solution**

The probabilities are $1/3$ for $m_l = 0$ with spin up, $2/3$ for $m_l = 1$ with spin down. Thus $\langle L_z \rangle = 2\hbar/3$, $\langle S_z \rangle = -\hbar/6$ and the sum is $\hbar/2$. The product vectors in this two-dimensional plane are each superpositions of both j's, as can be checked by inverting the two orthogonal combinations. A joint measurement of orbital and spin labels can therefore disturb a previously exact J^2. The total m remains fixed in every outcome, stronger than merely having the correct mean.

Determine a spin–orbit interval without an angular integral

Within $l = 1, s = 1/2$, take $H' = A\,\mathbf{L}\cdot\mathbf{S}/\hbar^2$ with $A/h = 8$ GHz. Find level shifts, degeneracies and their degeneracy-weighted centre. Check the result from a product-basis trace.

▼ **Work through the solution**

Use $2\mathbf{L}\cdot\mathbf{S} = J^2 - L^2 - S^2$. For $j = 3/2$ the factor is $+1/2$; for $j = 1/2$ it is -1. Shifts are $+4$ and -8 GHz, with degeneracies four and two; their interval is 12 GHz and their weighted centre is zero. The product-basis trace of each $L_i S_i$ is $\operatorname{tr} L_i \, \operatorname{tr} S_i = 0$, which predicts that centre independently. If a measured centroid is weighted by optical intensity instead of counting every substate once, it need not equal this trace centre.

Distinguish resolved lines from a coherent pulse

A short $q = 0$ pulse prepares the product final state $|0, \uparrow\rangle$ above. If its spectrum covers both j levels coherently, find the later return probability to that same product state for a level angular-frequency splitting Δ. What would an incoherent population mixture give?

▼ **Work through the solution**

The return amplitude is $(2/3)e^{-iE_{3/2}t/\hbar} + (1/3)e^{-iE_{1/2}t/\hbar}$. Its squared modulus is $5/9 + (4/9)\cos\Delta t$, between $1/9$ and 1. The corresponding incoherent mixture gives a constant $5/9$. At $t = 0$ both constructions share the same energy-measurement probabilities, yet only the coherent preparation returns with certainty to the original product state. Temporal resolution, controlled pulse phase and bandwidth sufficient to cover both

lines are part of the proposal. If the emitted or absorbed field records distinguishable frequencies, the observed beat can be suppressed.

Further reading

- Fließbach · Operatorenmethode, angular momentum algebra
- Cambridge · atoms in electromagnetic fields and electric dipole selection rules
- Course investigation · angular momentum and atoms

Quantum states, dynamics and measurement · 5. From a Coulomb orbit to the fine-structure multiplet

Separate the hydrogenic Schrödinger equation, identify the origin of its large degeneracy and evaluate the three dominant relativistic corrections. The cancellation among them predicts an equality that spectroscopy can then test.

Separate centre-of-mass and relative motion

A light particle of mass m_e and charge $-e$ interacts with a pointlike nucleus $+Ze$ of mass M. Use centre of mass $\mathbf{R} = (m_e\mathbf{r}_e + M\mathbf{r}_N)/(m_e + M)$ and separation $\mathbf{r} = \mathbf{r}_e - \mathbf{r}_N$. Expanding the sum of the two kinetic energies cancels all cross terms and gives a free centre-of-mass kinetic energy plus relative energy $p^2/(2\mu)$, with reduced mass $\mu = m_eM/(m_e + M)$. The relative potential is $-k/r,\ k = Ze^2/(4\pi\epsilon_0)$. Its rotational symmetry implies $[H, L_i] = 0$. Use a spherical harmonic Y_{lm} and write $\psi = u(r)Y_{lm}/r$, so the probability in a thin radial shell is $|u(r)|^2dr$ for angular harmonics of unit norm.

Substituting in the Laplacian separates the radial equation

$$-\frac{\hbar^2}{2\mu}u'' + \left[\frac{\hbar^2 l(l+1)}{2\mu r^2} - \frac{k}{r}\right]u = Eu.$$

For the regular three-dimensional solution $u(0) = 0$ and a bound state also decays at infinity. The centrifugal term is an effective energy associated with angular variation, not a separate force added by hand. For $E < 0$ define $\kappa = \sqrt{-2\mu E}/\hbar,\ \rho = 2\kappa r$, and $a = \hbar^2/(\mu k)$. The small-radius behaviour of u is r^{l+1}; the large-radius tail contains $e^{-\kappa r}$. Pull both out by $u = \rho^{l+1}e^{-\rho/2}v(\rho)$. Straight differentiation leaves

$$\rho v'' + (2l + 2 - \rho)v' + (1/\kappa a - l - 1)v = 0.$$

For $v = \sum c_j\rho^j$, coefficient matching gives

$$c_{j+1}/c_j = (j + l + 1 - 1/\kappa a)/[(j + 1)(j + 2l + 2)].$$

The regular series must terminate to avoid an exponential asymptotic part that defeats the decaying prefactor. Thus $1/(\kappa a) = n_r + l + 1 = n$, and $E_n = -\mu k^2/(2\hbar^2 n^2)$. The normalised radial polynomials are associated Laguerre functions. This gives $l = 0, \ldots, n - 1$ and $\sum_l(2l + 1) = n^2$ states before including spin. Rotational symmetry guarantees the degeneracy in m, but a different central potential can split levels of different l. The additional equality is specific to the Coulomb potential.

Calculate the radial moments needed for a correction

To isolate the leading relativistic corrections now take an infinitely heavy nucleus and set $m = m_e$, $a = \hbar^2/(mk)$. Including reduced mass in the leading energy is straightforward; a complete recoil correction to the fine structure requires additional interactions. A useful identity is $\langle 1/r\rangle = 1/(an^2)$, obtained by differentiating E_n with respect to k. The derivative of an eigenvalue equals the expectation of $\partial H/\partial k = -1/r$ for the normalised eigenstate, which follows by differentiating $H|\psi\rangle = E|\psi\rangle$ and cancelling terms using its adjoint. The same method applied to the radial centrifugal parameter yields $\langle r^{-2}\rangle = 1/[a^2n^3(l + 1/2)]$ at fixed radial node count. Treat l as a continuously varied radial parameter during differentiation, then evaluate at its physical integer value.

For $l > 0$, the stationary radial force identity $\langle dV_{\mathrm{eff}}/dr \rangle = 0$ with vanishing endpoint term gives

$$\langle r^{-3} \rangle = 1/[a^3 n^3 l(l+1/2)(l+1)].$$

Here V_{eff} is the potential in the radial equation. The step uses sufficient behaviour of u, u' at the origin. For $l = 0$, the integral of r^{-3} diverges in an s state and the neglected endpoint term need not vanish. Substituting $l = 0$ formally into this expression would therefore be invalid. The contact density of an s state needed below is $|\psi_{n00}(0)|^2 = 1/(\pi a^3 n^3)$, while states with $l > 0$ vanish at the origin.

Track the kinetic, spin–orbit and contact terms separately

The small parameter is $Z\alpha = k/(\hbar c)$, with α the fine-structure constant for $Z = 1$. Expand the relativistic dispersion about small momentum after subtracting rest energy: $p^2/(2m) - p^4/(8m^3c^2) + \cdots$. For an energy eigenfunction $p^2\psi = 2m(E_n + k/r)\psi$, so $\langle p^4 \rangle = 4m^2\langle (E_n + k/r)^2 \rangle$ follows by taking the squared norm of $p^2\psi$. Squaring the two operators as if p^2 and $1/r$ commute is unnecessary and would not be an operator identity. Inserting the first two radial moments gives

$$\Delta E_{\mathrm{kin}} = -C_n \left[\frac{1}{l+1/2} - \frac{3}{4n} \right],$$

$$C_n = \frac{mc^2(Z\alpha)^4}{2n^3}.$$

The Pauli expansion of the Dirac Hamiltonian adds $H_{\mathrm{so}} = (1/[2m^2c^2r])V'(r)\mathbf{L}\cdot\mathbf{S}$ and the Darwin contact $H_D = \hbar^2\nabla^2 V/(8m^2c^2)$. These are effective interactions through order $(Z\alpha)^4$; a purely nonrelativistic scalar equation does not generate the spin or contact term. For $V = -k/r$, $V' = k/r^2$ and $\nabla^2 V = 4\pi k\delta^3(\mathbf{r})$. Thus the contact shifts only s states by $+C_n$. For $l > 0$ insert the third inverse radial moment and the angular identity $\mathbf{L}\cdot\mathbf{S} = \hbar^2[j(j+1) - l(l+1) - 3/4]/2$. For the branch $j = l + 1/2$ the bracket is l; for $j = l - 1/2$ it is $-(l+1)$. Reducing fractions then gives

$$\Delta E_{\mathrm{so}} = C_n[1/(l+1/2) - 1/(j+1/2)].$$

For $l = 0$ the spin–orbit operator vanishes and the Darwin term supplies the required difference. The sum in either case is

$$\Delta E_{nj} = -C_n \left[\frac{1}{j+1/2} - \frac{3}{4n} \right].$$

At fixed n the answer depends on j but no longer separately on l. The cancellation cannot be inferred by merely attaching a phenomenological $\mathbf{L}\cdot\mathbf{S}$ splitting to a nonrelativistic p state. In particular it predicts identical energies for $2s_{1/2}$ and $2p_{1/2}$ in this approximation, while $2p_{3/2}$ lies higher. Radiative, recoil and finite-nuclear-size contributions can resolve some of these residual equalities. An experimentally resolved residual is evidence that the specified Hamiltonian omitted a physical term; it does not remove the cancellation inside that Hamiltonian.

Compare a small energy correction with a resolvable line

For hydrogen with $mc^2 \simeq 511\,000$ eV and $\alpha \simeq 1/137.04$, the $2p_{3/2}$–$2p_{1/2}$ interval is $C_2/2 \simeq 4.53 \times 10^{-5}$ eV or 10.95 GHz. Compare with the leading $n = 2$ binding energy 3.40 eV. A small correction to total binding can give a prominent splitting when a measurement subtracts nearly equal large energies. Resolving it in a Fourier-limited time-domain experiment requires a coherent observation time at least of order the inverse frequency difference, here around 0.09 ns. Actual linewidths, lifetimes and pulse shapes affect how two spectral components can be distinguished.

A useful numerical audit is to fit the radial spectrum with three successive physical models. First solve the nonrelativistic radial boundary problem with the correct near-origin power and check that 2s and 2p converge toward the same energy. Then evaluate expectation values of the three corrections on those numerical states. Finally combine terms. Converged binding energies alone do not guarantee a converged contact density, which samples one spatial endpoint. A coarse radial grid can imitate an anomalous splitting between states which the analytical correction predicts to be equal. Check the origin slope of u_{n0}, not only its norm or far tail.

Find where the electron is most likely to be detected

Confusion between spatial density and radial probability can lead to very different-looking descriptions of the same atom. For a hydrogenic ground state the normalised position-space wavefunction is $\psi_{100} = (\pi a^3)^{-1/2} e^{-r/a}$. Its density is largest at the centre, but a thin spherical shell at radius r contains probability $P(r)dr = 4r^2 e^{-2r/a} dr/a^3$. The shell volume tends to zero at the origin. Differentiating $\ln P$ gives $2/r - 2/a = 0$, so the most likely shell radius is a. The mean radius $\int rP(r)dr = 3a/2$ is different again because the shell distribution has a long tail. Neither number alone gives the probability that a finite spatial detector clicks.

Integrating the shell probability between zero and R gives

$$\Pr(r < R) = 1 - e^{-2R/a}\left(1 + \frac{2R}{a} + \frac{2R^2}{a^2}\right).$$

At the most likely shell radius $R = a$, the probability inside is $1 - 5e^{-2} = 0.3233$, substantially less than one half. At $R = 2a$ it is

$$1 - 13e^{-4} = 0.7619.$$

Taylor expansion at a small radius $R \ll a$ gives $4(R/a)^3/3$, equal to the ball volume times central density to leading order. This is the local scaling underlying finite-nucleus corrections to a contact-sensitive observable. Using the most likely radius as though the electron followed a thin circular orbit would give none of these cumulative probabilities, even though the same Coulomb length scale appears in the Bohr model.

For $2s$ and $2p$, both states have principal index two and the same leading energy. The $2s$ radial amplitude is proportional to $(2 - r/a)e^{-r/(2a)}$; its node at $2a$ separates regions of opposite amplitude sign. The $2p$ radial amplitude is proportional to $re^{-r/(2a)}$, with no positive-radius radial node. At short distance, the probability inside R scales as R^3 for a nonzero s amplitude and as R^5 for the p wave. A sufficiently localised weak perturbation near the nucleus can therefore split the nonrelativistic degeneracy even if it is rotationally symmetric. Rotational symmetry still protects equal shifts within the m multiplet. A radial perturbation and a field selecting a direction leave different patterns of protected degeneracy.

Diagonalise a degeneracy before expanding a weak electric field

The $n = 2$ Coulomb degeneracy creates a striking response to a weak uniform electric field along z. In the purely nonrelativistic model with spin ignored the relevant spatial states are $2s,\ 2p_0,\ 2p_{\pm1}$. The perturbation for an electron in electric field $\mathcal{E}\hat{z}$ can be written $e\mathcal{E}z$ up to an additive constant. Each unperturbed parity eigenstate has zero diagonal expectation of this odd operator. Applying ordinary nondegenerate first-order perturbation theory separately to all four states would therefore predict zero for every shift. But first-order perturbation theory in a degenerate manifold requires diagonalising the projected perturbation, including off-diagonal amplitudes.

The $z = r\cos\theta$ operator has $q = 0$ and changes parity, so within $n = 2$ it couples only $2s$ and $2p_0$. Its angular matrix element between Y_{00} and Y_{10} is $1/\sqrt{3}$. In a conventional phase convention the radial functions are $R_{20} = (2 - r/a)e^{-r/(2a)}/(2\sqrt{2}\,a^{3/2})$ and $R_{21} = (r/a)e^{-r/(2a)}/(2\sqrt{6}\,a^{3/2})$. The remaining radial integral is $\int r^3 R_{20}R_{21}\,dr$. Setting $s = r/a$ gives

$$a(2\cdot 4! - 5!)/(8\sqrt{3}) = -3\sqrt{3}a.$$

Multiplying the angular factor gives $-3a$. Rephasing one basis state would change this sign but leave all level shifts invariant.

The projected matrix has zero diagonal and off-diagonal $-3ea\mathcal{E}$. Its symmetric and antisymmetric combinations have linear shifts $\mp 3ea\mathcal{E}$, while $2p_{\pm1}$ remain unshifted to this order. The new states have definite first-order dipoles along opposite directions even though their zero-field basis states had zero dipoles. For $a = 5.29\times10^{-11}$ m and $\mathcal{E} = 1000$ V/m, an individual linear shift has magnitude 1.59×10^{-7} eV, corresponding to 38.4 MHz. This field is weak compared with the gross level separation, but whether the degenerate result is accurate in a real atom also depends on much smaller fine, hyperfine and radiative splittings. “Weak field” must specify the comparison energy.

A more precise two-state model retains a zero-field splitting Δ between the coupled states. After subtracting their mean the Hamiltonian is $(\Delta/2)\sigma_z + d\mathcal{E}\sigma_x$, with $d = 3ea$ in magnitude. The energies are $\pm\sqrt{(\Delta/2)^2 + d^2\mathcal{E}^2}$. At $d|\mathcal{E}| \ll |\Delta|$ the shift from each initial level is quadratic, $\pm d^2\mathcal{E}^2/|\Delta|$, while at $d|\mathcal{E}| \gg |\Delta|$ the slopes approach the linear-degenerate prediction. An observed crossover distinguishes a small unresolved internal splitting from a strictly zero projected dipole matrix element. Tracking energies over both positive and negative fields also tests for a stray dc field shifting the centre of the response.

Estimate a finite-size correction from the potential itself

The contact term derived from the Pauli expansion is not the same physical effect as giving the nucleus a finite charge radius. A simple electrostatic model separates the two scalings. Replace a point nucleus by a uniformly charged sphere of radius R_N. Inside it the electron potential becomes

$$V_{\mathrm{sph}}(r) = -k[3 - (r/R_N)^2]/(2R_N);$$

outside it is unchanged at $-k/r$. The perturbation $\delta V = k[1/r - 3/(2R_N) + r^2/(2R_N^3)]$ is positive for interior points and zero at the surface. The smoothed potential is less attractive than the point law near the centre.

For an s state and $R_N \ll a$, the unperturbed density varies little across the tiny sphere. First-order perturbation theory therefore gives

$$\delta E \simeq 4\pi k |\psi(0)|^2 \int_0^{R_N} \left(r - \frac{3r^2}{2R_N} + \frac{r^4}{2R_N^3} \right) dr$$

$$\frac{2\pi}{5}\, k |\psi(0)|^2 R_N^2 .$$

The first two radial integrals cancel and leave the $R_N^2/10$ contribution before multiplication by 4π. Equivalently write $\delta E = (2\pi/3) k |\psi(0)|^2 \langle r_N^2 \rangle$, since the uniform sphere has mean-square charge radius $3R_N^2/5$. The result is a positive shift, decreasing binding. Corrections to the constant-density approximation involve additional powers of R_N/a. A different microscopic charge distribution changes higher moments, while the leading small-size shift depends on its mean-square charge radius.

For a model $R_N = 1.0$ fm and the ordinary electronic hydrogen length $a = 5.29 \times 10^{-11}$ m, $R_N/a = 1.89 \times 10^{-5}$. For the ground state compare with $|E_1| = k/(2a)$ and $|\psi(0)|^2 = 1/(\pi a^3)$. The relative shift is $(4/5)(R_N/a)^2 = 2.86 \times 10^{-10}$, equivalent to roughly 3.9×10^{-9} eV. This illustrative uniform charge radius is input data of the model; do not equate it directly to a measured rms radius. For a heavier orbiting lepton a shorter Bohr radius increases the fractional sensitivity strongly. A complete precision spectrum would additionally include radiative and recoil effects and a realistic nuclear response. The present derivation isolates how one local perturbation scales and why s levels are especially sensitive.

Use dilation to audit stationary Coulomb energies numerically

A variational scaling argument gives a powerful independent check on a computed wavefunction. For any normalised trial state ψ, define the normalised rescaling $\psi_\lambda(\mathbf{r}) = \lambda^{3/2}\psi(\lambda \mathbf{r})$. The Coulomb potential is homogeneous of degree minus one, and the gradient acquires a factor λ. Its energy expectation is therefore

$$E(\lambda) = \lambda^2 \langle T \rangle + \lambda \langle V \rangle .$$

For an exact stationary bound state the energy functional is stationary with respect to norm-preserving variations, giving $2\langle T \rangle + \langle V \rangle = 0$ at $\lambda = 1$. Hence $\langle T \rangle = -E,\ \langle V \rangle = 2E$, consistent with the independently derived inverse-radius mean. The argument requires appropriate normalisation and boundary behaviour so the scaled trial function is allowed.

Suppose a numerical $1s$ reconstruction yields $\langle T \rangle = 13.2$ eV and $\langle V \rangle = -26.7$ eV. Its energy is -13.5 eV, yet its virial residual is -0.3 eV. Along the scale family the minimum occurs at $\lambda_* = -\langle V \rangle / (2\langle T \rangle) = 1.0114$, a slightly more concentrated state. The resulting energy reduction is only about 0.0017 eV, much smaller than the change in individual kinetic and potential terms. A small total-energy error can coexist with a noticeably incorrect radial scale because first-order energy errors cancel near a stationary solution. Monitoring separate expectation values, the virial residual and a contact-sensitive observable provides a more informative test of a numerical atomic calculation than reporting the total binding energy alone.

Distinguish orbital density from radial probability

Compare radial states at the same principal quantum number. Inspect their nodes before considering which additional interactions can lift their degeneracy.

Open this experiment online

Try the calculation

Nuclear charge Z

1

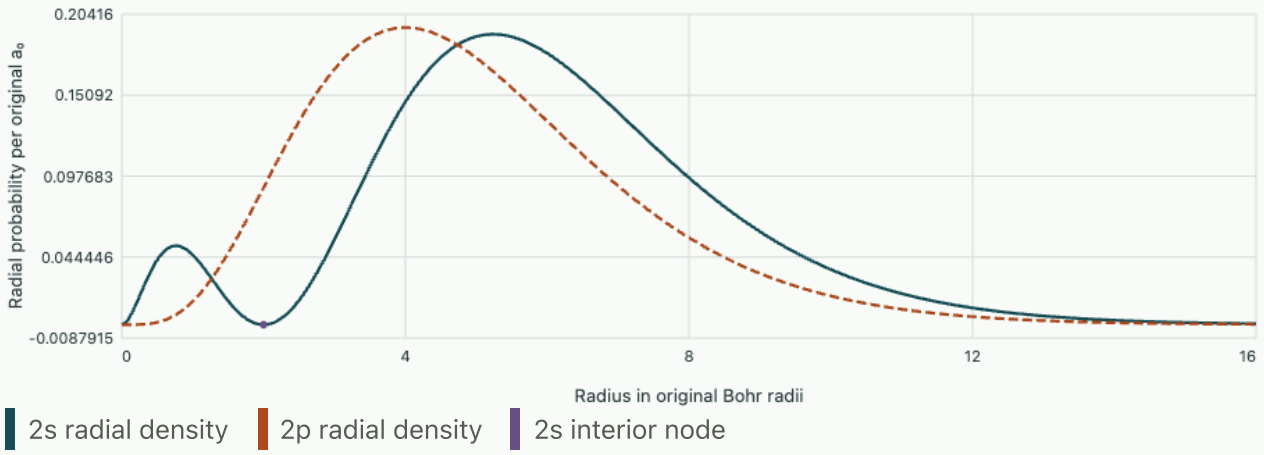

The 2s node is at 2 original Bohr radii. The most probable 2p radius is 4 original Bohr radii.

Appraise a finite-grid Coulomb spectrum from convergence and degeneracy

To compute the radial Schrödinger eigenvalues directly, work in dimensionless radius $s = r/a$, energy unit $\hbar^2/(2ma^2)$, and write $[-d^2/ds^2 + l(l+1)/s^2 - 2/s]u = \varepsilon u$. On a uniform grid $s_j = jh$, use $u_0 = 0$ and a far boundary $u_{N+1} = 0$. The diagonal approximation is $2/h^2 + l(l+1)/s_j^2 - 2/s_j$ and its neighbours are $-1/h^2$. All potential terms are evaluated at positive interior grid points. Inserting an arbitrary finite value for $1/s$ at the origin and treating it as an ordinary interior point changes the boundary approximation. A tridiagonal symmetric eigensolver returns orthogonal numerical vectors; physical radial norm is $h\sum_j |u_j|^2$, so normalise appropriately before reconstructing spatial densities.

For $l = 0$ on an infinite grid the lowest mode permits an unusually informative exact lattice check. Try $u_j = jq^j$ for $0 < q < 1$. The finite-difference kinetic row divided by this trial value is $[2 - q - q^{-1}]/h^2 - (q - q^{-1})/(jh^2)$. Adding potential $-2/(jh)$, the $1/j$ coefficient vanishes when $q^{-1} - q = 2h$, or $q = \sqrt{1+h^2} - h$. The eigenvalue is then

$$\varepsilon_h = [2 - 2\sqrt{1+h^2}]/h^2 = -2/[1 + \sqrt{1+h^2}].$$

It tends to the exact continuum ground value -1 as $-1 + h^2/4 + O(h^4)$. This analytic discrete answer separates algebraic solver correctness from discretisation error for the exact same potential and boundary treatment.

At $h = 0.20$ the infinite-grid ground value is approximately -0.99020; at $h = 0.10$ it is -0.99751. Their errors differ by nearly a factor of four as predicted. A finite far wall shifts both energies further; repeat first at fixed h and increasing boundary radius to isolate that effect. An excited $n = 2$ state has a more extended tail and requires a larger radius than the ground state for comparable truncation bias. Solving $l = 0$ and $l = 1$ on a small common box can yield a spurious nonrelativistic 2s–2p splitting. Check separately that both energies tend to $-1/4$ before interpreting a much smaller physical fine splitting. Subtracting independently computed energies amplifies absolute numerical biases even if each number appears visually close to its expected value.

The distinctive near-origin behaviour supplies a second check. Since $\psi_{100}(0)$ is proportional to the slope of the reduced radial function at zero, a first forward difference u_1/h typically incurs a first-order error even when the energy error is second order. For the lattice shape above $u_j \propto jq^j$, the first two positive grid values can be used to estimate the endpoint derivative by $(4u_1 - u_2)/(2h)$. Test its convergence under the physical radial normalisation. For a contact perturbation, squaring an endpoint derivative doubles its small relative error to leading order. This is why an atomic calculation should include an observable-specific refinement criterion in addition to an energy tolerance.

Use an eigenvalue benchmark to separate numerical and physical shifts

A report on the infinite-grid hydrogenic ground state quotes energy -0.99750 at $h = 0.10$ and attributes the difference from -1 to a short-range interaction. Predict the lattice-only change on refinement to $h = 0.05$, and propose a test that could reveal a real interaction independently.

▼ **Derive and check the result**

The lattice expression gives $-2/[1 + \sqrt{1.0025}] \simeq -0.999376$ at 0.05, so the discretisation alone accounts for the original difference at the stated accuracy. A local physical correction with fixed continuum strength should converge toward a fixed nonzero shift as $h \to 0$, while the lattice bias vanishes approximately as $h^2/4$. Compute control and perturbed energies on exactly the same grids and compare their difference under refinement, with consistent treatment of the perturbation's physical range. An unresolved delta perturbation represented as a fixed diagonal value on one grid point changes its effective continuum strength upon refinement and would produce another artificial shift. Cross-check with the first-order expectation of a resolved physical potential or with the independently converged contact density.

Use an isotope pair to cancel a common transition-energy reference

High-precision spectroscopy often obtains its most reliable insight from combinations of frequencies rather than absolute energies. Consider two hydrogenic model isotopes with the same nuclear charge, different reduced masses μ_1, μ_2 and different mean-square charge radii R_1^2, R_2^2. For a specified transition the leading point-Coulomb frequency can be written $C\mu$, since its n-dependent dimensionless numbers are fixed. A local nuclear-size perturbation shifts each participating s level by a factor proportional to $|\psi_{ns}(0)|^2 R_i^2 \propto \mu_i^3 R_i^2/n^3$, while a p level receives a much smaller finite-size correction in the same limit. The different powers of reduced mass and radial quantum number permit differences of transitions to cancel selected terms.

For illustration compare a 1s–2s optical frequency with a 1s–2p frequency within one model isotope, assuming the fine, radiative and differing upper-state structures have been independently accounted for in the desired order. Subtracting these frequencies cancels the entire common lower 1s energy and the gross equal 2s/2p Coulomb energy, leaving the upper 2s–2p differences. If instead one sought the large 1s contact contribution, subtracting these same two lines would remove precisely the effect of interest. An apparently convenient common reference can therefore eliminate physical signal as well as noise. Write the coefficient of each contributing level explicitly before selecting a frequency combination.

Suppose two independently fitted line frequencies ν_a, ν_b each have a 10 MHz common reference error and independent residual standard errors of 0.30 and 0.40 MHz. The covariance from the shared reference is $100\,\mathrm{MHz}^2$. The difference has variance $0.30^2 + 0.40^2 = 0.25\,\mathrm{MHz}^2$, hence 0.50 MHz standard error. Treating both total errors as independent instead would suggest roughly 14 MHz. Conversely, an error in a calibration scale proportional to frequency leaves a residual proportional to the frequency difference; its cancellation is only approximate. Reporting the provenance of uncertainties and their covariance can matter more than recording an extra digit of each absolute transition.

Recover a contact shift from two excited-state probes

In a diagnostic model a local perturbation gives an unknown positive 1s shift D, an ns shift D/n^3, and zero shifts of p levels. Relative to a matched reference, how do 1s–2s, 1s–2p and 1s–3s frequencies change? Which experimental difference removes the unknown common lower-state shift yet recovers D?

▼ Derive and check the result

The changes are $(-7D/8)/h,\ -D/h,\ (-26D/27)/h$, respectively. Subtracting the 1s–2p shift from the 1s–2s shift gives $D/(8h)$, since the common lower-state contribution cancels. For a measured difference of 0.90 ± 0.05 MHz the implied $D/h = 7.2 \pm 0.4$ MHz. The 1s–3s shift should then be -6.93 MHz in this stipulated model. If unavailable corrections also differ among the upper states, that value is not an absolute nuclear-radius extraction. The third line gives an independent test of the $1/n^3$ contact scaling before attributing all measured differences to the single parameter D.

Tutorial · measure degeneracy through independent calculations

Verify two inverse moments directly in the $2p$ state

Use $R_{21}(r) = re^{-r/(2a)}/(2\sqrt{6}\,a^{5/2})$ and normalised Y_{1m}. Evaluate $\langle r^{-2}\rangle$, $\langle r^{-3}\rangle$ without parameter differentiation.

▼ Work through the solution

The radial measure is r^2dr, so after inserting R the general moment is $(24a^5)^{-1}\int_0^\infty r^{4-q}e^{-r/a}dr$ for r^{-q}. For $q=2$ the integral is $2a^3$, giving $1/(12a^2)$. For $q=3$ it is a^2, giving $1/(24a^3)$. Both agree with the formulas for $n=2, l=1$. Dropping the radial Jacobian would change both powers and could falsely suggest singular behaviour. The result is independent of m because the operator is rotationally scalar.

Calculate the complete $n = 2$ shifts in units of C_2

Resolve the kinetic, spin–orbit and Darwin contributions for $2s_{1/2}$, $2p_{1/2}$ and $2p_{3/2}$. Which comparison isolates the effect of the spin–orbit term cleanly at this order?

▼ **Work through the solution**

For s, kinetic gives $-[2 - 3/8] = -13/8$ and Darwin $+1$, summing to $-5/8$. For both p states the kinetic value is

$$-2/3 + 3/8 = -7/24.$$

The spin–orbit corrections are $2/3 - 1 = -1/3$ for $j = 1/2$ and $2/3 - 1/2 = 1/6$ for $j = 3/2$, giving $-5/8$ and $-1/8$. Subtracting the two p levels removes their shared kinetic term and zero contact term. Their interval is $C_2/2$. Comparing 2s with either p instead involves two different radial wavefunctions and does not isolate the same correction.

Separate a reduced-mass shift from fine structure

Consider two otherwise identical hydrogenic nuclei with $M/m = 1836$ and 3672. Predict the ratio of gross transition frequencies and compare its fractional change with the order of a fine-structure correction. Which common interpretation would conflate them?

▼ **Work through the solution**

Gross transition frequencies scale as reduced mass $\mu = m/(1 + m/M)$. The heavier-to-lighter ratio is $(1 + 1/1836)/(1 + 1/3672) = 1.000272$ approximately. The order of a relative fine-structure correction for $Z = 1$ is $\alpha^2 \simeq 5.3 \times 10^{-5}$ with state-dependent coefficients. A change of isotope can shift a gross line by more than a particular resolved fine splitting, even with identical nuclear charge. Centroid measurements across isotopes and intervals within one multiplet are distinct observables. Recoil and nuclear magnetism require further inputs for precision predictions.

Further reading

- Fließbach · Dreidimensionale Probleme
- Cambridge · hydrogen, helium and fine structure
- Course investigation · perturbations and level mixing

Quantum states, dynamics and measurement · 6. When a magnetic field changes the quantum numbers of an atom

Weak-field Zeeman slopes label total-angular-momentum multiplets. At stronger fields those slopes curve and intensities migrate. Solve an entire orbital-plus-spin manifold and infer coupling constants from measurable line positions and polarisation.

Choose the Hamiltonian appropriate to a resolved orbital manifold

Suppose an isolated atomic p manifold with $l = 1, s = 1/2$ is sufficiently far from other orbital or principal levels to treat only its six substates. Take z along a uniform external magnetic field B. Work to first order in the electronic Zeeman operator and model spin–orbit coupling by the constant $A > 0$. With $b = \mu_B B$ and electron spin factor approximated by two,

$$H = \frac{A}{\hbar^2}\mathbf{L}\cdot\mathbf{S} + \frac{b}{\hbar}(L_z + 2S_z).$$

Both A and b are energies. The sign of the second term is positive for the electron because magnetic moment is opposite to angular momentum and field energy is minus moment dot field. Here $\mu_B = e\hbar/(2m_e) > 0$. The orbital manifold approximation requires magnetic splittings small compared with gaps to omitted levels. Neglected terms can include the difference of the actual spin factor from two, hyperfine interactions and the diamagnetic B^2 term. An “intermediate field” in this discussion means b comparable to A; it need not exceed any other electronic binding scale.

At zero field the coupled levels $j = 1/2, 3/2$ have energies $-A, +A/2$. At nonzero field the projection $m = m_l + m_s$ is still conserved because $[H, J_z] = 0$. For intermediate fields J^2 need not commute with the Zeeman term. Consequently the six-state problem reduces to two one-dimensional and two two-dimensional invariant blocks, even when the label j loses its exact meaning. Following a state by magnetic quantum number and overlap with a neighbouring field value is generally more reliable than sorting all computed energies from low to high. Crossings of different m can reorder a sorted list without exchanging the underlying states.

Derive the Landé slope from angular momentum projection

In a state with definite j, l, s, m, expectation of any vector operator proportional in direction to $\mathbf{J}$ within its irreducible multiplet can be obtained from its scalar projection on $\mathbf{J}$. For spin,

$$\frac{\langle S_z\rangle}{\hbar} = m\,\frac{j(j+1) + s(s+1) - l(l+1)}{2j(j+1)}.$$

One way to obtain the projection numerator is $2\mathbf{S}\cdot\mathbf{J} = J^2 + S^2 - L^2$. Since $L_z + 2S_z = J_z + S_z$, the energy slope for weak fields is $dE/db = g_j m$ with $g_j = 1 + [j(j+1) + s(s+1) - l(l+1)]/[2j(j+1)]$. A state $j = 0$ has vanishing linear vector moment and should be treated separately instead of dividing by zero. Within the p manifold the $j = 1/2$ factor is $2/3$ and the $j = 3/2$ factor $4/3$. For $m = +1/2$ these correspond to slopes $1/3$ and $2/3$ in energy units of b. Their sum is exactly one, a trace constraint that remains valid for the sum of the two eigenvalues in that block at every field.

If $A/h = 8.00$ GHz as in the previous angular exercise, the zero-field interval is 12.00 GHz. At $B = 20.0$ mT, $b/h \simeq 0.280$ GHz using $\mu_B/h = 13.996$ GHz/T. The two positive-m shifts to linear order are $+0.0933$ and $+0.1866$ GHz. Actual optical frequencies measure final minus initial energies; an initial $s_{1/2}$ level has its own Zeeman slope. Omitting it is a common reason a slope inferred from a line scan seems inconsistent with the expected Landé factor.

Diagonalise the mixed levels across the entire crossover

In the $m = +1/2$ product basis $(|0,\uparrow\rangle, |1,\downarrow\rangle)$, expand $\mathbf{L}\cdot\mathbf{S} = L_zS_z + (L_+S_- + L_-S_+)/2$. The diagonal spin–orbit elements are $0, -A/2$; the off-diagonal element is $A/\sqrt{2}$. The Zeeman term is diagonal with entries $b, 0$. Hence

$$H_+(b) = \begin{pmatrix} b & A/\sqrt{2} \\ A/\sqrt{2} & -A/2 \end{pmatrix}, \qquad E_\pm(b) = \frac{b - A/2}{2} \pm \frac{1}{2}\sqrt{(b + A/2)^2 + 2A^2}.$$

At $b = 0$ the square root is $3A/2$, recovering the coupled levels. For the $m = -1/2$ basis $(|0,\downarrow\rangle, |-1,\uparrow\rangle)$, replace b by $-b$. The stretched levels $|1,\uparrow\rangle$ and $|-1,\downarrow\rangle$ remain product states at every field with energies $A/2 \pm 2b$. Counting all levels and their traces provides an independent check: spin–orbit and Zeeman each have zero trace in the full six-dimensional space. At zero field both slopes in the $+$ block are finite and positive. At large positive b the upper slope tends to one and the lower to zero, approaching their uncoupled product labels.

In the negative-m block the two uncoupled energies would cross at $b = A/2$. The coupling opens a minimum gap $\sqrt{2}A$. Within the model, levels of the same m avoid one another whenever this matrix element is nonzero. Crossings between different m can remain exact because the uniform z-directed Hamiltonian does not mix them. A tilted perturbation could turn some of these into avoided crossings. Such symmetry dependence matters when trying to infer a coupling from the smallest visible distance between any two lines.

Obtain an eigenvector weight from an energy derivative

Suppose an initial s state has spin up. A weak $q = 0$ electric-dipole drive prepares the product basis state $|0,\uparrow\rangle$ in the $m = +1/2$ block. Its squared overlap with an exact eigenstate determines the angular part of that optical line's strength. Rather than computing a mixing angle with a branch ambiguity, differentiate the exact energies. Since $\partial H_+/\partial b = |0,\uparrow\rangle\langle 0,\uparrow|$, the eigenvalue derivative equals precisely that population:

$$w_\pm = \frac{dE_\pm}{db} = \frac{1}{2}\left[1 \pm \frac{b + A/2}{\sqrt{(b + A/2)^2 + 2A^2}}\right].$$

The weights sum to one as required by completeness. At zero field they are $2/3$ and $1/3$. At large positive field the upper line takes almost all this polarisation strength. In the negative block an analogous derivative with respect to $-b$ gives the population of $|0,\downarrow\rangle$. At its avoided crossing each eigenstate has equal weights and the two corresponding $q = 0$ intensities are equal for a selected initial spin down. The same experiment can test slopes measured from frequencies and overlaps measured through intensities, provided detection and radial coupling are common to both lines.

Set $b = A$ in the positive block. The level separation becomes $A\sqrt{4.25} = 2.06155A$, while the upper product weight is 0.8638. A naive extrapolation of the weak-field $g_j m$ slopes would give an interval $1.5A + A/3 = 1.833A$. The discrepancy is governed by the spin–orbit mixing scale rather than failure of the magnetic Hamiltonian itself. If the frequency analyser resolves only a single intensity-weighted centroid, completeness still constrains the first moment:

$$\sum_{\pm} w_{\pm} E_{\pm} = \langle 0, \uparrow | H_{+} | 0, \uparrow \rangle = b.$$

Thus the centroid of unresolved lines in this ideal selected excitation can move linearly even while each constituent bends.

Resolve a second angular-momentum hierarchy with the same algebra

The p-manifold example coupled orbital and electronic spin. In an s-state atom the dominant remaining weak angular coupling may be between an electronic $J = 1/2$ and a nuclear $I = 1/2$. This is a different scale and a different physical pair of degrees of freedom. To illustrate how a state label changes with field, take

$$H = \frac{A_h}{\hbar^2} \mathbf{I} \cdot \mathbf{J} + \frac{B}{\hbar} (g_J \mu_B J_z - g_I \mu_N I_z)$$

with positive hyperfine constant A_h and a positive model nuclear g_I. The minus before the nuclear-magneton term reflects the assumed positive nuclear magnetic moment. Actual sign and g-factor depend on the isotope. Here both A_h and the products of magnetic moments and field have energy units. At zero field the electronic and nuclear angular momenta combine to $F = 0, 1$. Their scalar-product shifts are $-3A_h/4$ and $+A_h/4$, respectively. A rotation of the entire coupled system leaves this zero-field Hamiltonian invariant, but independent rotations of just one constituent generally do not.

Within $m_F = 0$ order the product basis by electronic magnetic sign as $|\uparrow_J, \downarrow_I\rangle, |\downarrow_J, \uparrow_I\rangle$. Both diagonal hyperfine matrix elements are $-A_h/4$, off-diagonal exchange is $A_h/2$, and field changes the diagonal difference by $\delta = (g_J \mu_B + g_I \mu_N) B$. Hence the central pair of energies is

$$E_{\pm} = -A_h/4 \pm \tfrac{1}{2} \sqrt{A_h^2 + \delta^2}.$$

The two stretched $m_F = \pm 1$ levels sit at $+A_h/4 \pm (g_J \mu_B - g_I \mu_N) B/2$. Their slopes contain a difference of electronic and nuclear moments, while the central splitting depends on their sum. This sign distinction can be determined directly by comparing the relevant product states: in the anti-aligned pair the sign change reverses one angular momentum along with its opposite magnetic sign. A sign convention imported from electron spin alone would give an erroneous nuclear contribution.

The transition between the two central states has gap $\sqrt{A_h^2 + \delta^2}$ and zero linear slope at $B = 0$. Its quadratic shift is $\delta^2/(2A_h)$ for $|\delta| \ll A_h$. Measuring at a point with zero first derivative reduces sensitivity to slow field fluctuations without eliminating it. Suppose $A_h/h = 1.00$ GHz and $(g_J \mu_B + g_I \mu_N)/h = 28.0$ GHz/T in an illustrative species model. At a 1.0 mT bias the central gap rises by about 0.392 MHz. A 1 µT field error at that bias gives a first-order frequency error of about 784 Hz using the derivative of the quadratic

expansion. At zero mean bias, fluctuations with rms 1 μT instead give an average quadratic shift about 0.392 Hz. Stabilising the field mean and reducing its variance address different terms in the measured spectral position.

The magnetic drive operator for a small oscillating field along z is itself proportional in this block to σ_z. At zero bias the hyperfine eigenstates are the symmetric and antisymmetric superpositions, so this drive couples them off-diagonally despite the central transition frequency being stationary against a small static bias. The diagonal derivative and the off-diagonal transition matrix element are distinct quantities. At large bias the eigenstates approach product states and the same longitudinal drive becomes mostly diagonal, reducing its ability to flip between them. Thus a transition can have both a vanishing first-order shift and a nonzero resonant driving strength, consistent with the same Hamiltonian. Measuring both properties probes the field dependence of eigenvectors as well as eigenvalues.

Infer a field distribution from widths and centres

Spectra of spatially extended samples often contain a distribution of static magnetic fields. Consider a single weak-field transition with frequency $\nu(B) = \nu_0 + sB$, an intrinsic Gaussian line standard deviation w_0, and a sample field distributed approximately Gaussian with mean B_0 and standard deviation σ_B. A sum over independent source regions convolves the intrinsic line with the frequency distribution generated by sB. The observed centre is $\nu_0 + sB_0$, and variance is $w_0^2 + s^2\sigma_B^2$. A single linewidth cannot determine both an intrinsic width and a field inhomogeneity. Two lines with different known magnetic slopes or one calibrated zero-slope reference can separate the contributions.

For example two transitions with slopes 7.0 and 14.0 GHz/T have measured standard deviations 0.250 and 0.400 MHz. Subtract the squared widths to eliminate w_0^2:

$$\sigma_B^2 = \frac{(0.400)^2 - (0.250)^2}{(14000)^2 - (7000)^2}\ \mathrm{T}^2,$$

using MHz/T in the denominator. The result is $\sigma_B = 25.8\,\mu\mathrm{T}$, and the inferred intrinsic width is

$$\sqrt{0.250^2 - (7000\sigma_B)^2} = 0.173\,\mathrm{MHz}.$$

A third line with slope near zero supplies an out-of-sample check. If lines have significantly different lifetimes or power broadening, the assumption of a common w_0 should be relaxed; otherwise one can spuriously infer a magnetic gradient or even a negative intrinsic variance.

For a curved avoided-crossing line expand $\nu(B) \simeq \nu(B_0) + s(B - B_0) + q(B - B_0)^2/2$. A symmetric distribution shifts the measured mean by $q\sigma_B^2/2$ even when the mean field is known exactly. Near a stationary point the leading field broadening can be strongly asymmetric because a squared Gaussian variable has a one-sided distribution around its minimum. Fitting a symmetric line at such a point can shift its apparent peak or centre according to the fitting convention. This effect is especially relevant when locating a minimum gap from spectra at several nominal fields. A model integrating intensities over a measured spatial field distribution can be better constrained than independent fits of Gaussian centres, because all lines share the same underlying field profile.

An independent geometric test tilts the weak sensing light relative to the magnetic axis. For equal populations, a truly scalar average over all polarisations and emitted directions is insensitive to this coordinate change, but a selected q transition changes strength according to the decomposition of the laboratory polarisation. If a peak attributed to frequency splitting mainly changes its fitted centre while unresolved component intensities redistribute, some of its apparent shift may be an intensity-weighted centroid effect. Record spectra at two orthogonal polarisations, fit the same physical energy parameters and allow the predicted angular strengths to vary. A common energy correction that survives the intensity change carries a different interpretation from a change in the mixture of unresolved substates.

Predict state transfer during a field sweep

The negative-m avoided crossing of the p manifold also poses a time-dependent question. Does a changing field keep the atom in the instantaneous energy eigenstate, or in approximately the same product state? In the local two-state description subtract the trace and write $H = (\Delta(t)/2)\sigma_z + g\sigma_x$, where the difference of diabatic diagonal energies is Δ and the minimal adiabatic gap is $2|g|$. Introduce a mixing angle $\tan(2\vartheta) = 2g/\Delta$, chosen continuously through the crossing. Differentiating at fixed g gives

$$|\dot{\vartheta}| = |g\dot{\Delta}|/(\Delta^2 + 4g^2).$$

The moving eigenbasis itself couples the two instantaneous energy states with an energy scale $\hbar|\dot{\vartheta}|$.

At closest approach the order-of-magnitude adiabatic criterion becomes $\hbar|\dot{\Delta}|/(8g^2) \ll 1$. A slow sweep keeps occupation on an energy branch, which changes its product character across the avoided crossing. A fast sweep can leave the original product state approximately unchanged, moving population from one instantaneous energy branch to the other. Both descriptions refer to actual state transfer; the basis in which “remaining in the same state” is assessed is part of the observable. The criterion here follows directly by comparing the nonadiabatic derivative coupling with the instantaneous level separation. It estimates a regime but does not by itself give a transition probability.

For a strictly linear detuning passing from large negative to large positive values with constant g, the Landau–Zener result gives the diabatic survival probability $P_D = \exp[-2\pi g^2/(\hbar|\dot{\Delta}|)]$. Its applicability requires asymptotic start and end detunings well outside the coupling region, coherent passage and the indicated linearity over the crossing. In ordinary-frequency units take $g/h = 20$ kHz and $|d(\Delta/h)/dt| = 1.0 \times 10^{10}\ \mathrm{Hz/s}$. The exponent is $4\pi^2(2.0 \times 10^4)^2/10^{10} = 1.58$, giving $P_D \simeq 0.206$. Doubling the sweep rate halves the exponent and raises survival to about 0.454. Confusing angular with ordinary frequency in this conversion loses a factor 2π, large enough to change the predicted population qualitatively.

A double sweep forms a useful control. On entering and later leaving the avoided crossing, the two amplitudes propagated along different energy branches can accumulate different phases and interfere. Treating the two passages as independent classical jumps would miss oscillations versus waiting time between them. Strong dephasing or averaging over a broad field distribution can remove the oscillations while retaining the one-pass populations. Distinguish these mechanisms by calibrating single-passage transfer first and then varying the waiting time over several periods at multiple sweep speeds. This extends a static magnetospectroscopy model to a specified dynamical experiment without inferring the full time dependence from a static level diagram alone.

Resolve the crossover from coupled to uncoupled angular momenta

Scan the magnetic field across the spin–orbit scale. Compare the exact two-state spectrum at fixed magnetic quantum number with the weak-field shifts.

Open this experiment online

Try the calculation

Spin–orbit coefficient A (μeV) **48**

Selected magnetic field (T) **0.6**

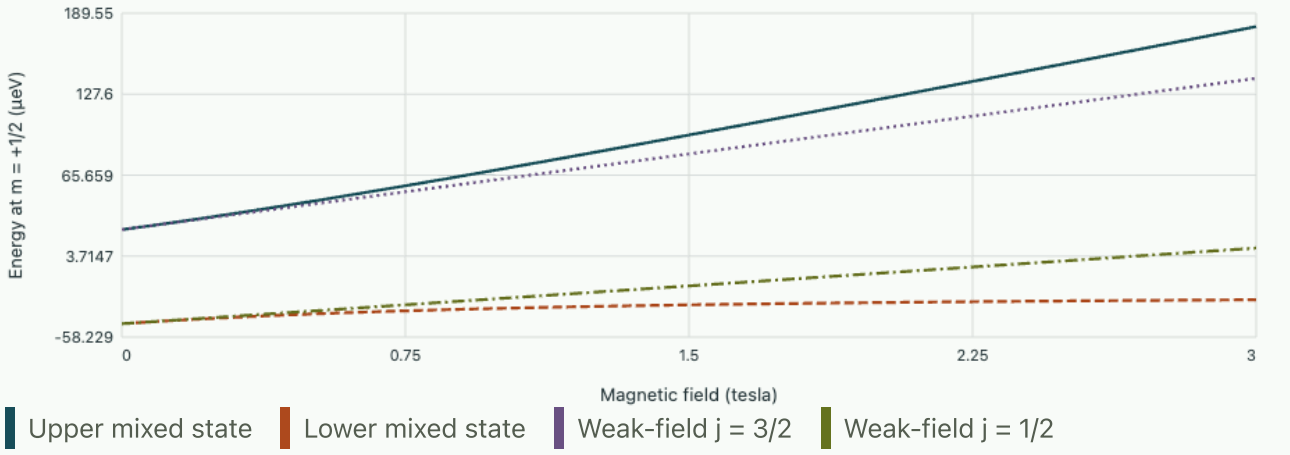

Selected energies -39.516 and 50.246 μeV. The ratio μB B/A is 0.72355.

Test the orientation of a magnetic sensor from resolved spin populations

The scalar energy spectrum of an isotropic spin-$1/2$ Zeeman Hamiltonian measures field magnitude. A transition energy by itself cannot determine the signed orientation of that field. Let $H = \hbar\gamma\mathbf{B}\cdot\boldsymbol{\sigma}/2$, with signed angular gyromagnetic ratio γ defined in the experiment. Its eigenvalues are $\pm\hbar|\gamma||\mathbf{B}|/2$. If instead one prepares spin up along a known laboratory z direction and observes its population after evolution time t, the rotation axis appears in the dynamics. Exponentiating $U = \cos(\omega t/2)I - i\sin(\omega t/2)\hat{\mathbf{B}}\cdot\boldsymbol{\sigma}$ for a consistent choice of $\omega, \hat{\mathbf{B}}$ gives

$$P_{\downarrow z}(t) = (1 - \hat{B}_z^2)\sin^2(\gamma|\mathbf{B}|t/2).$$

The oscillation frequency gives magnitude and its contrast gives the squared transverse fraction. It is insensitive to whether the field points slightly above or below a perpendicular plane. Incomplete preparation or pulse contrast can also reduce the observed oscillation amplitude. A calibration with a known transverse field tests that ambiguity before assigning every missing contrast to geometric tilt.

Suppose a resolved precession has angular frequency corresponding to 0.50 mT and a calibrated flip contrast of 0.36 for initial $+z$. Then $\sqrt{B_x^2 + B_y^2} = 0.30$ mT and $|B_z| = 0.40$ mT. Repeating for an initial $+x$ state with

calibrated contrast 0.64 yields $|B_x| = 0.30$ mT and thus $B_y = 0$ for these numbers. Neither experiment identifies the signs of B_x and B_z. To resolve them one can measure an early-time transverse quadrature with a phase reference. In Bloch notation $\dot{\mathbf{r}} = \gamma \mathbf{B} \times \mathbf{r}$, so the initially zero slope $\dot{r}_y = -\gamma B_x$ for $+z$ is sign sensitive. This derivative is independent of B_z at the initial instant, giving a direct check complementary to the energy splitting.

A second strategy adds a calibrated bias $\mathbf{b}_0$ in opposite directions and examines squared frequencies. Since $|\mathbf{B} + \mathbf{b}_0|^2 - |\mathbf{B} - \mathbf{b}_0|^2 = 4\mathbf{B} \cdot \mathbf{b}_0$, the difference isolates a signed projection even when the zero-field splitting reveals only magnitude. For a 0.20 mT known bias along z and a 0.40 mT signed z component the squared-field difference should be $0.32\,\mathrm{mT}^2$. Alternating the two bias signs across the measurement reduces slow field drift. A quadratic diamagnetic response, anisotropic g tensor or an uncalibrated bias axis can break this simple inversion; repeating in three independent directions tests the vector model and reveals which assumptions are being used.

Separate an anisotropic magnetic response from an unknown field direction

A crystal-fixed effective spin has diagonal $g = \mathrm{diag}(2.0, 2.0, 2.4)$. Derive the transition energy for a field rotated in the x, z plane by angle θ from z, and explain what an isotropic fit would infer.

▼ **Derive and check the result**

The Hamiltonian is $\mu_B \boldsymbol{\sigma} \cdot (g\mathbf{B})/2$. The splitting is $\mu_B B \sqrt{2.0^2 \sin^2\theta + 2.4^2 \cos^2\theta}$. At 45° the effective g is $\sqrt{4.88} = 2.209$; interpreting the measured splitting with constant $g = 2$ would overestimate field magnitude by about 10.5%. The spin also precesses around the effective axis $g\mathbf{B}$, whose polar angle satisfies

$$\tan\theta_{\mathrm{eff}} = (2.0/2.4)\tan\theta.$$

Measuring contrast from a calibrated spin preparation and rotating a field of known magnitude can therefore distinguish material anisotropy from a simple error in magnet alignment. One scalar spectral line at an unknown orientation cannot separate these effects.

Tutorial · reconstruct a magnetically mixed spectrum

Infer the coupling from an anticrossing

The two negative-m lines exhibit a minimum frequency separation of 11.31 GHz at 0.286 T. Infer A/h from each measurement. What calibration would be needed for a comparison of absolute optical line frequencies?

▼ **Work through the solution**

Minimum gap divided by h equals $\sqrt{2}A/h$, giving $A/h = 8.00$ GHz. The crossing condition $b = A/2$ gives

$$A/h = 2(13.996)(0.286) = 8.01\,\mathrm{GHz}.$$

Their agreement checks field calibration independently of a common frequency reference. Both final states must be viewed from the same initial state when using their frequency difference. Comparing absolute optical

frequencies also requires the zero-field transition energy and the field-dependent initial-level shift. Two lines originating in different initial states would have a field-dependent separation even if their final levels had a fixed interval.

Extract the polarisation weight using a finite difference

Near $b = A$, estimate the slope of the upper positive-m level using its exact values at $b/A = 0.95, 1.05$. Compare with the exact weight and predict the second line's strength.

▼ **Work through the solution**

In units of A the upper level is $(t - 0.5 + \sqrt{(t+0.5)^2 + 2})/2$ with $t = b/A$. Subtracting the values and dividing by 0.10 gives about 0.864. The analytical derivative is

$$[1 + 1.5/\sqrt{4.25}]/2 = 0.8638.$$

The second line carries 0.1362 of the total angular strength. Halving the centred step should reduce the derivative's truncation error by about four for this smooth curve. In an experiment frequency noise per point is instead amplified as inverse step. Compare several field steps to establish a practical slope uncertainty.

Use an unresolved centroid to detect a calibration problem

At $b = A$ a weak $q = 0$ experiment driven from a selected spin-up preparation records both positive-m final states. Its excitation-weighted centroid after restoring the initial-state energy is $0.78A$. Could the exact two-level mixing alone produce that centroid?

▼ **Work through the solution**

The sum rule requires $b = A$. Direct calculation gives energies approximately $1.2808A$ and $-0.7808A$; weights 0.8638 and 0.1362 give $1.000A$ within rounding. In this model mixing alone cannot explain $0.78A$. A missing weak line, saturation, polarisation contamination, different collected fluorescence efficiencies or a mistaken correction for the initial state can shift a measured centroid. Compare both individual integrated line areas at low intensity and reverse the field with the corresponding spin preparation. These tests use distinct symmetry and strength predictions rather than interpreting one centroid as a complete spectrum.

Further reading

- Cambridge · magnetic fields, Zeeman structures and optical transitions
- Course investigation · angular momentum and atoms
- Fließbach · Näherungsmethoden

Quantum states, dynamics and measurement · 7. Two electrons, exchange splitting and the cost of localisation

Build a two-electron state that respects exchange, derive its charge and spin correlations, and then solve a two-site model exactly. The two examples reveal when a triplet is favoured and when an apparently weak hopping lowers the singlet instead.

Construct states before introducing a force

A two-electron wavefunction has two spatial arguments and two spin labels. The labels 1 and 2 identify argument slots; they do not attach persistent tags to individual identical electrons. Exchanging all position and spin variables must multiply an electronic state by minus one. This is the fermionic exchange postulate. It constrains which vectors describe electrons even when their Hamiltonian contains no spin-dependent term. A measurement with no access to the ordering of slots is represented by an observable symmetric under the same exchange.

Let $a(\mathbf{r}), b(\mathbf{r})$ be normalised orthogonal spatial orbitals. The symmetric and antisymmetric spatial combinations are $\Phi_\pm = [a(1)b(2) \pm b(1)a(2)]/\sqrt{2}$. Combine Φ_+ with the spin singlet $(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)/\sqrt{2}$, and Φ_- with any of the three symmetric triplet spin states. For equal positions the antisymmetric spatial combination vanishes exactly. Its joint probability has an exchange hole at coincidence even for noninteracting same-spin fermions. When $a = b$, the antisymmetric combination vanishes identically and only the symmetric spatial state with opposite spins in a singlet is available. Occupation of the same spatial mode by two electrons is thus compatible with the exclusion principle if the full spin orbitals differ.

For two distinct orthogonal modes, both Φ_+ and Φ_- have the same one-particle number density $|a(\mathbf{r})|^2 + |b(\mathbf{r})|^2$, normalised to two electrons. To see this, integrate $|\Phi_\pm(\mathbf{r}_1, \mathbf{r}_2)|^2$ over $\mathbf{r}_2$ and multiply by two. Cross terms contain $\langle a|b\rangle = 0$. Coincidence and correlation measurements can therefore distinguish preparations that share their entire one-particle density. The spin-antisymmetric singlet can also give correlated opposite spin outcomes along any common measurement axis, while both individual spins are unpolarised. Position and spin comparisons probe different reduced descriptions of the same total state.

Separate the direct and exchange integrals

For an independent-orbital estimate take $H = h(1) + h(2) + v(1,2)$ with $ha = \epsilon_a a,\ hb = \epsilon_b b$, and a symmetric spin-independent repulsion. Expand the expectation of v in $\Phi_\pm$. Diagonal terms give the same direct integral and off-diagonal terms are conjugates:

$$J_{ab} = \iint |a(1)|^2 v(1,2) |b(2)|^2 \, d^3r_1 d^3r_2,$$

$$K_{ab} = \iint a^*(1) b(1) v(1,2) b^*(2) a(2) \, d^3r_1 d^3r_2.$$

For Coulomb repulsion K_{ab} is nonnegative: it is the positive Coulomb quadratic form of the transition density a^*b, seen for instance from its Fourier representation weighted by $4\pi/k^2$. Orthonormality only makes the spatial integral of the transition density vanish, not its weighted self-interaction. The expectation energies are $\epsilon_a + \epsilon_b + J_{ab} \pm K_{ab}$. In this fixed pair of spatial orbitals the triplet is lower by $2K_{ab}$. If the orbitals have

disjoint support the exchange integral vanishes, although the direct repulsion need not. A splitting caused by spatial symmetry does not require an extra microscopic spin–spin repulsive force.

Set $J_{ab} = 0.80$ eV and $K_{ab} = 0.070$ eV in an illustrative excited configuration. Its Coulomb shifts are 0.87 eV and 0.73 eV. A weighted average over all four spin states is

$$J - K/2 = 0.765\,\mathrm{eV}.$$

Optical line strengths generally do not populate those four spin states equally. The energy difference is a statement about this restricted configuration; another spatial occupation can place a singlet below the entire triplet manifold. Predicting a ground-state spin therefore requires admitting the competing occupations and the one-body and correlation energies they carry.

Allow hopping and double occupation

Now take two orthonormal localised orbitals L, R of equal one-body energy. A fermion hops between them at amplitude $-t$, and placing opposite spins on the same site costs $U > 0$. Define $H = -t\sum_\sigma (c^\dagger_{L\sigma}c_{R\sigma} + c^\dagger_{R\sigma}c_{L\sigma}) + U\sum_i n_{i\uparrow}n_{i\downarrow}$. This is a specified two-site Hubbard model, not a complete microscopic Coulomb Hamiltonian. Its direct intersite exchange integral is set to zero here, so its spin ordering need not follow the fixed-orbital calculation above. Position, spin and field units enter through the experimentally or computationally determined t, U.

In the two-electron sector a state with two up spins must occupy one orbital on each site. Hopping would create two up spins in one orbital and therefore annihilates this state. Rotational symmetry implies the same energy zero for all three triplets. Define the singly occupied singlet $|S\rangle$ and symmetric doubly occupied state $|D_+\rangle = (|\uparrow\downarrow, 0\rangle + |0, \uparrow\downarrow\rangle)/\sqrt{2}$. A consistent fermionic ordering gives an off-diagonal $-2t$ between them; each direction and spin contributes and the normalisations leave the factor two. Changing the phase of D_+ reverses that matrix element but leaves the eigenvalues unchanged. The orthogonal antisymmetric doublon remains at energy U. The relevant singlet block is

$$H_S = \begin{pmatrix} 0 & -2t \\ -2t & U \end{pmatrix}, \qquad E_\pm = \frac{U \pm \sqrt{U^2 + 16t^2}}{2}.$$

The ground energy $E_- < 0$ for $t \neq 0$. At $U \gg |t|$ expand the square root to obtain

$$E_- \simeq -4t^2/U + 16t^4/U^3 + \cdots.$$

The singlet–triplet splitting $J_{\mathrm{eff}} = E_T - E_- \simeq 4t^2/U$ can be represented in the singly occupied subspace by $J_{\mathrm{eff}}(\mathbf{S}_L \cdot \mathbf{S}_R/\hbar^2 - 1/4)$. The singlet can lower its energy by virtual double occupation while spin symmetry blocks the triplets. The inferred spin coupling scales as the square of hopping and inversely with the energy cost. This mechanism is commonly called antiferromagnetic superexchange.

Determine a correlation without reading an eigenvector

The fraction of probability with double occupation in the ground state is $\langle \partial H/\partial U \rangle$. Differentiating its energy gives

$$D = \tfrac{1}{2}[1 - U/\sqrt{U^2 + 16t^2}].$$

For positive large U, $D \simeq 4t^2/U^2$. Comparing the derivative with the explicit eigenvector provides a useful independent test of signs and normalisations. From $Es = -2td$ one obtains

$$d/s = -E/(2t).$$

Normalising gives $D = E^2/(4t^2 + E^2)$, identical to the derivative expression for either lower-branch parameter value.

For $t = 10$ meV and $U = 100$ meV, $\sqrt{U^2 + 16t^2} = 107.70$ meV and the exact splitting is 3.85 meV. The leading 4.00 meV expression is about 3.9% high. Double occupation is only 3.58%, yet forbidding it at the outset would predict exactly zero splitting. A virtually occupied state can matter appreciably in a low-energy Hamiltonian even when a projective charge measurement rarely finds it occupied. At $U = 0$ the ground energy is $-2|t|$, corresponding to two opposite spins in the bonding one-electron orbital; the double-occupation fraction is one half. The large-U expression diverges as $U \to 0$ because its perturbative premise disappears. The exact two-state matrix remains finite.

Derive an effective screening charge for a helium-like ground state

A two-electron atom provides a continuous-space counterpart to the two-site model. In atomic units set $\hbar = m_e = e^2/(4\pi\epsilon_0) = 1$ and take an infinitely heavy nucleus of charge Z. The nonrelativistic Hamiltonian is

$$H = \sum_{i=1}^{2}[-\nabla_i^2/2 - Z/r_i] + 1/r_{12}.$$

The two one-electron ground orbitals may be identical in space because opposite spins can form an antisymmetric singlet. Use a trial product of two normalised hydrogenic functions $u_\zeta(\mathbf{r}) = (\zeta^3/\pi)^{1/2}e^{-\zeta r}$ with a variable positive exponent ζ. This single parameter lets the spatial cloud expand in response to mutual repulsion. Choosing the exponent equal to Z without testing it would impose the noninteracting size on an interacting problem.

A single orbital gives $\langle -\nabla^2/2\rangle = \zeta^2/2$ and $\langle 1/r\rangle = \zeta$, either by direct integration or by hydrogenic scaling. Two electrons thus contribute $\zeta^2 - 2Z\zeta$ to the expectation. To derive the repulsion term, average $1/|\mathbf{r}_1 - \mathbf{r}_2|$ over two independent uniform angles at fixed radii. Gauss's shell result, or a direct change of variable in the angular integral, gives $1/\max(r_1, r_2)$. The radial probability is

$$P(r) = 4\zeta^3 r^2 e^{-2\zeta r}.$$

Symmetry between the two electron labels permits restricting one integral to $r_2 < r_1$, then doubling:

$$J = 2\int_0^\infty \frac{P(r)}{r}\,dr \int_0^r P(s)\,ds,$$

$$\int_0^r P(s)ds = 1 - e^{-2\zeta r}(1 + 2\zeta r + 2\zeta^2 r^2).$$

The first part of the resulting r integral gives 2ζ. The three exponential subtraction terms sum to $11\zeta/8$, leaving $J = 5\zeta/8$. Their evaluation uses $\int_0^\infty r^n e^{-ar}dr = n!/a^{n+1}$. Both the dependence linear in ζ and the positive sign follow from scaling an inverse-distance repulsion; only the numerical coefficient needed the angular and radial integration. Replacing the expectation of $1/r_{12}$ by $1/\langle r_{12}\rangle$ would generally give a different coefficient because inverse distance is nonlinear.

Consequently $E(\zeta) = \zeta^2 - 2(Z - 5/16)\zeta$ is an upper bound on the ground energy. It minimises at $\zeta_* = Z - 5/16$, with $E_{\mathrm{var}} = -(Z - 5/16)^2$ hartree for positive ζ_*. For helium $Z = 2$, $\zeta_* = 1.6875$, giving -2.84766 hartree, compared with -2.75 hartree at $\zeta = 2$. The one-electron ion has energy -2 hartree, so this variational value lies below the threshold for an electron plus the ion and demonstrates binding within the model. The trial ionisation energy 0.84766 hartree is a lower bound on the exact ionisation energy when that threshold is known exactly. An upper bound on the neutral energy gives a lower bound on the energy required to remove an electron.

The interpretation of $5/16$ as a screening correction is limited. At short radii each electron feels nearly the full nuclear singularity, while far outside it sees the remaining net charge. One constant exponent cannot reproduce both limits and cannot describe how the two electron positions avoid one another at fixed one-particle radius. The product trial permits nonzero joint weight at electron coincidence. Multiplying by a factor involving r_{12} introduces genuine correlation beyond independent orbitals. Optimising a broader variational family containing this product can only lower or retain the variational energy, provided its integrals and normalisation are computed consistently. Monitoring energy together with the pair-separation distribution tests whether the enlargement captures the physics of repulsion rather than merely adjusting a one-body radius.

Infer a local two-electron boundary condition from the Coulomb singularity

Near coincidence of two unlike-spin electrons, the Coulomb repulsion diverges like $1/r_{12}$. The wavefunction itself can remain finite at coincidence, but its derivative must cancel this singularity in the Schrödinger equation. Consider a fixed smooth centre coordinate and average the wavefunction over the direction of the interelectron separation r, away from simultaneous nuclear coincidence. Assume an expansion $\bar{\Psi}(r) = A + Br + \cdots$. In atomic units the reduced mass is one half, so the relative spherical kinetic operator is $-[d^2/dr^2 + (2/r)d/dr]$ for the locally s-wave part. Its leading singular term is $-2B/r$. Adding $\bar{\Psi}/r$ gives coefficient $A - 2B$ of $1/r$. A finite energy eigenvalue requires this coefficient to vanish:

$$\left.\frac{d\bar{\Psi}}{dr}\right|_0 = \frac{1}{2}\,\bar{\Psi}(0).$$

The positive cusp slope means the spatial amplitude rises as the electrons move apart from contact, consistent with repulsion. The corresponding electron–nucleus cusp has opposite sign and coefficient set by the nuclear charge and reduced mass. In a same-spin state the scalar contact amplitude vanishes by

antisymmetry; this s-wave formula becomes $0 = 0$ and a leading higher-partial-wave analysis is needed. Forgetting either angular averaging or exchange symmetry can turn an instructive local condition into a false statement about every directional derivative of every wavefunction.

A product of independent $1s$ orbitals is smooth in r_{12} at a generic fixed centre and has no required linear cusp in the spherical average. A trial factor $1 + r_{12}/2$ in atomic units can impose the first-order unlike-spin cusp while leaving additional variational freedom to model long-range behaviour. Merely enforcing the cusp does not solve the entire two-electron problem. It constrains a local ratio of derivative to value, while energy depends on a six-dimensional integral. In variational Monte Carlo one can diagnose this distinction through the local energy $H\Psi/\Psi$ near coincidence. A cusp-mismatched trial may have a singular local energy even though the volume element $r^2 dr$ keeps its integrated average finite. Imposing the cusp can reduce variance and improve short-distance pair observables more dramatically than it changes the energy expectation.

Use one-body occupation measurements to detect missing correlations

A single determinant of orthonormal spin orbitals χ_i has a one-particle reduced density matrix $\gamma = \sum_{\text{occupied}} |\chi_i\rangle\langle\chi_i|$, normalised to the electron number N_e. Thus $\gamma^2 = \gamma$ and its eigenvalues are zero or one. This is an algebraic signature of a single pure Slater determinant in the spin-orbital convention. For a correlated pure fermion state, natural occupations can become fractional while remaining between zero and one and summing to N_e. Ensemble mixing can produce fractional occupations too, so knowing whether the entire state is pure matters if one wants to interpret fractional one-particle populations as evidence of correlation.

Consider the normalised two-electron singlet $|\Psi\rangle = \sqrt{0.8}\,|a\uparrow, a\downarrow\rangle + \sqrt{0.2}\,|b\uparrow, b\downarrow\rangle$, with each ket a determinant. A one-body operator changes at most one occupied spin orbital, so its matrix element between these two doubly different determinants vanishes. The four natural spin-orbital occupations are $0.8, 0.8, 0.2, 0.2$. Their sum is two while $\operatorname{tr}\gamma^2 = 1.36$, smaller than the determinant value two. An incoherent mixture of the two determinants at the same weights has the same one-body matrix, although it no longer has their joint coherence. A two-body pair-transfer operator that converts both a electrons into b electrons, plus its adjoint, can reveal a relative phase between these branches. Thus a complete one-body spectrum still need not determine a two-body state.

The symmetric Hubbard dimer ground state has another direct relation between energy derivatives and its one-body coherence. Let $T = \sum_\sigma (c^\dagger_{L\sigma} c_{R\sigma} + c^\dagger_{R\sigma} c_{L\sigma})$. The energy derivative gives

$$\langle T\rangle = -\partial E_-/\partial t = 8t/\sqrt{U^2 + 16t^2}.$$

At $U = 0$ and $t > 0$, $\langle T\rangle = 2$ because both electrons occupy the bonding orbital. At large U/t this coherence approaches zero even though the exact finite-t ground state remains a total spin singlet. In the bonding/antibonding spatial basis, the sum of both populations is two and their difference is $\langle T\rangle$. For the example $U = 100,\ t = 10$ meV this gives spatial occupations roughly 1.371 and 0.629. Dividing by two gives the spin-orbital occupations for each spin in this singlet. Site populations alone are one on each site by symmetry at every U/t; their apparent constancy completely misses this changing form of correlation.

Break site symmetry and measure the onset of charge transfer

Add a static site-energy imbalance with energies $-\Delta/2$ on L and $+\Delta/2$ on R. The singly occupied configuration still has total site energy zero, while double occupations cost $U-\Delta$ and $U+\Delta$. For $|t| \ll U - |\Delta|$ the singly occupied singlet gains the second-order shift $-2t^2/(U-\Delta) - 2t^2/(U+\Delta)$, since its matrix element to either double occupation has magnitude $\sqrt{2}|t|$. The triplets remain blocked. The effective exchange splitting is $4t^2U/(U^2-\Delta^2)$. At zero imbalance it reduces to the earlier $4t^2/U$; small imbalance changes it quadratically, consistent with exchanging left and right sites.

For $U = 100,\ t = 5,\ \Delta = 40$ meV, the estimate gives 1.19 meV compared with 1.00 meV at zero bias. Increasing imbalance can make one virtual double occupation especially inexpensive, changing the gap before the mean site charge substantially relocates. Near $\Delta = U$ the denominator vanishes and the perturbative model predicts a divergence. The true three-state singlet matrix instead has an avoided crossing of singly and doubly occupied configurations. A sweep through that region should retain both double occupations and compute site charge from $-2\,\partial E/\partial\Delta$, with the sign matched to the specified imbalance. Here an apparent runaway effective spin coupling marks the failure of eliminating charge degrees of freedom, rather than an arbitrarily large spin energy in the actual finite Hamiltonian.

Measure singlet–triplet exchange without a spin force

Change the on-site repulsion in a two-electron dimer. Check whether exchange inferred from the spectrum follows the virtual-hopping estimate.

Open this experiment online

Try the calculation

On-site energy U (meV) **6**

Hopping t (meV) **1**



Triplet energy 0. Lower singlet energy -0.60555 meV. Estimated gap 4t²/U = 0.66667 meV.

See exchange interference directly in a two-port experiment

A lossless balanced beam splitter for two orthonormal single-particle modes can be chosen to act on creation operators as $a^\dagger \to (c^\dagger + d^\dagger)/\sqrt{2},\ b^\dagger \to (c^\dagger - d^\dagger)/\sqrt{2}$. Bring two identical particles to the device, one per input, with wave-packet timing and internal states controlled. The input $a^\dagger b^\dagger|0\rangle$ contains no explicit interaction. Nevertheless the coincidence statistics after the device depend on exchange because two paths can lead to the same unordered output configuration. This provides a way to test exchange far from the Coulomb energy calculations of the previous sections.

For identical spinless bosons multiply the two sums and use commuting creation operators. The cross terms cancel, leaving $[(c^\dagger)^2 - (d^\dagger)^2]|0\rangle/2 = (|2,0\rangle - |0,2\rangle)/\sqrt{2}$. Both particles leave the same port and coincidence probability vanishes. The $\sqrt{2}$ in a normalised doubly occupied Fock state is essential. For identical fermions in the same internal state the repeated creation terms vanish and the cross terms add because $d^\dagger c^\dagger = -c^\dagger d^\dagger$. The outgoing state has one particle in each port with unit probability. Neither result requires an attractive or repulsive force at the splitter. Their interpretation turns on whether the alternatives remain indistinguishable and on the symmetry of the total input.

Now give the two particles internal wave packets $|u\rangle, |v\rangle$, which may encode time, frequency or spin, and suppose the detectors report output ports but do not resolve internal labels. One coincidence alternative puts u in c and v in d; the other interchanges them. The interference cross term contains the product $\langle u|v\rangle\langle v|u\rangle = |\langle u|v\rangle|^2$. At a balanced splitter the coincidence probability is $[1 - |\langle u|v\rangle|^2]/2$ for initially independently prepared bosons and $[1 + |\langle u|v\rangle|^2]/2$ for corresponding fermions. At zero overlap both reduce to one half, as expected for two distinguishable independent particles each taking a random output.

Take transform-limited Gaussian wavepacket intensities of common time variance σ_t^2 with delay τ between their centres and no carrier mismatch. Integrating their two normalised amplitude envelopes gives

$$|\langle u|v\rangle|^2 = \exp[-\tau^2/(4\sigma_t^2)].$$

For $\sigma_t = 0.50$ ps and delay 1.00 ps this is

$$e^{-1} = 0.368.$$

The predicted bosonic coincidence fraction is 0.316 and the fermionic value 0.684. A narrow time gate can itself identify the particles differently at different delays and must be included in the counting model. Measuring the single-output rates together with coincidences tests splitter balance and detection bias before converting dip or peak contrast into an overlap estimate.

Infer distinguishability with an unbalanced splitter

A calibrated splitter has intensity transmittance $T = 0.60$, reflectance $R = 0.40$ and negligible loss. Two photons from separate inputs give coincidence probability 0.184 at equal central arrival time. Infer their internal squared overlap. Would equal-time arrival alone have guaranteed perfect bunching?

▼ Derive and check the result

For indistinguishable paths the two coincidence amplitudes have opposite signs with magnitudes T, R. A nonresolving detector gives

$$P_c = T^2 + R^2 - 2TR|\langle u|v\rangle|^2.$$

Substitution yields squared overlap $(0.52 - 0.184)/0.48 = 0.70$. Even perfect internal overlap at this unbalanced splitter would leave $(T - R)^2 = 0.040$ coincidences. Equal arrival time controls only one possible distinguishing label; orthogonal polarisation or spectral mismatch can still reduce overlap. A far-delay reference should give $T^2 + R^2 = 0.52$ for distinguishable, otherwise comparable photons. This independent control helps separate exchange interference from lost counts or a miscalibrated splitter.

Tutorial · discriminate spatial exchange from virtual hopping

Normalise two nonorthogonal orbitals before predicting coincidences

Let a, b each have norm one but overlap $s = 0.30$. Find the normalisations of their symmetric and antisymmetric products. Does the antisymmetric state still vanish at common position?

▼ **Work through the solution**

The inner product $\langle ab|ba\rangle = |s|^2 = 0.09$. The normalisation denominators are

$$\sqrt{2(1 + 0.09)} = 1.476,$$

$$\sqrt{2(1 - 0.09)} = 1.349.$$

At identical positions the antisymmetric numerator cancels pointwise for every s. When the orbitals become identical its norm tends to zero and there is no independent antisymmetric state of that pair. For nonzero overlap the one-body cross terms in the earlier density calculation no longer integrate away, so those equal-density formulas must be rederived. Orthonormalising local orbitals changes the parameters of a corresponding Hubbard model as well.

Infer a hopping amplitude from a measured singlet gap

A calibrated two-site sample has $U = 120$ meV and a singlet–triplet gap 5.0 meV. In the model above find t exactly, predict the double-occupation probability and compare with a large-U inference.

▼ **Work through the solution**

Writing $J = (\sqrt{U^2 + 16t^2} - U)/2$ gives $4t^2 = J(U + J)$, hence $|t| = 12.50$ meV. The large-U estimate $\sqrt{JU}/2 = 12.25$ meV is close but measurably different. The double-occupation fraction is

$$\tfrac{1}{2}[1 - U/(U + 2J)] = J/(U + 2J) = 0.03846.$$

Gap alone cannot distinguish signs of t in a two-site model because rephasing one orbital reverses the sign. Measuring charge fluctuations independently tests whether the observed spin gap can be explained by this same virtual process.

Predict coherent double-occupation oscillations

Start suddenly in the singly occupied singlet of the symmetric two-site model. Find the probability of observing either doubly occupied site after time τ and its maximum at $U = 100,\ t = 10$ meV. Predict the triplet control.

▼ **Work through the solution**

Subtract $U/2$ times identity from the two-dimensional Hamiltonian. The traceless matrix squares to $(U^2 + 16t^2)I/4$. Exponentiating gives

$$P_D(\tau) = 16t^2 \sin^2[\sqrt{U^2 + 16t^2}\,\tau/(2\hbar)]/(U^2 + 16t^2).$$

Its maximum is

$$1600/11600 = 0.138.$$

This is larger than the ground-state admixture because the sudden preparation is a coherent superposition of both energy eigenstates. A triplet evolves with no doubly occupied component in this ideal Hamiltonian. Observing double occupation after triplet preparation would test spin mixing, state-preparation fidelity or additional orbitals.

Further reading

- Fließbach · Mehrteilchensysteme
- Cambridge · electron exchange, helium and determinants
- Course investigation · exchange of identical particles

Quantum states, dynamics and measurement · 8. From a finite pulse to a continuum transition rate

Switch a weak perturbation on for a measured duration. Its Fourier transform determines the amplitude of a discrete transition. Summing over a dense set of final states reveals where a time-independent transition rate comes from and where it fails.

Remove the known free evolution before approximating

Let $H_0|n\rangle = E_n|n\rangle$ and $H(t) = H_0 + V(t)$. Expand $|\psi(t)\rangle = \sum_n c_n(t) e^{-iE_n t/\hbar}|n\rangle$. Substituting into the time-dependent Schrödinger equation cancels the free phase derivatives and leaves an exact equation $i\hbar \dot{c}_f = \sum_n e^{i(E_f - E_n)t/\hbar} V_{fn}(t) c_n(t)$. Suppose the initial state is $|i\rangle$ and the perturbation is appreciable only between 0 and T. Iterating the integral equation once, replacing $c_n(t)$ on the right by δ_{ni}, yields

$$c_f^{(1)}(T) = -\frac{i}{\hbar}\int_0^T e^{i\omega_{fi}t} V_{fi}(t)\, dt,$$

$$\omega_{fi} = (E_f - E_i)/\hbar.$$

For $f \neq i$ the leading transition probability is $|c_f^{(1)}|^2$. It is second order in the perturbation even though the amplitude is first order. Conservation of total probability to this order requires a second-order correction to the initial amplitude. Simply appending all positive final probabilities while keeping the initial probability exactly one would fail that test. The approximation is controlled when total probability removed from the initial state is small over the observation duration and higher-order pathways are negligible.

Let the resonant component of a drive have matrix element $V_{fi}(t) = v e^{-i\omega t}$ for $0 < t < T$. For a real cosine perturbation the two exponentials each carry half the real peak amplitude, so v refers to just the near-resonant part. Define $\Delta = \omega_{fi} - \omega$. Direct integration gives

$$P_f^{(1)} = \frac{|v|^2 T^2}{\hbar^2}\operatorname{sinc}^2(\Delta T/2),$$

$$\operatorname{sinc} x = \sin x / x.$$

At resonance the amplitude grows linearly in time while its phase remains coherent. A detuned contribution winds around in the complex plane and largely cancels. The first zeros occur at ordinary detuning $\pm 1/T$ in hertz. The full width at half height is approximately $0.886/T$ in ordinary frequency since $\sin x/x = 1/\sqrt{2}$ at $x \simeq 1.3916$. Doubling duration quadruples the line-centre probability for one final state but halves the linewidth, a distinction central to interpreting a continuum rate.

Sum over final states before taking a long-time limit

A large box provides a discrete set of final levels with spacing δE; at fixed finite time many can lie inside the pulse's energy bandwidth. Let $\rho(E)$ be the density of the accessible final states per joule, including degeneracy exactly once. Replace the sum by an integral with $E_* = E_i + \hbar\omega$:

$$P_{\text{out}} \simeq \frac{T^2}{\hbar^2}\int dE\, \rho(E)|v(E)|^2 \operatorname{sinc}^2\left[\frac{(E - E_*)T}{2\hbar}\right].$$

If $\rho|v|^2$ changes little across the dominant width and the line is well away from thresholds, take it outside the integral. Substitute $u = (E - E_*)T/(2\hbar)$, and use $\int_{-\infty}^{\infty} (\sin u/u)^2 du = \pi$. The area of the squared Fourier transform grows linearly, giving

$$\Gamma = \frac{P_{\mathrm{out}}}{T} \simeq \frac{2\pi}{\hbar} |v(E_*)|^2 \rho(E_*).$$

The identity for the sinc area also follows by Parseval from the energy in the rectangular time window. This is Fermi's golden rule with ρ per energy. If a supplied density is per unit angular frequency,

$$\rho_\omega = \hbar \rho_E;$$

failing to change that measure introduces a factor of $\hbar$. The units are a strong check: energy squared times states per energy divided by action is inverse seconds.

Several limits have been taken in a definite order. The perturbation lasts long enough to resolve the scale on which $\rho|v|^2$ varies but short enough that many levels still fit in its bandwidth. A window of times must exist between a reservoir correlation time and a recurrence time of order $2\pi\hbar/\delta E$. In addition $\Gamma T \ll 1$ for the first-order depletion argument. Taking $T \to \infty$ at fixed finite box instead resolves individual levels and can produce coherent oscillations or recurrences. Extending the constant rate to a full exponential decay uses a further memoryless approximation and renewed weak transitions of the remaining population. A linear short-time fit by itself does not establish exponential dynamics at every time.

Keep matrix elements and density normalisation consistent

For illustration choose $|v| = 5.00 \times 10^{-6}$ eV and $\rho = 20.0\ \mathrm{eV}^{-1}$ for box-normalised accessible modes dense enough on the pulse scale, giving $\Gamma \simeq 4.77 \times 10^{6}\ \mathrm{s}^{-1}$ as a formal product. But the implied mean level spacing $1/\rho = 0.050$ eV corresponds to a recurrence scale of only 8.27×10^{-14} s. Using a nanosecond pulse with this literal finite density violates the continuum criterion, irrespective of how small ΓT looks. This numerical mismatch is a useful trap: reproducing a golden-rule number to many digits does not ensure the system supports irreversible decay in the chosen time range.

Now enlarge the quantisation volume by 10^6 for continuum-normalised physical coupling. The number of accessible states per energy grows by 10^6, while each squared matrix element shrinks by 10^6 for a localised initial state and individually unit-normalised extended final states. The same product gives the same Γ, but spacing is now 5.0×10^{-8} eV and the recurrence scale 82.7 ns. At $T = 2.0$ ns, $\Gamma T = 0.00954$ and roughly $h/(T\delta E) \simeq 41$ spacings fit up to the first positive zero. Whether density and coupling are locally smooth on that bandwidth remains an additional condition. The scaling removes explicit box dependence from a physical rate while clarifying why direct sums at finite volume can require larger spaces.

Use phase space to predict the shape of a continuum edge

For free relative motion in three dimensions, count single scalar internal channels in a cubic periodic volume $\mathcal{V}$. One state occupies volume $(2\pi)^3/\mathcal{V}$ in wave-vector space. The shell between $k, k + dk$ thus contains $\mathcal{V} 4\pi k^2 dk/(2\pi)^3$. From $E = \hbar^2 k^2/(2m)$ it follows that $\rho(E) = \mathcal{V} m k/(2\pi^2\hbar^2)$, which vanishes like $\sqrt{E}$ at threshold. In one dimension the corresponding free density diverges like $1/\sqrt{E}$; the geometrical shell

measure is different. A threshold suppression observed in one spectrum could come from available states, the coupling matrix element, or both.

If a local transition couples chiefly to a final orbital partial wave l, its regular low-k free radial factor near the source scales as k^l when the potential is sufficiently short ranged and no threshold resonance changes the law. Squaring and multiplying by the three-dimensional density gives

$$\Gamma_l \propto k^{2l+1} \propto E^{l+1/2}.$$

For a p-wave channel doubling the excess energy increases this ideal rate by $2^{3/2} = 2.83$. Long-range Coulomb fields alter the final-state radial behaviour and can invalidate the short-range threshold power. Very close to any edge a finite pulse straddles energies where the density changes rapidly or vanishes. Replace the local golden-rule product by the finite-window integral to model the measured onset.

Compare pulse shapes at equal area and equal energy

The rectangular pulse is useful because its Fourier transform is elementary. Its abrupt start and end also generate sidelobes that can excite nearby unwanted transitions. For a general narrowband drive write its resonant matrix element $v(t)e^{-i\omega t}$. First order gives

$$P_f = |\int v(t)e^{i\Delta t}dt|^2/\hbar^2.$$

At zero detuning only the complex pulse area $\int v(t)dt$ matters. After summing over a flat continuum, Parseval gives $\int P_f(\Delta)d\Delta = (2\pi/\hbar^2)\int |v(t)|^2 dt$, so it is the temporal integral of squared coupling that controls total spectral weight. Two pulses with the same area can have different continuum probabilities; two with the same integrated intensity can have different resonant peak heights.

For a Gaussian envelope $v(t) = v_0 e^{-t^2/(2\tau^2)}$ on the full line, complete the square in its Fourier integral. The amplitude is $v_0\sqrt{2\pi}\tau e^{-\Delta^2\tau^2/2}$, so the transition profile has no sidelobes and has full half-height angular width $2\sqrt{\ln 2}/\tau$. The time intensity is $|v_0|^2 e^{-t^2/\tau^2}$, with temporal full half-height width $2\sqrt{\ln 2}\tau$. Multiplying this by the ordinary-frequency full width gives

$$2\ln 2/\pi = 0.441.$$

Replacing an amplitude width by an intensity width changes that product by $\sqrt{2}$; specifying the plotted quantity is necessary when comparing a spectrum with temporal photodetector data.

Take a Gaussian with $\tau = 0.50$ ms. Its spectral intensity full width is 530 Hz and pulse intensity full temporal width 0.833 ms. An unwanted line 1.0 kHz away receives a relative first-order probability $e^{-(2\pi 1000\tau)^2} \simeq 5.17 \times 10^{-5}$. A nominally same-duration rectangular pulse may place an ideal zero near some precisely chosen detuning, but neighbouring frequencies generally see sidelobes. Which pulse gives smaller unwanted excitation depends on what quantity is fixed, where the extra lines lie and how much uncertainty there is in their positions. A shaped envelope can suppress spectral leakage without improving every possible resolution metric.

A pair of pulses illustrates how coherently adding amplitudes differs from broadening a single envelope. Let identical, short pulses of area A be separated by start-time difference T_s, and let the second have a

controlled drive phase ϕ. At a detuning small compared with each pulse bandwidth, the amplitude is proportional to $A[1 + e^{i(\Delta T_s + \phi)}]$. Its squared magnitude has fringes $2|A|^2[1 + \cos(\Delta T_s + \phi)]$ inside the broader single-pulse line. Their ordinary-frequency spacing $1/T_s$ can be much smaller than the inverse individual pulse duration. Randomising the relative phase removes fringes while keeping the two single-pulse energies unchanged. The observed fine spectral structure is therefore a measurement of coherent time separation as well as energy difference.

Derive depletion and the energy shift from the same kernel

The first-order rate gives initial population loss in a useful time window. To see which extra assumptions allow exponential decay, retain a single initial discrete level coupled to many initially empty final levels with time-independent couplings v_f in a rotating frame. Let $c_i(0) = 1,\ c_f(0) = 0$. In the interaction picture $\dot{c}_f = -(i/\hbar)v_f e^{i\omega_{fi}t}c_i(t)$. Integrate this equation exactly in terms of the unknown c_i and insert the result in its partner equation for $\dot{c}_i$. Assuming final levels do not couple directly to one another gives

$$\dot{c}_i(t) = -\int_0^t K(t - t')c_i(t')\,dt',$$
$$K(\tau) = \hbar^{-2}\sum_f |v_f|^2 e^{-i\omega_{fi}\tau}.$$

No irreversibility has been assumed in this elimination. The previous discrete amplitudes are encoded in a memory kernel and can feed back into the initial level. At very short times $c_i(t) \simeq 1 - K(0)t^2/2$ when the second moment is finite, yielding a quadratic survival loss. At long times a sharply structured or bounded continuum can produce nonexponential decay. A finite number of exactly known final levels evolves unitarily under the full Hamiltonian and can show recurrences; treating the memory as identically zero after any arbitrarily short interval would discard this possibility.

When the spectral coupling varies slowly over the line and extends across a sufficiently broad band, phase cancellation makes $K(\tau)$ decay over a correlation time τ_c. For τ_c short relative to changes in c_i, write $c_i(t - \tau) \approx c_i(t)$ inside its memory integral after the initial transient, and extend its upper limit to infinity. The regulated integral $\int_0^\infty e^{-i\omega\tau - \eta\tau}d\tau = (\eta + i\omega)^{-1}$ tends to $\pi\delta(\omega) - i\,\mathcal{P}(1/\omega)$ as a distribution. Its real part gives

$$\dot{c}_i = -\Gamma c_i/2 + \cdots,$$

with the same golden-rule population rate Γ. The imaginary part gives a coherent energy shift proportional to $\mathcal{P}\sum_f |v_f|^2/(E_i - E_f)$, requiring the physical bandwidth and a consistent bare energy. This shows how an absorptive linewidth and a dispersive shift can emerge from different parts of the same elimination, without claiming a perfectly flat infinite continuum supplies a physically meaningful absolute shift.

For a symmetric model with constant spectral strength inside $-W < \omega_{fi} < W$, the kernel is proportional to $\sin(W\tau)/\tau$. Narrowing the band increases its memory time and can delay the onset of an approximately constant rate. Moving the discrete level close to a band edge also removes the cancellation between positive and negative detunings responsible for a vanishing symmetric shift. An experiment that varies the carrier across such an edge should fit the full finite-window spectrum before inferring a time-independent linewidth or level shift. A numerically cheap sum over known discrete final energies provides a check: double the

spectral cutoff, halve level spacing and track the early quadratic regime separately from any intermediate exponential segment.

Ask whether frequent projective tests inhibit a transition

For an initial pure state with finite energy variance under a time-independent H, expand the survival amplitude $A(t) = \langle i|e^{-iHt/\hbar}|i\rangle$ to second order. The linear term $-it\langle H\rangle/\hbar$ is a phase. Squaring gives

$$P_i(t) = 1 - (\Delta H)^2 t^2/\hbar^2 + O(t^3),$$

with still higher-order control depending on available moments. Define $t_Z = \hbar/\Delta H$. In the model with one initial level coupled only off-diagonally to the final levels, $(\Delta H)^2 = \sum_f |v_f|^2$, agreeing with the memory-kernel expansion. If the needed moments diverge in an idealised continuum, this short-time expansion requires a physical high-frequency regularisation.

Now make ideal instantaneous projective tests of whether the system is still in $|i\rangle$ every $\tau = T/N$, selecting only histories that have passed every test. Conditional on survival, each projection prepares exactly the same initial state. The joint survival probability is $P_i(\tau)^N \simeq [1 - (T/Nt_Z)^2]^N \to 1$ as $N \to \infty$ for fixed total duration T. This is a mathematically precise limiting model of the quantum Zeno effect. It is not obtained by simply looking at a stored record more often after a single final measurement; the successive physical interventions reset the conditional state.

With a resonantly coupled two-state Hamiltonian $H = v\sigma_x$, the survival against a single interval is $\cos^2(v\tau/\hbar)$. Choose $vT/\hbar = \pi/2$; unobserved survival at T is zero. For ten equally spaced perfect tests the joint survival is

$$\cos^{20}(\pi/20) = 0.7805.$$

If the instrument has 2% independent loss per successful test, the recorded joint fraction falls to

$$0.7805(0.98)^{10} = 0.638.$$

A finite-duration detector may also change H or broaden spectral coupling rather than implement an instantaneous projection. The ideal repeated-measurement result gives a comparison baseline, while predicted observed counts must include the physical instrument and conditioning.

Derive the sweep-rate scale for a transition through resonance

For a chirped weak drive let its instantaneous detuning from one discrete transition change linearly, $\Delta(t) = \dot{\Delta}\, t$, crossing zero at the middle of a long pulse. The phase accumulated in the first-order integral is $\int^t \Delta(t')dt' = \dot{\Delta}\, t^2/2$, not $\Delta(t)t$. For approximately constant coupling v around the crossing and sufficiently far endpoints the oscillatory Gaussian integral gives $|c_f^{(1)}|^2 \simeq 2\pi|v|^2/(\hbar^2|\dot{\Delta}|)$. A slower passage accumulates a larger resonant amplitude, but this weak-coupling formula cannot exceed probability one. At large transfer the full two-level avoided-crossing dynamics replaces it.

For $|v|/\hbar = 2\pi \times 200$ Hz and an ordinary-frequency sweep $d(\Delta/2\pi)/dt = 10^8$ Hz/s, the estimate is 0.0158. Changing only the sign of the sweep gives the same first-order probability in this symmetric idealisation, while the phase of the transition amplitude changes. Coherently driving a second passage can convert that phase difference into an observable interference pattern. Missing endpoints, asymmetric pulse shaping and additional nearby lines can spoil the simple symmetry. Comparing both sweep directions and checking that the result scales as inverse speed over a low-transfer range are concrete tests of the stationary-phase approximation instead of assuming every apparent line crossing signifies complete adiabatic transfer.

Separate a transition frequency from coherence loss

Vary detuning and shot-to-shot spread while reading both pulse quadratures. Determine whether a change of population comes from a shifted carrier or a fading phase.

Open this experiment online

Try the calculation

Mean cyclic detuning (Hz) **500**

Shot-to-shot rms spread (Hz) **80**

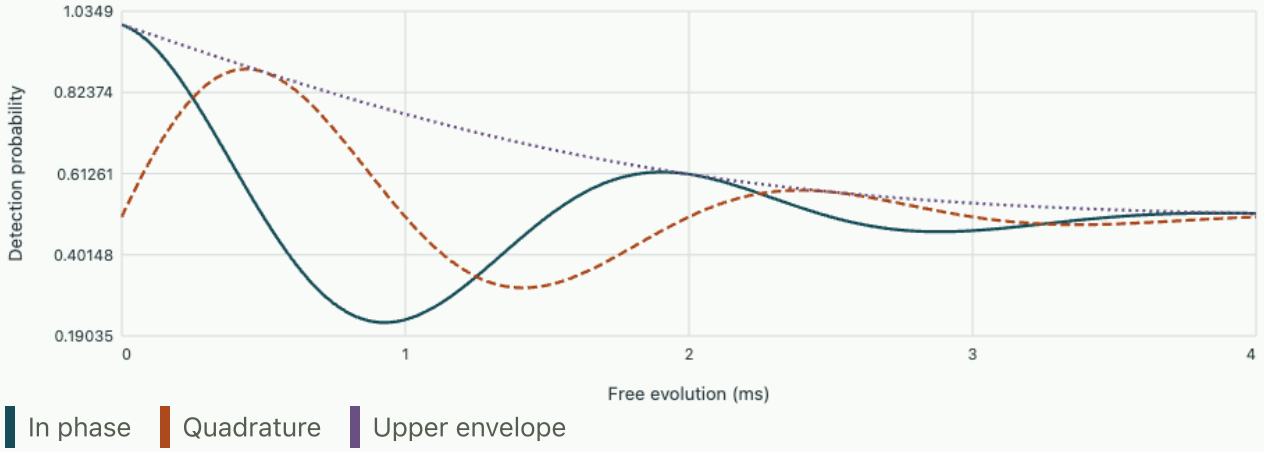

Contrast after 1 ms = 0.53455. The two probabilities at 0.50 ms are 0.5 and 0.87729.

Derive the spectral distinction between a phase-coherent pulse and stationary noise

A fluctuating perturbation is often characterised by a power spectral density rather than a prescribed pulse phase. Suppose $V_{fi}(t) = d_{fi}F(t)$ with F a real classical, zero-mean stationary random process and the coupling weak enough for first-order amplitudes in each realisation. Averaging the modulus squared produces a double time integral involving the autocorrelation $C_F(\tau) = \langle F(t+\tau)F(t)\rangle$:

$$\langle P_f(T)\rangle = \frac{|d_{fi}|^2}{\hbar^2}\int_{-T}^{T}(T-|\tau|)C_F(\tau)e^{i\omega_{fi}\tau}\,d\tau.$$

The factor $T - |\tau|$ counts the length of pairs of times in $[0, T]$ separated by τ. If T is long compared with the correlation time and depletion remains weak, divide by T to find rate $|d_{fi}|^2 S_F(\omega_{fi})/\hbar^2$, where $S_F(\omega) = \int_{-\infty}^{\infty} C_F(\tau) e^{i\omega\tau} d\tau$ is the two-sided angular-frequency spectrum. No sum over a continuum of final quantum states was involved here. A continuum of drive frequencies supplies a different physical route to an approximately linear transition probability of one resolved final state.

For exponentially correlated noise $C_F(\tau) = F_{\mathrm{rms}}^2 e^{-|\tau|/\tau_c}$, integration gives

$$S_F(\omega) = 2F_{\mathrm{rms}}^2 \tau_c / [1 + \omega^2 \tau_c^2].$$

Fix F_{rms} and vary correlation time. At a nonzero fixed transition frequency, extremely rapid fluctuations have little integrated spectral weight there because their fixed variance is spread over a very broad band. Very slow fluctuations concentrate their weight near zero frequency. Differentiating the spectral expression gives a maximum excitation rate at $\tau_c = 1/|\omega_{fi}|$. This nonmonotonic dependence cannot be inferred just by quoting the field's total rms amplitude. Unit conventions are also critical: changing to ordinary-frequency spectral densities and a one-sided instrument display requires the corresponding Jacobian and factor of two.

A real classical stationary process has an even spectrum $S_F(-\omega) = S_F(\omega)$, so this model drives upward and downward transitions equally per available source population when matrix elements are conjugate. A quantum environment at thermal equilibrium can have unequal positive and negative frequency unsymmetrised spectra. Its absorption and emission probabilities encode detailed balance. Thus an observed difference between calibrated upward and downward rates cannot generally be explained by merely drawing samples from a real classical noise spectrum of the same symmetrised power. Distinguishing a driven classical source from an equilibrium quantum bath requires precise statements about which correlations enter the transition probability.

Infer a fluctuation timescale from a rate ratio

Two transitions with equal coupling strengths see the same exponentially correlated classical field at ordinary frequencies f and $2f$. Their measured weak transition rates obey $\Gamma(2f)/\Gamma(f) = 0.40$. Infer the dimensionless correlation frequency and, for $f = 50.0$ kHz, the correlation time. What feature of the acquisition would invalidate this inference?

▼ **Derive and check the result**

Let $x = (2\pi f \tau_c)^2$. The ratio is $(1 + x)/(1 + 4x) = 0.40$, giving $0.60 = 0.60x$ and $x = 1$. The correlation time is

$$1/(2\pi f) = 3.18\,\mu\mathrm{s}.$$

At the lower frequency the noise spectrum lies at half its zero-frequency value; the higher line samples one fifth of that value. If the exposure time is comparable to or shorter than τ_c, the triangular window in the double integral still influences both rates and the infinite-time Lorentzian ratio is inadequate. Likewise unequal matrix elements or a finite observation band must be corrected separately. Measuring a third transition at another frequency supplies an independent test of the exponential-correlation model rather than simply determining its fitted timescale.

Tutorial · choose a pulse duration and check the perturbation

Measure a frequency width without confusing two line shapes

A rectangular pulse lasts 4.00 ms with resonant coupling $|v|/\hbar = 25.0\,\mathrm{s}^{-1}$. Find the line-centre probability, separation of the first zeros and full half-height width in hertz.

▼ **Work through the solution**

At resonance first-order probability is

$$(25 \times 0.004)^2 = 0.0100.$$

The first zeros are 250 Hz on either side, giving a 500 Hz separation. The full half-height width is

$$0.886/0.004 = 221.5\,\mathrm{Hz}.$$

A Lorentzian fit with its width interpreted as a spontaneous lifetime would mistake pulse broadening for decay. Doubling T while keeping coupling fixed is a diagnostic: the first-order peak would rise fourfold and the width halve. At large peak probability use the full two-level dynamics rather than extrapolating the quadratic indefinitely.

Verify the breakdown against exact two-level dynamics

Assume an isolated resonantly driven two-state system with rotating-frame Hamiltonian off-diagonal v and zero detuning. What replaces the first-order expression when $T = \pi\hbar/(2|v|)$? At which small pulse area is the fractional error approximately one percent?

▼ **Work through the solution**

The resonant Hamiltonian squares to $|v|^2 I$, so the exact probability is $\sin^2(|v|T/\hbar)$. At the stated pulse time it is unity while the first-order result is $(\pi/2)^2 = 2.47$, impossible as a probability. For pulse area $a = |v|T/\hbar$,

$$\sin^2 a = a^2 - a^4/3 + \cdots.$$

The relative discrepancy is approximately $a^2/3$, reaching 1% at $a \simeq 0.173$. The pulse can already make a measurable few-percent transition while first-order theory remains reasonably accurate. A dense continuum requires checking the total across all levels as well.

Infer an accessible density from the observed time dependence

At long-enough but weak-depletion times an ensemble loses 2.0% population in 5.0 ns. The squared box matrix element is calibrated as $2.5 \times 10^{-17}\,\mathrm{eV}^2$. Infer ρ, estimate the recurrence scale and test whether this interpretation is internally consistent.

▼ **Work through the solution**

The slope estimate is $4.0 \times 10^{6}\ \mathrm{s}^{-1}$. With $\hbar = 6.5821 \times 10^{-16}\,\mathrm{eV\,s}$ the inferred $\rho = \Gamma\hbar/(2\pi|v|^{2}) = 1.68 \times 10^{7}\,\mathrm{eV}^{-1}$. The corresponding spacing is 5.97×10^{-8} eV and recurrence scale $h/\delta E \simeq 69$ ns. A 5 ns observation fits within that scale, with about 14 spacings to the first zero. Two percent is compatible with weak depletion. Smoothness of the density and a single accessible channel still require independent spectral checks. Repeating the measurement at several durations reveals whether the inferred slope is stable.

Further reading

- Fließbach · Näherungsmethoden
- MIT 8.322 · time-dependent perturbation theory and scattering
- Course investigation · driven transitions and finite pulse spectra

Quantum states, dynamics and measurement · 9. Inferring a short-range interaction from elastic scattering

A beam of neutral particles encounters a weak, localised static potential. Derive differential cross section from current, extract phase shifts by radial matching and compare exact low-energy scattering with the Born approximation and its failure near a bound state.

Define an observable from incident and outgoing current

Consider elastic scattering from a real, short-range, spherically symmetric potential in three dimensions. In a two-body problem use the relative coordinate and reduced mass m; here the target may simply be a prescribed static potential. At a large distance the stationary solution for a plane wave incident along z has the form $\psi \simeq e^{ikz} + f(\theta)e^{ikr}/r$. The amplitude f has dimensions of length in this convention. Incident current density is $\hbar k/m$ for unit plane-wave amplitude. The outgoing spherical component has radial current density $\hbar k|f(\theta)|^2/(mr^2)$ to leading order. Multiplying by detector area $r^2 d\Omega$ and dividing by incident flux gives

$$d\sigma/d\Omega = |f(\theta)|^2.$$

Normalised wave packets and finite detectors supply a physical limiting interpretation for the plane waves.

The incident and scattered components also interfere. Away from a vanishingly narrow forward region the far-field detector resolves the outgoing scattered beam in this derivation. Integrated forward interference accounts for removed incident flux and leads to the optical theorem. Dropping the cross term everywhere and then demanding total current conservation would incorrectly add a stream of scattered particles without depleting the forward beam. For different outgoing and incident wave numbers an additional ratio of group velocities enters the differential flux; the simple squared amplitude here assumes elastic kinematics.

Match a regular radial solution to free waves

Expand a plane wave in Legendre polynomials as $e^{ikz} = \sum_l (2l+1) i^l j_l(kr) P_l(\cos\theta)$. Outside a short-range potential the reduced radial function in each partial wave is a combination of spherical Bessel and Neumann solutions, with large-distance behaviour proportional to $\sin(kr - l\pi/2 + \delta_l)$. Fix the incoming part equal to that in the incident plane wave. Then the ratio by which the outgoing coefficient is changed is $S_l = e^{2i\delta_l}$. For a real one-channel Hamiltonian conservation of each partial wave's radial current gives $|S_l| = 1$, so δ_l is real modulo π. It is the change in radial phase at the same energy, not a separate potential-dependent energy shift of the free particles far from the target.

Separating the excess outgoing part gives

$$f(\theta) = \sum_{l=0}^{\infty} \frac{2l+1}{k} e^{i\delta_l} \sin\delta_l \, P_l(\cos\theta),$$

$$\sigma_{\mathrm{el}} = \frac{4\pi}{k^2} \sum_l (2l+1) \sin^2\delta_l.$$

The second formula follows from orthogonality of Legendre polynomials in the angular integral. At $\theta = 0$, $P_l(1) = 1$ and $\operatorname{Im}(e^{i\delta_l}\sin\delta_l) = \sin^2\delta_l$, so $\sigma_{\mathrm{el}} = 4\pi \operatorname{Im} f(0)/k$ in this purely elastic model. The differential intensity at some other angle can result from cancellations between waves, even though their contributions to the integrated cross section add separately. A narrow angular detector cannot replace full angular integration without a model of the partial waves.

Derive and interpret a scattering length

Take an attractive spherical well $V = -V_0$ for $r < R$, zero outside. At zero incident energy the regular s-wave reduced solution inside is $u = A\sin(\kappa r)$,

$$\kappa = \sqrt{2mV_0}/\hbar.$$

Outside $u = B(r - a)$; extrapolating this line to zero defines the scattering length a. Matching the logarithmic derivative at R gives $\kappa\cot(\kappa R) = 1/(R - a)$, or

$$a = R\left[1 - \frac{\tan u}{u}\right],$$

$$u = \kappa R.$$

A regular zero-energy extrapolation node can lie outside the range of the potential. It is an intercept characterising the exterior solution and need not mark an actual hard surface. For weak attraction $u \ll 1$ one finds $a/R = -u^2/3 - 2u^4/15 + \cdots$, so a is negative. As u crosses $\pi/2$ the first s-wave bound state appears at threshold and a diverges and changes sign. A large positive scattering length can indicate a shallow bound state for a short-range attraction, whereas a hard repulsive sphere also has positive length. Its sign alone cannot reconstruct an arbitrary potential.

At nonzero k, regular interior solution has wave number $q = \sqrt{k^2 + \kappa^2}$. Match $q\cot(qR) = k\cot(kR + \delta_0)$ to compute an exact phase modulo π. For $kR \ll 1$ and away from singular length scales, the effective range expansion $k\cot\delta_0 = -1/a + r_e k^2/2 + \cdots$ describes the first energy dependence. Writing $f_0 = 1/(k\cot\delta_0 - ik)$ retains unitarity automatically. Neglecting the effective range gives

$$\sigma_0 = 4\pi a^2/(1 + k^2a^2).$$

The zero-energy result $4\pi a^2$ is justified only if $k|a| \ll 1$ as well as range effects being small. Taking a to infinity at fixed k gives the finite bound $4\pi/k^2$ for one partial wave. The apparent infinite cross section at the threshold arose from interchanging limits.

Turn a potential into a first-order Fourier amplitude

The outgoing free Green function obeys $(\nabla^2 + k^2)e^{ikr}/r = -4\pi\delta^3(\mathbf{r})$. Rearrange Schrödinger's equation into a free wave plus an integral of this kernel times $V\psi$. At large observation distance $|\mathbf{r} - \mathbf{r}'| \simeq r - \hat{\mathbf{r}}\cdot\mathbf{r}'$. Replacing the wave inside the integral by the incident plane wave gives the first Born amplitude

$$f_B(\mathbf{k}',\mathbf{k}) = -\frac{m}{2\pi\hbar^2}\int e^{-i(\mathbf{k}'-\mathbf{k})\cdot\mathbf{r}'}V(\mathbf{r}')\,d^3r',$$

$$|\mathbf{k}'-\mathbf{k}| = 2k\sin(\theta/2).$$

For the square well the angular integral is $4\pi\sin(Qr)/(Qr)$,

$$Q = |\mathbf{k}'-\mathbf{k}|.$$

Integrating r^2dr gives

$$f_B = (2mV_0/\hbar^2)[\sin(QR) - QR\cos(QR)]/Q^3.$$

At $Q \to 0$ the bracket is $(QR)^3/3 + \cdots$, so $a_B = -f_B(0) = -u^2R/3$. This reproduces the first weak-well term but misses the bound-state pole entirely. For a real central potential $f_B(0)$ is real. Its first-order imaginary part is zero, while the integrated Born cross section is already second order in V. The forward imaginary contribution needed by the optical theorem first appears at that same second order. Applying the theorem as though a first-order amplitude were exact would mix orders inconsistently.

Extract a phase numerically without fitting an arbitrary sinusoid

For a finite-range real potential, one of the most useful laboratory calculations integrates each radial equation from its regular origin to a matching radius r_m outside the interaction. Use the reduced radial functions $\hat{j}_l(z) = zj_l(z)$ and $\hat{n}_l(z) = zn_l(z)$; their large-z limits are $\sin(z - l\pi/2)$ and $-\cos(z - l\pi/2)$. The exterior solution can be written $u = A[\hat{j}_l(kr)\cos\delta_l - \hat{n}_l(kr)\sin\delta_l]$. The interior ODE supplies u_m, u'_m up to an arbitrary common normalisation. Form the matching determinants rather than estimating a phase by reading one oscillation maximum from a plot.

Taking the radial derivative and eliminating A gives

$$\tan\delta_l = \frac{k\hat{j}'_l(kr_m)u_m - \hat{j}_l(kr_m)u'_m}{k\hat{n}'_l(kr_m)u_m - \hat{n}_l(kr_m)u'_m},$$

where primes on the hatted functions differentiate their dimensionless argument. Both numerator and denominator remain finite if u_m happens to lie close to a node; a direct logarithmic derivative u'_m/u_m can instead amplify roundoff severely. Use the two-argument arctangent before reducing the result modulo π. A reported phase can jump by π under a continuously varying energy while $\sin^2\delta_l$ and S_l remain continuous. Unwrap branches by checking neighbouring energies together with the desired physical reference, especially when estimating an energy derivative of the phase.

For $l = 0$, $\hat{j}_0 = \sin z, \hat{n}_0 = -\cos z$, choose a simple controlled test: a hard sphere of radius R imposes $u(R) = 0$. With $u_m = 0$ at $r_m = R$, the ratio reduces to $-\tan kR$, giving $\delta_0 = -kR \pmod{\pi}$. At $kR = 0.40$, the exact contribution is $4\pi R^2[\sin(0.40)/0.40]^2 = 0.9478(4\pi R^2)$, while the zero-energy length estimate gives $4\pi R^2$. The finite-k effective-range expansion of the hard sphere follows from $k\cot\delta_0 = -k\cot kR = -1/R + k^2R/3 + \cdots$, so

$$a = R,\ r_e = 2R/3.$$

Comparing this term with the previous denominator retaining a alone explains why their leading corrections at finite k need not agree.

A general radial solver should separately vary the first finite radius, step size, matching radius and maximum partial wave $l_{\max}$. Near a nonsingular origin the regular solution behaves up to normalisation as r^{l+1}; initialize both value and derivative with this local solution or a more accurate series. For Coulomb or more singular long-range interactions the exterior wavefunctions must be changed rather than matching to free Bessel functions at an arbitrary finite radius. Relative error near a zero of differential intensity can be meaningless as a convergence measure. Report absolute angular error there along with integrated current, compare several scattering directions and retain the exact free result $V = 0$ as a phase-zero test.

Relate a rapidly varying phase to a wave-packet time delay

A narrow incident energy packet accumulates an extra phase from an elastic partial wave relative to free propagation. The S-matrix phase is $2\delta_l(E)$. In a stationary-phase arrival-time calculation the derivative of phase with respect to energy shifts the envelope peak, giving the Wigner time delay $\tau_l = 2\hbar\, d\delta_l/dE$ when the potential varies slowly across the packet bandwidth and the outgoing branch is well resolved. A derivative of the observable scattering intensity is not enough; $\sin^2\delta$ loses the branch and sign information required for a delay. Causality and reference distances must be treated with care, particularly when a finite detector counts a sum of waves or a long-range potential changes even the asymptotic reference.

A narrow isolated resonance can be represented approximately by a slowly varying background plus a phase with $\cot\delta_{\mathrm{res}} = (E_r - E)/(\Gamma/2)$. Choose the continuous branch increasing through $\pi/2$. Then

$$\sin^2\delta_{\mathrm{res}} = (\Gamma/2)^2/[(E-E_r)^2 + (\Gamma/2)^2],$$
$$d\delta_{\mathrm{res}}/dE = (\Gamma/2)/[(E-E_r)^2 + (\Gamma/2)^2].$$

With nearly constant k and zero background the partial cross section has full energy width Γ and its peak reaches the unitarity limit. The associated peak delay in this convention is $4\hbar/\Gamma$, whereas the population lifetime inferred from an isolated pole with amplitude $e^{-\Gamma t/(2\hbar)}$ is $\hbar/\Gamma$. Both involve the same width, but comparing them as if they were identical measurements would introduce a factor of four.

For $\Gamma = 0.10$ meV and $\hbar = 0.6582\ \mathrm{meV\,ps}$, the pole lifetime is 6.58 ps and the peak partial-wave delay is 26.3 ps in this model. A finite incoming energy spread comparable to 0.10 meV samples a changing phase derivative; it can distort the shape of the outgoing packet and invalidate the simple peak-delay interpretation. A background phase also causes interference: $\sin^2(\delta_{\mathrm{bg}} + \delta_{\mathrm{res}})$ may exhibit an asymmetric feature or a dip while the resonant phase still evolves by almost π. Observing the peak position of a single angular intensity curve is therefore a poor standalone estimator of the resonance energy when background amplitudes are appreciable.

Symmetrise the scattering amplitude for identical particles

If two identical particles emerge into opposite directions in their centre-of-mass frame, detecting one at θ and the other at $\pi - \theta$ does not assign unique incoming particle labels. For a definite spatial symmetry the indistinguishable alternatives interfere. For spinless identical bosons the amplitude is $f(\theta) + f(\pi - \theta)$; for a spatially antisymmetric channel it is $f(\theta) - f(\pi - \theta)$. A differential cross section defined as the number of unordered final pairs per incident flux can be integrated over one hemisphere, or over the whole sphere with a

factor $1/2$. If an instrument instead records the count of individual outgoing particles over all angles it uses a different event convention. State the convention before quoting integrated enhancement or suppression factors.

Since $P_l(-\cos\theta) = (-1)^l P_l(\cos\theta)$, the symmetric combination retains only even l and the antisymmetric one only odd l. At very low energy short-range s-wave scattering is typically dominant. For spinless identical bosons with $f \simeq -a$, the unordered-pair total becomes

$$2\pi|2a|^2 = 8\pi a^2;$$

for distinguishable particles it is $4\pi a^2$. Fully spin-polarised identical fermions instead have antisymmetric spatial motion and no s-wave channel. Low-energy collisions can be suppressed without a repulsive force. If internal spin states differ or are mixed, the allowed symmetric and antisymmetric spin fractions must be tracked before assigning a measured cross section.

For two unpolarised spin-$1/2$ identical fermions with equal independent spin populations and a spin-independent potential, the spin singlet fraction is $1/4$ and permits symmetric spatial motion; the triplet fraction $3/4$ requires antisymmetric spatial motion. Under the stated unordered-pair convention the low-energy total from the s-wave singlet is $2\pi a^2$ after spin averaging. Preparing a particular unlike-spin product and tagging the two final spin species changes the counting experiment, and one should compute its spin-resolved amplitudes rather than reusing the unpolarised factor. A detector that can label internal states can distinguish alternatives that a spin-insensitive detector combines.

See which potential features low-energy data can determine

At low momentum a finite set of parameters such as a, r_e summarises scattering over a narrow energy band. This does not usually identify a complete potential $V(r)$. An illustrative degeneracy is already visible within spherical square wells: $\tan u/u$ has poles and multiple branches, so different well strengths at the same range can give the same scattering length while supporting different numbers of bound states. A single very-low-energy total cross section loses even the sign of that length. Including a measured bound-state count, phase-sensitive interference or higher partial waves can remove some of these ambiguities. The intended parameterisation and the measured energy interval define the inverse problem.

To extract a, r_e from phase-sensitive s-wave data, plot $Y = k\cot\delta_0$ against $X = k^2$. In the effective-range window a straight-line fit has intercept $-1/a$ and slope $r_e/2$. Suppose values in nm units are $(X, Y) = (0.010, -0.190), (0.040, -0.160), (0.090, -0.110)$. They lie on $Y = -0.200 + 1.00X$, giving $a = 5.00$ nm and $r_e = 2.00$ nm. Repeat at a fourth moderately larger k to test curvature from higher terms, then repeat the fit omitting that point. A near-zero intercept makes the extracted a highly sensitive to measurement noise even if individual phase points are accurate. Reporting the intercept and its uncertainty directly can be more informative than a huge indistinct scattering length with asymmetric error bars.

Infer low-energy scattering from a model well

Vary the well strength through a near-threshold state. Compare the measured scattering length with the geometrical range of the potential.

Open this experiment online

Try the calculation

Well phase KR **0.8**

Incoming wave number kR **0.1**

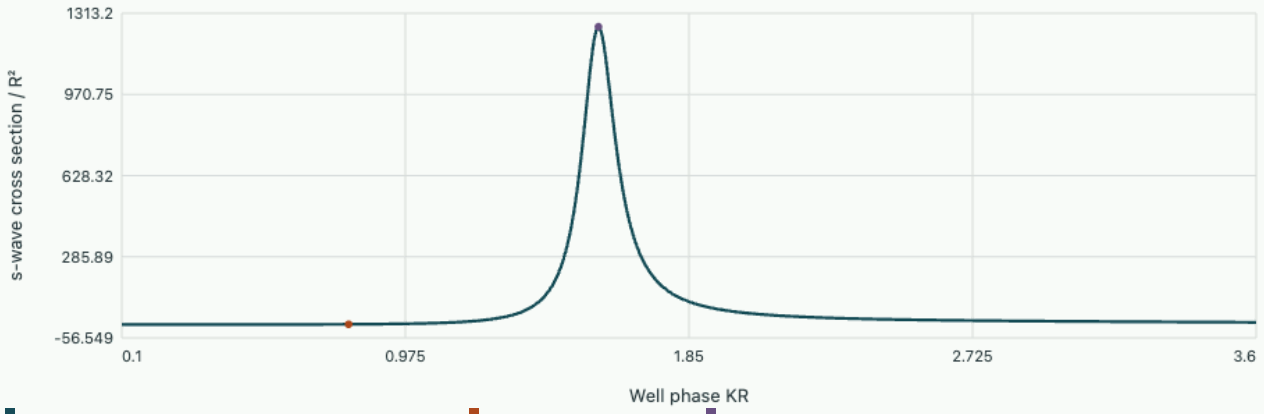

Leading low-energy cross section | Selected depth | First threshold

At KR = 0.8 the cross section is 1.035 R². The finite-k limit at resonance is 1257 R².

Derive an effective range from a zero-energy solution

The scattering length captures the intercept of a zero-energy exterior s-wave. A second coefficient can be related to how the interior wave differs from that extrapolated straight line. For a finite-range real potential choose the dimensionless zero-energy radial solution $u_0(r) \to 1 - r/a$ outside the range. A convenient finite-k exterior normalisation is $u_k = \cos kr + \cot\delta_0 \, \sin kr$; its zero-energy limit gives the same straight line when $k\cot\delta_0 \to -1/a$. Two regular radial equations at energies $0, \hbar^2k^2/(2m)$ give a Wronskian derivative $d(u_0u_k' - u_0'u_k)/dr = -k^2u_0u_k$. Integrate to a radius beyond the potential. The origin contribution vanishes for regular solutions.

Expand the exterior wave using $k\cot\delta_0 = -1/a + r_ek^2/2 + \cdots$. To order k^2 it is $u_k = v + k^2w$, with $v = 1 - r/a$ and $w = -r^2/2 + r^3/(6a) + r_er/2$. In the exterior Wronskian $vw' - v'w$ simplifies to $r_e/2 - \int_0^r (1 - s/a)^2 ds$. Equating to the integrated differential equation gives Bethe's effective range identity

$$r_e = 2\int_0^\infty \left[\left(1 - \frac{r}{a}\right)^2 - u_0(r)^2\right] dr.$$

The integrand vanishes identically outside the true range, so the formal infinite limit is harmless. Both solutions must use the specified equal exterior normalisation; rescaling one independently would leave a divergent difference rather than a material coefficient. The proof also shows why a single scattering length cannot fix an effective range: the former fixes only the exterior line, while r_e measures an integral of the true interior probability relative to its extrapolation.

For a hard sphere of radius $R = a$, set $u_0 = 0$ inside. The integral gives $2\int_0^R (1 - r/R)^2 dr = 2R/3$, agreeing with the direct expansion of $-k\cot kR$. For the attractive square well inside radius R, write $u_0 = (1 - R/a)\sin(\kappa r)/\sin(\kappa R)$. The integral of its square uses

$\int_0^R \sin^2(\kappa r)dr = R/2 - \sin(2\kappa R)/(4\kappa)$. Combining with the scattering-length matching condition reduces the result to $r_e = R - R^3/(3a^2) - 1/(\kappa^2 a)$. This expression should not be evaluated at exactly zero scattering length using naive division; the adopted expansion and exterior normalisation then have to be reconsidered. Close to a diverging length in this finite well one instead finds $r_e \to R$, a natural range scale even while $|a|$ is much larger.

Distinguish two short-range interactions from their energy-dependent phases

A hard sphere and a different finite-range well happen to have equal scattering length $a = 5.0$ nm but effective ranges 3.33 and 5.00 nm. Predict the difference of $k\cot\delta_0$ at $k = 0.10\ \mathrm{nm}^{-1}$. Which measured uncertainty would resolve it at about three standard deviations?

▼ **Derive and check the result**

The shift between the two predictions is

$$\Delta(k\cot\delta) = \tfrac{1}{2}(5.00 - 3.33)(0.10)^2 = 0.00835\,\mathrm{nm}^{-1}.$$

Their values are approximately -0.18335 and $-0.1750\,\mathrm{nm}^{-1}$. If the theoretical curves and common length are accurately specified, a measurement uncertainty below roughly $0.0028\,\mathrm{nm}^{-1}$ would distinguish their difference at three standard deviations. If comparing two independent equally uncertain experimental values, require each to be below $0.00835/(3\sqrt{2}) = 0.00197\,\mathrm{nm}^{-1}$. The expansion error from fourth and higher powers of k must also be smaller than the target difference. Repeating at multiple wave numbers tests that the observed difference scales as k^2 before treating it as evidence for a distinct effective range.

Read a spatial length scale from an angular diffraction envelope

For a weak smooth isotropic potential $V(r) = V_g \exp[-r^2/(2R^2)]$, the first Born amplitude follows from three one-dimensional Gaussian Fourier integrals. In the convention used above,

$$f_B(Q) = -\sqrt{2\pi}\frac{mV_gR^3}{\hbar^2}\exp(-Q^2R^2/2).$$

Its differential intensity relative to the forward value is $\exp(-Q^2R^2)$, with $Q = 2k\sin(\theta/2)$. Forming this ratio cancels the unknown scalar prefactor and fixes the product kR from the angular width. In contrast, a single absolute forward intensity measures $|V_g|^2R^6$ and cannot separate radius and strength. Neither intensity-only measurement fixes the sign of V_g at this order. Thus even in an exactly specified two-parameter family, different angular data have different inverse information.

Suppose $k = 0.40\ \mathrm{nm}^{-1}$ and the intensity at 60° is half that at zero angle. Then $Q = k$ at 60° and $R = \sqrt{\ln 2}/k = 2.08$ nm. At 90°, $Q = \sqrt{2}k$, so the predicted ratio is $1/4$. A third angle tests the Gaussian shape rather than merely refitting its radius. An unmodelled isotropic background would lift the ratio at large angle and bias the inferred radius low if one took logarithms of raw intensity without fitting that background. Use counts divided by incident flux, detector solid angle and efficiency before fitting. Differences of

measurements at several wave numbers can check whether the same spatial R accounts for all angular envelopes.

The integrated cross section can also be evaluated without numerical quadrature. Since $Q^2R^2 = 2(kR)^2(1-\cos\theta)$, integrate over $u = \cos\theta$ to obtain

$$\sigma_B = 2\pi|f_B(0)|^2[1 - e^{-4(kR)^2}]/[2(kR)^2].$$

At $kR \to 0$ the ratio of the exponential difference to $2(kR)^2$ tends to two, yielding the isotropic $4\pi|f_B(0)|^2$. At $kR \gg 1$ the angular envelope narrows and the integrated cross section falls roughly as $1/(kR)^2$ relative to forward intensity. This is a geometrical focusing of the weak-scattering distribution and does not imply that the physical range of the potential changes with incident energy.

Test first Born consistency after fitting a length

For the Gaussian model a fitted low-energy length magnitude is $0.15R$. Infer the dimensionless strength $\eta = m|V_g|R^2/\hbar^2$ predicted by the Born formula. How would a comparison at a second field preparation test whether the amplitude fit is still perturbative?

▼ **Derive and check the result**

The model gives $|a_B|/R = \sqrt{2\pi}\eta$, hence $\eta = 0.0598$. This is small compared with unity, a useful first consistency check for weak low-energy scattering by a smooth single-scale potential. It does not rule out other interactions or state-dependent channels absent from this parametric model. If a controlled preparation doubles the potential amplitude without changing range, first Born predicts four times the differential intensity at every angle and the same normalised angular envelope. Failure of either scaling after calibration can reveal higher-order scattering or a changing spatial form. The test also illustrates why fitting an angular width alone is insufficient to establish that the Born approximation is controlled.

Tutorial · decide which approximation the data support

Compare exact and perturbative lengths at two well strengths

For $R = 2.00$ nm, compute a and a_B for $u = 0.50$ and 1.40. At $k = 0.020\ \mathrm{nm}^{-1}$, determine whether $4\pi a^2$ is a good low-energy cross-section estimate.

▼ **Work through the solution**

At $u = 0.50$, $a/R = 1 - \tan(0.50)/0.50 = -0.0926$, giving -0.185 nm versus -0.167 nm in Born approximation. Here $kR = 0.040$ and $k|a| \simeq 0.0037$. At $u = 1.40$, $\tan 1.40 \simeq 5.798$, giving $a = -6.28$ nm against $a_B = -1.31$ nm. Now $k|a| \simeq 0.126$; the leading unitarity denominator alone reduces $4\pi a^2$ by about 1.6%, before effective-range effects. Weak range resolution and weak scattering are separate criteria. A particle can have wavelength much larger than the target while Born approximation performs poorly.

Reconstruct interference between two partial waves

Only $l = 0, 1$ contribute, with $\delta_0 = \pi/4, \delta_1 = -\pi/6$. Find $k^2 d\sigma/d\Omega$ forward and backward and $k^2\sigma_{\mathrm{el}}$. Verify the optical theorem with the forward value.

▼ **Work through the solution**

In units of $1/k$, the s coefficient is $(1 + i)/2$, and the p coefficient multiplying $\cos\theta$ is $-3\sqrt{3}/4 + 3i/4$. Forward amplitude is roughly $-0.799 + 1.25i$, giving dimensionless intensity 2.20. Backward it is $1.799 - 0.25i$, giving 3.30. The integrated cross section is

$$4\pi[1/2 + 3(1/4)] = 5\pi.$$

The same value follows from the forward imaginary part $4\pi(1.25) = 5\pi$, whereas the forward intensity 2.20 cannot itself be substituted in the optical theorem. Phase-sensitive information survives in angular asymmetry.

Determine what a single low-energy cross section cannot tell you

An elastic isotropic sample gives $\sigma = 400\pi\ \mathrm{nm}^2$ at sufficiently low k. Infer the scattering length magnitude and propose a sign-sensitive check. What changes at $k = 0.1\ \mathrm{nm}^{-1}$?

▼ **Work through the solution**

The zero-energy relation gives $|a| = 10$ nm. The integrated cross section alone admits both signs, and the finite-k expression with zero effective range still depends only on a^2. Interference with a known reference scattering amplitude, a measured phase shift, or independent evidence for a shallow bound state together with range constraints can supply additional information. At the specified finite wave number $k|a| = 1$ the predicted cross section is $200\pi\ \mathrm{nm}^2$ when the effective range remains negligible. Fitting it with the zero-energy relation would return 7.07 nm and obscure the energy dependence.

Further reading

- Fließbach · Dreidimensionale Probleme and scattering
- Cambridge · partial waves, optical theorem and the Born expansion
- MIT 8.322 · scattering lectures 17–22

Quantum states, dynamics and measurement · 10. Distinguishing entanglement from ordinary shared randomness

Two distant analysers sometimes agree even when individual outcomes look random. Derive which correlations shared classical information can produce, compute the quantum prediction and identify what a finite table of coincidence counts can establish.

Build the full probability table for two quantum systems

Each of two separated laboratories receives one two-state system from a joint preparation. For orthonormal local bases the joint basis vectors are $|00\rangle, |01\rangle, |10\rangle, |11\rangle$. A pure state is a complex matrix of coefficients C_{ij} with $\sum_{ij} |C_{ij}|^2 = 1$. If $C_{ij} = a_i b_j$ the state is a product, and all local measurements have independent probabilities. A sum of product terms is always possible but may not factorise into one product. Singular-value decomposition $C = UDV^\dagger$ gives the Schmidt form $\sum_k s_k |u_k\rangle|v_k\rangle$, with orthonormal local bases and $s_k \geq 0,\ \sum s_k^2 = 1$. For two qubits two nonzero Schmidt numbers imply entanglement of the pure state.

Measure only the first system and sum over a complete set of results on the second. Its state is $\rho_A = CC^\dagger$, with eigenvalues s_k^2. For a globally pure state the local purity is therefore

$$\operatorname{tr} \rho_A^2 = \sum s_k^4;$$

it is one only for a product. If the global state is mixed, reduced-state impurity alone does not certify entanglement. A fifty-fifty classical mixture of $|01\rangle$ and $|10\rangle$ already has completely random local measurement results along z and perfect opposition along that axis. A second incompatible measurement arrangement is needed.

Take the singlet $|\Psi^-\rangle = (|01\rangle - |10\rangle)/\sqrt{2}$. The reduced state of either system is $I/2$. Both local means vanish for any analyser direction. Direct calculation gives $\langle \sigma_i \otimes \sigma_j \rangle = -\delta_{ij}$, or equivalently use rotational invariance and the action of each equal Pauli pair on the singlet. Measuring eigenvalues $s, t = \pm 1$ along $\mathbf{a}, \mathbf{b}$ uses product projector $(I + s\mathbf{a}\cdot\boldsymbol{\sigma}) \otimes (I + t\mathbf{b}\cdot\boldsymbol{\sigma})/4$, giving

$$P(s,t \mid \mathbf{a},\mathbf{b}) = \frac{1 - st\,\mathbf{a}\cdot\mathbf{b}}{4},$$

$$E(\mathbf{a},\mathbf{b}) = \sum_{s,t} stP = -\mathbf{a}\cdot\mathbf{b}.$$

Aligned analysers give complete anticorrelation; orthogonal axes give independent equiprobable results; opposite axes give complete agreement. The marginal $\sum_t P(s,t) = 1/2$ is independent of the remote axis. Conditioning on its outcome can sharpen predictions for the other system, but that conditional dataset requires the remote recorded information. For linear photon polarisation analysers the corresponding Bloch angle is twice the polariser angle, a detail that changes physical dial settings even though the probability algebra is the same.

Derive a four-setting inequality from local response functions

Suppose each emitted pair carries information λ drawn from a common distribution independent of the experimenters' analyser choices. A local shared-randomness model supplies outcomes $A, A' \in \{\pm 1\}$ at the

first location and $B, B' \in \{\pm 1\}$ at the second, depending on the local choice and on λ. Local probabilistic response can be treated by including additional random variables in this construction or by bounding their means between minus and plus one. Only one setting is actually selected at each location per run. The model nevertheless assigns responses that can be averaged under the same λ distribution for all four possible pairs.

For each assignment $A(B + B') + A'(B - B') = \pm 2$, since exactly one bracket vanishes and the other is ± 2. Averaging and using the triangle inequality yields the CHSH bound

$$|S| = |E(a, b) + E(a, b') + E(a', b) - E(a', b')| \leq 2.$$

This algebra does not assert that arbitrary measured quantum observables have pre-existing joint outcomes for incompatible settings. It specifies what a particular class of local shared-variable descriptions predicts. It also relies on comparing the same underlying source ensemble in all four settings. Outcome-dependent selection, untracked setting-dependent losses or a changing preparation can invalidate the link between the measured correlations and the inequality. Randomising settings and recording the acquisition and treatment of every event are physical parts of a meaningful experiment.

For the singlet choose $\mathbf{a} = z,\ \mathbf{a}' = x,\ \mathbf{b} = (z + x)/\sqrt{2},\ \mathbf{b}' = (z - x)/\sqrt{2}$. The first three correlations are $-1/\sqrt{2}$, the last $+1/\sqrt{2}$, giving $S = -2\sqrt{2}$. Neither local outcome becomes predictable. What exceeds the bound is a weighted combination of correlations observed in four independently chosen arrangements. A product state will never exceed it, nor will a classical probabilistic mixture of product states, since the correlations of separable preparations admit such local models.

Derive a quantum upper bound and a visibility threshold

Let A, A', B, B' now denote Hermitian operators with squared value identity at their respective locations. Construct $\mathcal{B} = A \otimes (B + B') + A' \otimes (B - B')$. Multiplying out, using the squared identities and retaining cross terms gives

$$\mathcal{B}^2 = 4I - [A, A'] \otimes [B, B'].$$

Each commutator has operator norm at most two, so $\|\mathcal{B}\| \leq 2\sqrt{2}$. The singlet settings above reach this bound. If the two candidate observables commute at either end the cross term vanishes and this construction allows no violation. Thus the upper bound and the need for local incompatibility follow from the same operator product.

Mix the singlet with uniformly random pairs,

$$\rho_v = v|\Psi^-\rangle\langle\Psi^-| + (1 - v)I/4.$$

Every product correlation becomes $-v\,\mathbf{a} \cdot \mathbf{b}$; the tested CHSH settings violate only for $v > 1/\sqrt{2}$. Entanglement itself persists farther into the noise. Transposing just the second subsystem of the singlet projector gives eigenvalues $+1/2, +1/2, +1/2, -1/2$. The smallest partial-transpose eigenvalue of the mixture is therefore $(1 - 3v)/4$. Every separable state's partial transpose is positive, since ordinary transposition preserves positivity on each product factor. Thus $v > 1/3$ certifies entanglement even where this CHSH experiment no longer violates. This argument establishes a sufficient witness from the negative eigenvalue without assuming failure of one witness proves separability.

Report correlation uncertainty using the counted events

For one arrangement write $N_{++}, N_{+-}, N_{-+}, N_{--}$ and N their sum. Estimate $E = (N_{\text{same}} - N_{\text{opposite}})/N$. Each recorded product outcome is ± 1 with variance $1 - E^2$; for independent trials a plug-in standard error is $\sqrt{(1 - \hat{E}^2)/N}$. Separate correctly randomised runs can produce independent estimation errors in the four setting groups. The variance of their signed sum then adds. Marginal balance and full tables offer checks hidden by reporting only one aggregate number. Two tables with the same E can have quite different marginal responses or efficiency problems.

Suppose each setting has 400 accepted pairs and same-result counts are $60, 64, 56, 340$, respectively. The correlations are $-0.70, -0.68, -0.72, +0.70$. Thus $S = -2.80$ and the independent-count plug-in deviation is about 0.071. That comparison is conditional on the sample construction and the independent-trial model. A full experimental Bell test also treats losses, timing, source memory and randomness of settings. A simulation drawing directly from the four Born distributions verifies the arithmetic of the prediction; agreement of its simulated S is not an independent laboratory test of the assumptions.

Use a product-basis measurement to estimate a two-qubit witness

A witness is an observable whose expectation is constrained for every separable preparation but can cross that limit for some entangled states. Consider $W = I/2 - |\Psi^-\rangle\langle\Psi^-|$. To understand the half, maximise squared singlet overlap over a product input $|a\rangle|b\rangle$. Contracting the first spin leaves an unnormalised vector on the second of norm $1/\sqrt{2}$, because the singlet's reduced state is $I/2$. Cauchy–Schwarz then bounds squared product overlap by $1/2$, with equality for appropriately opposite spinors. Any mixture of product states obeys the same bound by linearity. A measured negative $\langle W\rangle$ therefore certifies entanglement without requiring the whole state to be pure or exactly a white-noise mixture.

The singlet projector equals $[I - \sum_i \sigma_i \otimes \sigma_i]/4$, so only three equal-axis correlation settings are needed to measure this witness. Suppose independent experiments yield

$$E_{xx} = -0.55,\ E_{yy} = -0.63,\ E_{zz} = -0.74.$$

The singlet fraction is $F_s = (1 + 1.92)/4 = 0.730$ and $\langle W\rangle = -0.230$. If each correlation has standard uncertainty 0.04, independent random errors give a singlet-fraction uncertainty $\sqrt{3}(0.04)/4 = 0.0173$. Misoriented axes can bias multiple settings together, so independent statistical errors alone are not a calibration error budget. This test does not require verifying an exact isotropic visibility parameter; its bound holds across the entire separable set provided the measurements implement the stated Pauli observables.

A witness is tailored to a target. The maximally entangled $(|00\rangle + |11\rangle)/\sqrt{2}$ has zero singlet overlap, giving positive $\langle W\rangle$ rather than a violation. A known local rotation converts that state to the singlet and would make the same physical analyser scheme sensitive. Without such preparation knowledge, one can measure the full three-by-three correlation tensor and independently choose optimal axes for a separate validation run. Choosing axes to maximise a noisy estimate and then quoting naïve significance on the same data biases the result upward. Calibration, model selection and independent validation serve different roles even in an elementary two-qubit experiment.

Find the best CHSH settings from a measured correlation tensor

For a general two-qubit state define the real 3×3 matrix $T_{ij} = \mathrm{tr}[\rho(\sigma_i \otimes \sigma_j)]$. Its product correlation along unit directions is $\mathbf{a}^T T \mathbf{b}$, irrespective of the local Bloch vectors. The CHSH expression becomes

$$S = \mathbf{a}^T T(\mathbf{b} + \mathbf{b}') + \mathbf{a}'^T T(\mathbf{b} - \mathbf{b}').$$

The sum and difference of two unit vectors are perpendicular. Write them as $2\cos\theta\,\mathbf{u}$ and $2\sin\theta\,\mathbf{v}$ for orthonormal $\mathbf{u}, \mathbf{v}$ and $0 \le \theta \le \pi/2$. For fixed $\mathbf{u}, \mathbf{v}$, choose $\mathbf{a}$ parallel to $T\mathbf{u}$ and $\mathbf{a}'$ parallel to $T\mathbf{v}$ when these are nonzero. The remaining scalar maximisation gives $2\sqrt{\|T\mathbf{u}\|^2 + \|T\mathbf{v}\|^2}$.

The maximum of this sum of quadratic forms over orthonormal $\mathbf{u}, \mathbf{v}$ is the sum of the two largest eigenvalues of $T^T T$. One can see it by expanding both unit vectors in the eigenbasis: each squared projection weight summed over $\mathbf{u}, \mathbf{v}$ lies in $[0, 1]$, their sum is two, and the largest possible weighted sum fills the two largest eigenvalues. Thus $S_{\max} = 2\sqrt{s_1^2 + s_2^2}$ with s_i the singular values of T. A specific state permits a violation by projective spin measurements exactly when $s_1^2 + s_2^2 > 1$. If the singular values are near-degenerate, optimal axes may vary markedly under noise even though the best attainable S changes little.

For a concrete physical Bell-diagonal mixture take $\rho = 0.75|\Psi^-\rangle\langle\Psi^-| + 0.15|\Psi^+\rangle\langle\Psi^+| + 0.10|\Phi^+\rangle\langle\Phi^+|$, where $\Psi^\pm = (01 \pm 10)/\sqrt{2}$, $\Phi^+ = (00 + 11)/\sqrt{2}$ in ket shorthand. The respective correlation triples are $(-1,-1,-1), (1,1,-1), (1,-1,1)$. Their weighted sum is $(-0.50, -0.70, -0.80)$, giving best CHSH magnitude $2\sqrt{0.8^2 + 0.7^2} = 2.126$. If a particular apparatus sampled only the x, y plane its best value would be $2\sqrt{0.5^2 + 0.7^2} = 1.720$. The state and source are identical in both calculations; changing which axes are accessible changes what the experiment can demonstrate. A fixed set of settings optimised for an isotropic singlet might similarly underestimate an anisotropic source.

Derive the noise threshold for separability in a symmetric family

The negative partial-transpose eigenvalue proved entanglement for $v > 1/3$ in the white-noise singlet family. For this special family one can also construct a separable preparation exactly at that boundary. Average product states whose Bloch vectors point in opposite directions:

$$\rho_{\mathrm{avg}} = \int (I + \mathbf{n}\cdot\boldsymbol{\sigma}) \otimes (I - \mathbf{n}\cdot\boldsymbol{\sigma})\, d\Omega/(16\pi).$$

Symmetry gives $\langle n_i \rangle = 0$ and $\langle n_i n_j \rangle = \delta_{ij}/3$. Expanding yields $\rho_{\mathrm{avg}} = [I \otimes I - \sum_i \sigma_i \otimes \sigma_i/3]/4$, precisely the Werner singlet mixture with $v = 1/3$. Even a discrete uniform selection of the six signed Cartesian axes has the same first and second moments here. Mixture with $I/4$, itself separable, then constructs every $0 \le v \le 1/3$.

The full picture of this specific preparation family is therefore quantitative. Below and including one third it has a separable construction. Above one third it is entangled. Up to the CHSH threshold $1/\sqrt{2}$, its correlations still fail to exceed 2 in the test optimised over two spin directions per side. The gain from demanding a stronger operational certificate is a stricter threshold and potentially more demanding statistics. Failing a CHSH test in that intermediate interval does not erase the entanglement witness already derived. Conversely witnessing entanglement with trusted calibrated Pauli devices is a different experimental claim from testing a class of local causal models with explicitly addressed device and sampling assumptions.

For $v = 0.50$, the singlet fraction is $(1 + 3v)/4 = 0.625$ and the smallest partial-transpose eigenvalue is -0.125. The maximal CHSH value is $2\sqrt{2}v = 1.414$. No improvement in count precision can move these ideal mean correlations across 2 without changing the state or measurement scenario. At $v = 0.72$ the maximal value is about 2.036, technically above the threshold but small enough that repeat count and systematic uncertainties matter substantially. This comparison distinguishes identifiability from statistical power: one experiment cannot average its way through an exact theoretical bound, while a narrow permitted violation can become detectable with better data and calibration.

Compute remote conditional states while preserving no signalling

After a local measurement with operators M_a on system A, the unnormalised conditional state of B is $\tilde{\rho}_{B|a} = \mathrm{tr}_A[(M_a \otimes I)\rho_{AB}(M_a^\dagger \otimes I)]$. Its trace is the probability of the local outcome; dividing by that trace gives the conditional state. Summing over all outcomes and using $\sum_a M_a^\dagger M_a = I$ under the partial trace returns ρ_B. The result holds for any complete local instrument, including a nonprojective measurement with forgotten outcomes. A remote observer lacking the classical label sees exactly the same state as before. Conditioning becomes operational only when the two records can be compared.

For the singlet, if A measures direction $\mathbf{a}$ sharply and obtains $+1$, B's conditional Bloch vector is $-\mathbf{a}$; outcome -1 gives $+\mathbf{a}$. In the white-noise mixture those vectors shrink to $\mp v\mathbf{a}$, while each outcome remains equally probable. If the remote partner groups trials by A's announced result and measures along a fixed $\mathbf{b}$, it recovers the same conditional probabilities as the original correlation formula. Without that grouping the two opposite vectors average to zero for every $\mathbf{a}$. This calculation distinguishes a physical method for remotely preparing an ensemble conditional on communication from a method for sending a chosen bit by merely varying the local measurement angle.

A delayed readout changes none of these predicted coincidences because local operators $M_a \otimes I$ and $I \otimes M_b$ commute at the level of this separated systems model. The order in which a data analyst applies conditional probability does change which denominator appears in a finite table. In a test with imperfect detectors, report whether a conditional rate divides by all local triggers, all heralds or only coincidences. Setting-dependent missing events can make a conditional table look no-signalling or strongly correlated for purely sampling reasons. A transparent acquisition protocol preserves the correspondence between the Born joint probabilities being predicted and the ensemble included in their statistical test.

Test which preparations and settings exceed a CHSH bound

Vary analyser angle and visibility independently. Compare the expected joint statistic with the fluctuation scale of a finite counting experiment.

Open this experiment online

Try the calculation

Preparation visibility **0.8**

Symmetric analyser angle (degrees) **45**

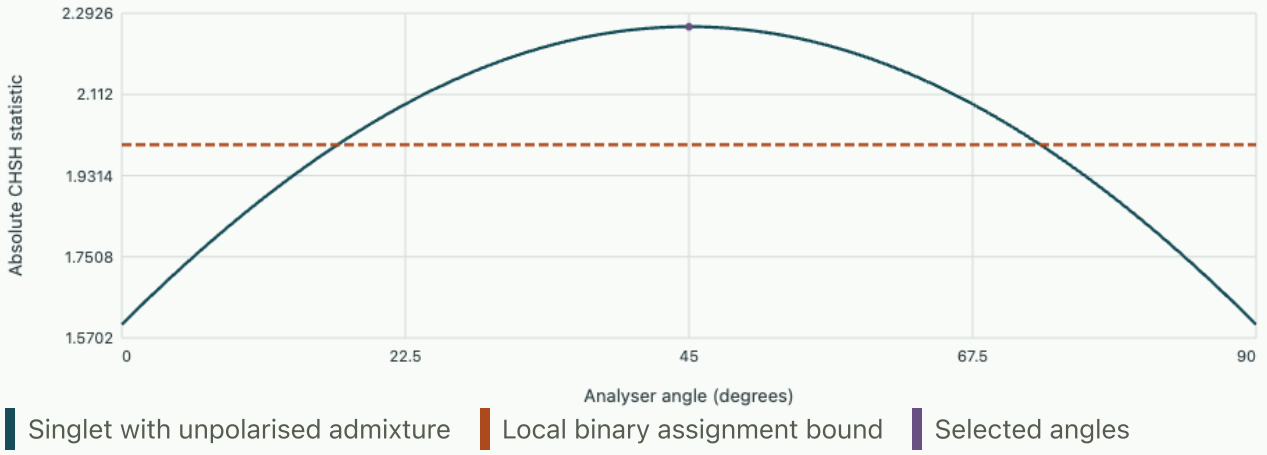

Predicted |S| = 2.2627. With these angles the visibility required to reach two is 0.70711.

Build an entangled pair by conditioning two independent sources

Entanglement is a property of a joint preparation, not a permanent label attached to an individual particle. A four-qubit example illustrates how changing which outcomes are retained can create new conditional pair correlations. Prepare A, B in $|\Phi^+\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$ and independently prepare C, D in the same Bell state. Before any further measurement the reduced preparation of A, D is $I/4$ because the two sources are independent and each outer qubit is individually maximally mixed. A separate central apparatus receives B, C and measures in the Bell basis. The correct prediction after this operation depends on whether its recorded result is available.

Expand the initial state in the order $A, D; B, C$, then use the identity expressing products $|00\rangle, |11\rangle$ as sums and differences of $\Phi^\pm$ and $|01\rangle, |10\rangle$ as sums and differences of $\Psi^\pm$. Regrouping gives

$$|\Phi^+\rangle_{AB}|\Phi^+\rangle_{CD} = \tfrac{1}{2} \sum_{Q\in\{\Phi^+,\Phi^-,\Psi^+,\Psi^-\}} |Q\rangle_{AD}|Q\rangle_{BC}$$

for the consistent real Bell basis used here. The norm stays one because the four central labels are orthogonal. Each central result occurs with probability one quarter. Given a particular result Q, the two outer systems are jointly in Q, even though they came from independent sources and did not directly interact. Ignoring the central result leaves an equally weighted Bell mixture $I/4$, the same outer reduced state as before. Conditioning, or applying a known local Pauli correction using the central label, is necessary to expose a consistent set of Bell correlations in the outer records.

For example, central outcomes Φ^+ and Φ^- both predict equal outer z outcomes, whereas Ψ^+ and Ψ^- predict opposite. In the x basis their signs follow a different pattern. A central device reporting only the parity of its z outcomes distinguishes the Φ from the Ψ subspace but cannot distinguish their relative phases. Conditioned on even parity alone the outer state is $(|\Phi^+\rangle\langle\Phi^+| + |\Phi^-\rangle\langle\Phi^-|)/2$, identical to the classical

mixture of $|00\rangle$ and $|11\rangle$. The very same perfect z correlation therefore has different implications for a resolved Bell label and a mere parity label. What the central device physically measures is essential background for interpreting the outer data.

Quantify how central-label mistakes change outer Bell contrast

Suppose the central Bell outcome is reported correctly with probability q, and otherwise is replaced by each of the three wrong labels with equal probability. The outer parties apply Pauli corrections based on the reported label to target Φ^+. Derive the effective visibility and threshold for a two-setting Bell violation.

▼ **Derive and check the result**

Because every correct label yields the target and wrong Pauli corrections map it to the other three orthogonal Bell states, the corrected outer ensemble is Bell diagonal with target weight q and equal wrong weights $(1-q)/3$. Comparing with $v|\Phi^+\rangle\langle\Phi^+| + (1-v)I/4$ gives $q = (1+3v)/4$, so $v = (4q-1)/3$. In the useful regime $q \geq 1/4$ the maximal CHSH value is $2\sqrt{2}(4q-1)/3$, above two when $q > (1+3/\sqrt{2})/4 = 0.7803$. At $q = 0.85$ the visibility is 0.80 and the best value is about 2.263. Central efficiency and central label fidelity affect different things: if losses are independent, lower efficiency reduces the number of heralded events, while wrong labels reduce conditional contrast. In a realistic inference, setting dependence of losses or an imperfect outer measurement must also be checked before assigning all observed contrast loss to central reporting errors.

Compare local entropy, global entropy and shared classical information

For a density matrix ρ with eigenvalues p_i, define von Neumann entropy in bits by $S(\rho) = -\sum_i p_i \log_2 p_i$, interpreting the zero-eigenvalue term by continuity. For a pure bipartite state the two reduced spectra agree by Schmidt decomposition, so both reduced entropies equal the binary or multi-outcome entropy of the squared Schmidt coefficients. This is an operationally meaningful measure of entanglement for a pure bipartite preparation. For a mixed state, local entropy also contains ordinary uncertainty due to forgotten preparation records and cannot by itself quantify entanglement.

A singlet has $S(AB) = 0,\ S(A) = S(B) = 1$. The equally weighted unentangled anticorrelation mixture $[|01\rangle\langle 01| + |10\rangle\langle 10|]/2$ has $S(AB) = 1$ and the same local entropies. Completely uncorrelated white noise $I/4$ gives $S(AB) = 2$ and still the same local values. Their quantum mutual informations $S(A) + S(B) - S(AB)$ are respectively 2, 1 and 0 bits. The two bits for a singlet do not mean each party receives two freely readable classical bits in a single spin measurement. With aligned local measurements the observed binary records share one bit, while the full quantum mutual information is a property of the joint density operator across all possible probes.

Infer entanglement entropy and contrast for a pure pair

A calibrated pure source has squared Schmidt coefficients $0.90, 0.10$. Find local entropy and compare the strongest CHSH value in its Schmidt basis with that of an incoherent mixture of its two Schmidt product states. Which part of the global description changes after completely dephasing one subsystem in the Schmidt basis?

▼ **Derive and check the result**

The pure-pair entanglement entropy is the binary value $-0.9\log_2(0.9) - 0.1\log_2(0.1) = 0.469$ bits. Write $\cos^2\theta = 0.9$, giving $\sin 2\theta = 2\sqrt{0.09} = 0.60$, and the best CHSH value $2\sqrt{1+0.60^2} = 2.332$. Complete dephasing removes off-diagonal coherence between the two global Schmidt products and leaves their same classical weights. The marginals and local entropies stay exactly the same. The global entropy rises from zero to 0.469 bits; the correlation tensor has only its longitudinal component of unit magnitude and hence best CHSH value two. Comparing these examples shows precisely which physical information is lost when a preparation record or coherent relative phase becomes inaccessible, and why an unchanged local entropy can coexist with a changed experimentally accessible joint correlation.

Tutorial · determine exactly what correlations identify

Distinguish a singlet from classical anticorrelation

Compare the singlet with an equally weighted mixture of $|01\rangle, |10\rangle$. Compute $E(z,z), E(x,x)$ and $E(z,x)$ for both. Which measurement distinguishes them?

▼ **Work through the solution**

The singlet gives $-1, -1, 0$ respectively. In the mixture each prepared product has perfect z opposition, but individual states measured along x have zero mean. Product outcomes factor on each component, so $E(x,x) = 0$ and $E(z,x) = 0$. Measuring just $E(z,z)$ cannot distinguish the two; $E(x,x)$ does. Distinguishing two named hypotheses with one extra setting does not establish a CHSH violation against all local shared-randomness models. The hypotheses and tested claim determine the required data.

Determine a partially entangled state's best planar settings

Let $|\psi\rangle = \cos\theta|00\rangle + \sin\theta|11\rangle$. Use first-side settings z, x and second-side $\mathbf{b}, \mathbf{b}' = (\pm\sin\alpha, 0, \cos\alpha)$, with the plus sign assigned to $\mathbf{b}$. Optimise S.

▼ **Work through the solution**

The correlation tensor in the x, z plane has entries $T_{zz} = 1,\ T_{xx} = \sin 2\theta$ and vanishing cross terms. Therefore $S = 2[\cos\alpha + \sin 2\theta\ \sin\alpha]$. Maximisation gives $\tan\alpha = \sin 2\theta$ and $S_{\max} = 2\sqrt{1+\sin^2 2\theta}$ for these directions. At $\theta = \pi/12$ it is $\sqrt{5} = 2.236$, at a product endpoint it is 2, and at the maximally entangled midpoint $2\sqrt{2}$. The optimising dial settings depend on the state; a poorly chosen apparatus can miss an available violation.

Estimate the needed pair count for a visibility-limited test

Take $v = 0.80$ at the ideal singlet CHSH settings and assume equal independent counts per setting. Estimate the number per setting for a nominal five-standard-error separation from 2. What limits this calculation?

▼ **Work through the solution**

The predicted excess is $2\sqrt{2}v - 2 = 0.2627$. Each product mean has magnitude $v/\sqrt{2}$, so

$$\operatorname{sd} S = 2\sqrt{(1 - v^2/2)/N} = 1.649/\sqrt{N}.$$

Requiring the excess to exceed five such deviations gives $N \gtrsim 986$ pairs per setting. This is a planning estimate using Gaussian independent-sample errors and known visibility. It is not a complete finite-sample hypothesis test and gives no protection against systematic errors or outcome-dependent postselection. Report the four realised counts along with acquisition, losses and uncertainty rather than presenting the planned count as guaranteed evidence.

Further reading

- Cambridge · quantum foundations, Bell inequalities and density matrices
- Course investigation · spin correlations in composite states
- Fließbach · Mehrteilchensysteme

Quantum states, dynamics and measurement · 11. How an unobserved environment changes a quantum state

Model a qubit and a two-state environment by a joint unitary. Derive channels, the continuous-time master equation and the relation between longitudinal and transverse relaxation. A recorded event and an unread detector produce different conditional predictions.

Derive loss of coherence from a larger unitary system

A qubit prepared in $\alpha|0\rangle + \beta|1\rangle$ interacts with an environment initially in $|e\rangle$. A joint unitary correlates $|0, e\rangle \to |0, e_0\rangle,\ |1, e\rangle \to |1, e_1\rangle$, with both environmental states normalised. The joint output is a pure state. Its qubit density matrix follows by tracing over a complete environment basis. Populations stay $|\alpha|^2, |\beta|^2$. The off-diagonal entry is $\alpha\beta^*\langle e_1|e_0\rangle$. Orthogonal environment records remove coherence from all local qubit observations; identical records leave its magnitude intact. Partial distinguishability gives a continuous range in between. Energy need not flow in this controlled dephasing model. An observed reduction in fringe contrast can arise from entanglement with unobserved degrees of freedom even when total evolution is unitary.

Consider instead amplitude loss from an excited qubit to an initially empty environment mode:

$$|0, 0_e\rangle \mapsto |0, 0_e\rangle, \qquad |1, 0_e\rangle \mapsto \sqrt{1-p}|1, 0_e\rangle + \sqrt{p}|0, 1_e\rangle.$$

The two output states remain orthogonal and unit-normalised, so this map on the prepared environment subspace can be extended to a unitary on the full joint space. If environment state k is unobserved, sum the associated unnormalised alternatives in density-matrix form:

$$\rho' = \sum_k K_k \rho K_k^\dagger.$$

Here $K_0 = |0\rangle\langle 0| + \sqrt{1-p}|1\rangle\langle 1|$ and $K_1 = \sqrt{p}|0\rangle\langle 1|$. Since $\sum K_k^\dagger K_k = I$, trace is preserved. The excited population obeys $\rho'_{11} = (1-p)\rho_{11}$ and the coherence $\rho'_{01} = \sqrt{1-p}\rho_{01}$. Averaging state vectors before forming outer products would introduce interference between orthogonal environmental records that have already been traced out.

Take a controlled infinitesimal limit

In a memoryless approximation each small interval dt meets a fresh environment record with vanishing probability of a previously lost excitation returning. Let jump operators L_j have units $1/\sqrt{\mathrm{s}}$, set $K_j = \sqrt{dt}\, L_j$ for events and $K_0 = I - dt(iH/\hbar + \sum L_j^\dagger L_j/2)$. Expanding the completeness relation shows trace preservation through first order in dt. Expand $\sum K_k \rho K_k^\dagger - \rho$, divide by dt and keep the finite terms:

$$\frac{d\rho}{dt} = -\frac{i}{\hbar}[H, \rho] + \sum_j \left(L_j \rho L_j^\dagger - \tfrac{1}{2} L_j^\dagger L_j \rho - \tfrac{1}{2} \rho L_j^\dagger L_j \right).$$

The first term of each dissipative bracket adds probability after an event and the anticommutator subtracts probability from the no-event branch. Their traces cancel. This construction is completely positive, meaning that the same local channel remains positive even when tensored with identity on an unused entangled

ancilla. Positivity only on isolated input states would be too weak: transposition maps density matrices to density matrices on a single system, yet partial transposition can produce a negative eigenvalue when applied to half an entangled pair. The Lindblad form embodies the relevant stronger constraint for the assumed Markovian dynamics.

Independent loss and pure dephasing can be represented by $L_1 = \sqrt{\Gamma_1}|0\rangle\langle 1|$,

$$L_\phi = \sqrt{\Gamma_\phi/2}\,\sigma_z.$$

The factor one half in the dephasing operator fixes our convention for Γ_ϕ. Direct substitution gives $\dot\rho_{11} = -\Gamma_1\rho_{11}$, while $\dot\rho_{01}$ acquires $-(\Gamma_1/2 + \Gamma_\phi)\rho_{01}$ in addition to its Hamiltonian phase. Thus

$$T_1 = \Gamma_1^{-1},$$

$$T_2^{-1} = (2T_1)^{-1} + \Gamma_\phi.$$

At zero temperature with these mechanisms, $T_2 \le 2T_1$. This upper bound is tied to definitions and assumptions. A broad ensemble with static detuning spread, a driven nonequilibrium sample or a refocused pulse sequence may yield an observed envelope with a different shape and timescale. Fitting every contrast decay with one exponential and interpreting that number as the homogeneous T_2 can assign an unphysical dephasing rate.

Update a state when an individual event is recorded

If an ideal detector reports branch j, the probability is $p_j = \operatorname{tr} K_j\rho K_j^\dagger$, and the normalised conditional output is $K_j\rho K_j^\dagger/p_j$. The conditional evolution depends on which environment observable is measured; summing outcomes gives the same physical unconditioned channel. In the loss model, detecting one emission places the qubit in its ground state. Absence of an emission also carries information. At finite duration t, no emission acts by $K_0 = |0\rangle\langle 0| + e^{-\Gamma_1 t/2}|1\rangle\langle 1|$ for the simple undriven zero-temperature case.

From initial $(|0\rangle + |1\rangle)/\sqrt{2}$, at $t = T_1 \ln 2$ the no-event probability is

$$(1 + 1/2)/2 = 3/4.$$

Conditional excited probability is

$$(1/2)/(1 + 1/2) = 1/3.$$

With the event branch in the ground state, the unconditioned excited probability is $(3/4)(1/3) = 1/4$, agreeing with exponential decay of its initial one half. Repeatedly renormalising only no-event trajectories and calling their mean the ensemble population would overestimate survival. If the event detector misses photons, “no click” groups genuine no-emission with undetected emissions and its conditional state is generally mixed.

Use heating and cooling to infer the reservoir

Allow upward and downward transitions at rates $\Gamma_\uparrow, \Gamma_\downarrow$. Rate equations give

$$\dot{p}_1 = \Gamma_\uparrow(1 - p_1) - \Gamma_\downarrow p_1,$$

$$p_{1,\infty} = \frac{\Gamma_\uparrow}{\Gamma_\uparrow + \Gamma_\downarrow},$$

$$p_1(t) = p_{1,\infty} + [p_1(0) - p_{1,\infty}]e^{-(\Gamma_\uparrow + \Gamma_\downarrow)t}.$$

Both upward and downward rates contribute half their sum to coherence decay. If the reservoir is in thermal equilibrium, detailed balance gives

$$\Gamma_\uparrow/\Gamma_\downarrow = e^{-\hbar\omega/(k_B T)}.$$

The population relaxation time measures their sum and equilibrium population measures their ratio, allowing a substantive check of a proposed bath temperature. The relation is a thermodynamic assumption in addition to the two-state rate model. A pumped bath can have the same measured relaxation time with an entirely different stationary population.

For a 5.00 GHz transition, $\hbar\omega/k_B = h(5\,\mathrm{GHz})/k_B \simeq 0.240$ K. A stationary excited fraction 0.10 would imply $T = 0.240/\ln 9 = 0.109$ K under equilibrium. With population relaxation 20 µs the sum of rates is $5.0 \times 10^4\,\mathrm{s}^{-1}$, giving 5000 upward and $45\,000$ downward transitions per second for the respective source populations. Verified independent upward and downward preparations test more than fitting the decay of an initially excited state alone.

Discriminate static detuning, Markovian dephasing and coherent error

A Ramsey experiment prepares an equal superposition, allows free evolution for time t, then converts relative phase to a population by a final analysis pulse. If the detuning δ is fixed in each run but varies randomly between runs, then the ensemble coherence is multiplied by the characteristic function $C(t) = \int p(\delta)e^{-i\delta t}d\delta$. Here δ is an angular frequency. For a Gaussian detuning of mean δ_0 and standard deviation σ,

$$C = e^{-i\delta_0 t}e^{-\sigma^2 t^2/2}.$$

A uniform detuning offset shifts fringe phase without reducing the contrast of a fully coherent, identically repeated preparation. A Lorentzian distribution yields an exponential envelope which can resemble homogeneous Markovian dephasing in a simple time scan. The fitted functional form alone may therefore fail to identify the physical cause of lost contrast.

An ideal instantaneous π pulse about x halfway through free evolution reverses the effect of a static z-detuning on relative phase: $e^{-i\delta t\sigma_z/4}\sigma_x e^{-i\delta t\sigma_z/4} = \sigma_x$, since σ_x anticommutes with σ_z. All static offsets refocus at the echo. Independent white Markovian dephasing accumulated in each half is uncorrelated and cannot be undone in this way. The deterministic Hamiltonian error and a rapidly fluctuating environment can share a short Ramsey contrast while giving different echo contrasts. If the pulse itself has finite duration or an angle error, the algebraic cancellation is no longer exact; repeating at two pulse powers tests that particular assumption.

Take $\sigma/2\pi = 20.0$ kHz and negligible damping. At $t = 10.0\,\mu$s the free Ramsey contrast from Gaussian detuning alone is

$$e^{-(1.257)^2/2} = 0.454.$$

An ideal echo returns unity. Add independent $T_1 = 80\,\mu$s, $T_\phi = 50\,\mu$s. The homogeneous coherence time is

$$[1/(160\,\mu\text{s}) + 1/(50\,\mu\text{s})]^{-1} = 38.1\,\mu\text{s}.$$

The expected echo contrast is now $e^{-10/38.1} = 0.769$ and the ordinary Ramsey contrast 0.349. A single fitted T_2^* cannot display these three separate timescales. Measuring populations, Ramsey fringes and echoes on consistent ensembles gives three complementary observables with independently interpretable limits.

For a time-dependent classical detuning $\delta(t)$, pulse sequences choose a modulation $y(t) = \pm 1$ multiplying the noise phase. The phase error is

$$\phi = \int y(t)\delta(t)dt.$$

For stationary zero-mean Gaussian noise, $\langle e^{-i\phi}\rangle = e^{-\langle\phi^2\rangle/2}$ and the variance is a double integral of the noise correlation weighted by $y(t)y(t')$. Equivalently it is the spectral integral of the power density weighted by $|\int y(t)e^{i\omega t}dt|^2$, with the normalisation set by the chosen two-sided Fourier convention. A single echo has zero weight at $\omega = 0$ because its two half integrals cancel, while keeping finite-frequency sensitivity. Families of pulse spacings can therefore probe noise timescales rather than simply fitting larger and larger numbers to a nominal coherence time.

Find the damped saturation curve from optical Bloch equations

A continuously driven qubit has an additional route to separate population and coherence relaxation. In a rotating frame let $H = \hbar(\Omega\sigma_x + \Delta\sigma_z)/2$, with $z = p_e - p_g$, equilibrium $z_{\text{eq}} = -1$ at zero temperature, and transverse damping $1/T_2$. Choose a consistent Bloch basis so the commutator gives the cross product of $(\Omega, 0, \Delta)$ with (x, y, z). The equations are

$$\dot{x} = -\Delta y - x/T_2,$$
$$\dot{y} = \Delta x - \Omega z - y/T_2,$$
$$\dot{z} = \Omega y - (z - z_{\text{eq}})/T_1.$$

At steady state the first equation gives

$$x = -\Delta T_2 y.$$

The second then gives

$$y = -\Omega T_2 z/(1 + \Delta^2 T_2^2).$$

Inserting into the third gives

$$z = z_{\mathrm{eq}}/[1 + \Omega^2 T_1 T_2/(1 + \Delta^2 T_2^2)].$$

Thus the excited fraction for $z_{\mathrm{eq}} = -1$ is

$$p_{e,\infty} = \frac{\Omega^2 T_1 T_2}{2[1 + \Delta^2 T_2^2 + \Omega^2 T_1 T_2]}.$$

A weak resonant drive gives a population proportional to its power, and a strong continuous drive saturates toward one half rather than preparing a permanently inverted state. Coherent π-pulse inversion is a transient unitary process; continuous steady balance with relaxation gives a different preparation. The full angular-frequency width at half of the on-resonance excitation is $2\sqrt{1 + \Omega^2 T_1 T_2}/T_2$, so increasing power broadens the line even at unchanged intrinsic T_2.

Choose $T_1 = 20$ µs, $T_2 = 16$ µs and resonant $\Omega/2\pi = 10$ kHz. The saturation parameter is

$$s = \Omega^2 T_1 T_2 = 1.26.$$

The stationary excited fraction is $s/[2(1 + s)] = 0.279$, and the ordinary-frequency full width is

$$\sqrt{1 + s}/(\pi T_2) = 29.9\,\mathrm{kHz}.$$

At very low drive it would be 19.9 kHz. A measured linewidth that increases approximately as the square root of power at high drive can have this coherent saturation origin without any increase in the environment's pure-dephasing rate. Comparing the zero-power linewidth extrapolation with an independent time-domain T_2 is a stronger test than fitting a single apparent line broadening.

Connect an individual event stream to an ensemble lifetime

The lifetime T_1 of one excited qubit is a distribution of observed emission times across many trials. With perfect detection, no drive and initially excited state, survival $S(t) = e^{-\Gamma t}$ implies waiting-time density $w(t) = -dS/dt = \Gamma e^{-\Gamma t}$. The mean waiting time and standard deviation are both $1/\Gamma$. A no-click trajectory stays in the excited state conditionally until its first emission when the initial preparation is certainly excited. Its unnormalised weight decays exponentially, but assigning every no-click trajectory a continuously decreasing normalised excitation would confuse conditional and unconditioned statements. For an initial superposition the absence record instead changes the normalised state as derived earlier.

Real acquisitions may stop each trial at a fixed deadline T_{stop}. Treating all no-click trials as though they emitted exactly at the deadline biases lifetime inference. For n observed emissions at times t_j and m censored trials with no emission through T_{stop}, the likelihood is proportional to $\prod_{j=1}^{n} \Gamma e^{-\Gamma t_j} \prod_{k=1}^{m} e^{-\Gamma T_{\mathrm{stop}}}$. Optimising gives

$$\hat{\Gamma} = n/(\sum_j t_j + m T_{\mathrm{stop}}).$$

The numerator is actual events and the denominator all time observed at risk. For 40 detections totalling 480 µs and 10 additional censored runs each observed 30 µs, the estimate is $40/780 = 0.0513\,\mu\mathrm{s}^{-1}$, or

$T_1 \simeq 19.5\,\mu s$. Discarding the censored runs instead would report 12 µs from the mean detected delay. This statistical correction follows directly from the same physical survival law used for the quantum channel.

Detection inefficiency can cause a different data problem. In a single-excitation undriven trial, a missed decay leaves a qubit in its ground state and unable to emit again. The probability of no registered event up to t is $1 - \eta[1 - e^{-\Gamma t}]$ at constant efficiency η, with a nonzero late plateau $1 - \eta$. It is not $e^{-\eta\Gamma t}$, which would describe thinning a renewable stationary Poisson event stream. Long-time survival of the recorded count and independent final-state readout can distinguish a missed emission from a genuinely long-lived excitation. This example shows why the detector model must specify whether events can repeat after one has been missed.

Audit a reconstructed Pauli noise channel

Suppose three independently prepared Bloch components appear to shrink by factors $\lambda_x, \lambda_y, \lambda_z$ after a calibrated waiting time, with no displacement of the centre. A diagonal unital qubit channel can be written $\mathcal{E}(\rho) = \sum_{\nu=0,x,y,z} p_\nu \sigma_\nu \rho \sigma_\nu$ for a physical Pauli mixture. Expanding its action on the three axes and inverting gives

$$p_0 = (1 + \lambda_x + \lambda_y + \lambda_z)/4,$$
$$p_x = (1 + \lambda_x - \lambda_y - \lambda_z)/4,$$

with cyclic partners for p_y, p_z. Complete positivity demands all four numbers be nonnegative. Checking each $|\lambda_i| \le 1$ only checks lengths along single Bloch axes and, for this diagonal map, positivity on a single qubit. It does not guarantee that the reconstructed operation is consistent on half of an entangled system.

For $\lambda = (0.9, 0.9, 0.5)$, every axis appears to shrink, yet $p_z = (1 - 0.9 - 0.9 + 0.5)/4 = -0.075$. Thus the exact diagonal unital channel inferred from those numbers is not completely positive. Finite measurement noise, neglected affine offsets or changing preparation could be responsible. Compare the fit with uncertainty regions and a constrained process model; simply forcing each measured contraction coefficient into $[-1, 1]$ cannot repair this violation. As a constructive control, $\lambda = (0.8, 0.7, 0.6)$ gives $p_0 = 0.775,\ p_x = 0.125,\ p_y = 0.075,\ p_z = 0.025$, all valid and summing to one. A simulator using these four random Pauli operations on entangled inputs then supplies a direct check of the estimated channel beyond single-qubit positivity.

Test how a coherence measurement depends on an ensemble

Compare quasistatic frequency spread with homogeneous decay in a Ramsey sequence. Identify which further measurement could distinguish the mechanisms.

Open this experiment online

Try the calculation

Mean cyclic detuning (Hz) **500**

Shot-to-shot rms spread (Hz) **80**

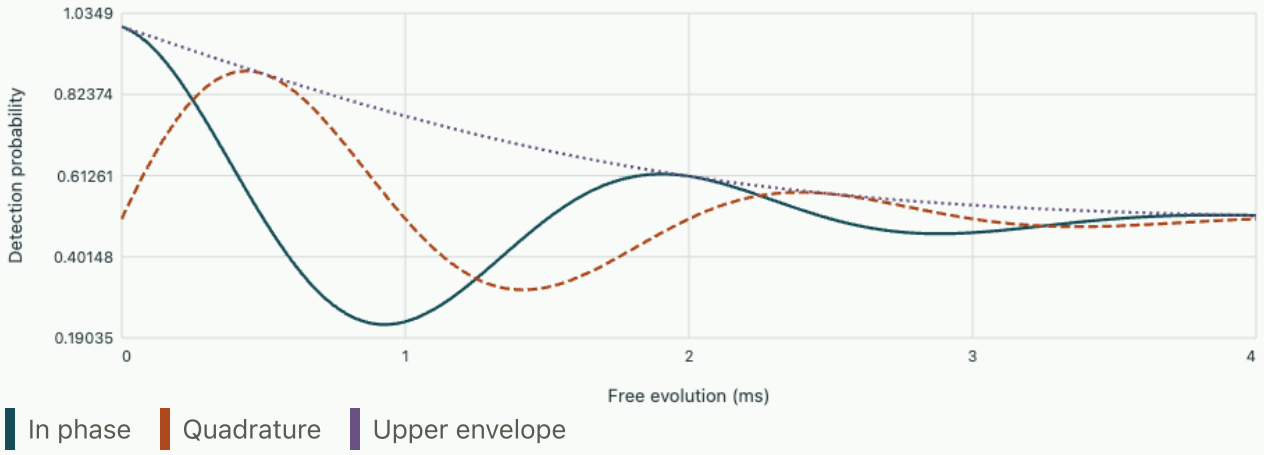

Contrast after 1 ms = 0.53455. The two probabilities at 0.50 ms are 0.5 and 0.87729.

Prove when two different event descriptions give the same channel

A given unconditioned quantum operation does not determine a unique set of environmental records. Suppose a set of Kraus operators K_j describes a channel. If an environment detector resolves an alternative orthonormal basis related by unitary $U = (u_{\alpha j})$, its new system operators are $\tilde{K}_\alpha = \sum_j u_{\alpha j} K_j$. Summing the unnormalised conditional outputs gives

$$\sum_{\alpha,j,k} u_{\alpha j} u^*_{\alpha k} K_j \rho K_k^\dagger = \sum_j K_j \rho K_j^\dagger .$$

This calculation explains why environment-dependent state trajectories can be quite different while every unconditioned prediction on the system agrees. A simulator's choice of individual jump records only describes a particular monitoring scheme when the environment observable actually corresponds to those records.

Consider complete dephasing in the z basis. One representation is $K_0 = |0\rangle\langle 0|,\ K_1 = |1\rangle\langle 1|$, an ideal path measurement with its result forgotten. Rotate the two environment labels by the Hadamard matrix to obtain

$$\tilde{K}_+ = (I/\sqrt{2}),\ \tilde{K}_- = (\sigma_z/\sqrt{2}).$$

These describe equally likely identity or phase-flip operations. For input $|+x\rangle$, the first recorded instrument conditionally gives $|0\rangle$ or $|1\rangle$; the second gives $|+x\rangle$ or $|-x\rangle$. Both produce $I/2$ if their outcome is ignored. For initial $|0\rangle$ the first record is deterministic, while the second remains equally random and leaves the state unchanged up to a phase. Measuring only the system after discarding the labels cannot tell which set was used, but correlations with a retained environmental measurement can immediately tell them apart.

The same distinction is useful for a partially dephasing environment. If $|0\rangle, |1\rangle$ become correlated with environment states of real overlap c, the qubit's transverse Bloch vector contracts by c. For $-1 \leq c \leq 1$ a random phase-flip representation has probabilities $p_I = (1+c)/2,\ p_Z = (1-c)/2$. For $c < 0$ a fringe phase reversal accompanies partial contraction; a reported positive "visibility" loses that sign. One can separate a deterministic phase rotation from the magnitude of random or entangling dephasing if the phase reference is maintained. Over independent identical intervals a real c multiplies repeatedly; the coherence after n steps is c^n. If the same randomly chosen phase flip is instead held fixed throughout all intervals, successive errors are correlated and the even-step coherence can return to one. A single-step channel tomography cannot identify the environmental memory shared across steps.

Tell apart a memoryless phase-flip channel and a fixed hidden phase

Both devices individually implement $p_Z = 0.20$ over a calibrated interval for freshly prepared qubits. In device A the phase-flip draws at successive intervals are independent. In device B one draw is held fixed for both consecutive intervals. Predict $+x$ readout after one and two intervals for an initial $|+x\rangle$.

▼ **Derive and check the result**

One interval leaves $x = 1 - 2p_Z = 0.60$, so both devices give

$$P(+x) = 0.80.$$

Independence in A gives $x = 0.60^2 = 0.36$ after two intervals, or $P(+x) = 0.68$. For B, two identical phase flips multiply to identity whether the hidden choice was identity or Z, hence $P(+x) = 1$. In both devices the one-step process applied to an independently prepared ensemble is consistent with the measured completely positive map. Composing that map twice implicitly assumed fresh, independent environment conditions; a memory experiment tests exactly that assumption. Observing a revival thus need not imply an energy input, a negative probability or a nonunitary failure in the larger coupled system.

Distinguish dipole coherence from an excited-state population lifetime

Consider a single initially excited emitter in a memoryless zero-temperature environment with radiative population decay Γ and no additional dephasing. The probability density of emitting at time $t \geq 0$ is $\Gamma e^{-\Gamma t}$. A normalised single-photon temporal mode consistent with this ideal undistinguished radiative channel is $g(t) = \sqrt{\Gamma}\, e^{-\Gamma t/2} e^{-i\omega_0 t}$ on $t \geq 0$. Its field amplitude decays twice as slowly as the population probability. Use the Fourier convention $\tilde{g}(\omega) = (2\pi)^{-1/2} \int g(t) e^{i\omega t} dt$. The half-line exponential integral gives

$$|\tilde{g}(\omega)|^2 = \frac{\Gamma}{2\pi[(\Gamma/2)^2 + (\omega - \omega_0)^2]}.$$

It integrates to one in angular frequency. At a detuning of $\Gamma/2$ the intensity has half its peak, so the full angular-frequency linewidth is Γ, or $\Gamma/(2\pi)$ in hertz. Using the population decay directly as a field amplitude would predict a linewidth twice too large and the wrong wave-packet envelope. The transform-limited result requires a single radiative mode in the sense of a pure emitted temporal state and absence of additional

phase randomisation. Extra pure dephasing broadens the spectrum while population decay may remain unchanged.

For $\Gamma^{-1} = 10.0$ ns the ideal linewidth is 15.9 MHz. Add independent homogeneous pure dephasing at $\Gamma_\phi = 2.0 \times 10^7\ \mathrm{s}^{-1}$. The first-order dipole correlation decays at $\Gamma/2 + \Gamma_\phi$, giving a full line width $(\Gamma/2 + \Gamma_\phi)/\pi = 22.3$ MHz. The photon detection time distribution can still have 10 ns decay. A measured line twice its nominal transform-limited width might indicate pure dephasing, but spatially varying frequencies, slow source drift or spectral resolution can also broaden the observed distribution. A lifetime measurement alone does not calibrate those effects.

The acquisition window has an independently calculable consequence. If only photons arriving before time T are retained, their normalised temporal mode is a truncated exponential divided by $\sqrt{1 - e^{-\Gamma T}}$ in this ideal pure-emission model. Its Fourier amplitude gains factor $1 - e^{-[\Gamma/2 - i(\omega - \omega_0)]T}$ in the numerator. At $T \ll 1/\Gamma$ the shape approaches the squared sinc of a short rectangular temporal window. Selecting only early photons can thus broaden the spectrum even though the emitter and its environment are unchanged. A time gate, an energy filter and unconditioned counting describe different optical preparations and different uncertainty budgets.

Infer the transform-limited fraction from a linewidth and test the model

A homogeneous line has measured weak-drive full width 25.0 MHz and population lifetime 10.0 ns. Compute T_2 and the implied pure dephasing rate assuming ideal Lorentzian broadening. Which second measurement would discriminate this inference from static sample inhomogeneity?

▼ **Derive and check the result**

With angular half width $1/T_2$, the frequency full width is $1/(\pi T_2)$, so $T_2 = 12.73$ ns. Subtracting $1/(2T_1) = 5.00 \times 10^7\ \mathrm{s}^{-1}$ from $1/T_2 = 7.85 \times 10^7\ \mathrm{s}^{-1}$ gives

$$\Gamma_\phi = 2.85 \times 10^7\ \mathrm{s}^{-1}.$$

The coherence relative to the radiative limit is

$$T_2/(2T_1) = 0.637.$$

Compare a time-domain echo, a single-emitter linewidth or spectra over short acquisition segments with the ensemble line. A recoverable component of contrast or slowly wandering line centres would indicate a static or slow contribution instead of assigning every megahertz to homogeneous phase damping.

Tutorial · distinguish different sources of coherence loss

Recover a pure-dephasing rate and assess a conflicting fit

A population measurement gives

$$T_1 = 30.0\,\mu\mathrm{s};$$

homogeneous Ramsey contrast gives

$$T_2 = 24.0\,\mu\mathrm{s}.$$

Find T_ϕ. What would a reported $T_2 = 80\,\mu$s imply in the same model?

▼ **Work through the solution**

Subtract $1/(60\,\mu\mathrm{s})$ from $1/(24\,\mu\mathrm{s})$ to get $\Gamma_\phi = 0.0250\,\mu\mathrm{s}^{-1}$, hence $T_\phi = 40.0\,\mu$s. Replacing 24 by 80 yields a negative pure dephasing rate, incompatible with the specified independent time-homogeneous channels. Check the fitted envelope, uncertainties, thermally shifted equilibrium population and whether the two times were measured under comparable conditions. A negative fitted number is evidence about the model or data interpretation, not a physically negative random-event rate.

Calculate an inefficient no-click update

The initial excited probability is one half, the decay probability for an initially excited qubit during the window is 0.40, and a true emitted excitation is detected with efficiency 0.75. Find the probability of no click, its conditional excited fraction and the unconditioned population.

▼ **Work through the solution**

A real emission occurs with probability $0.5 \times 0.40 = 0.20$; a click is seen with probability 0.15, so no click has probability 0.85. Remaining excited and no click occurs with probability $0.5(1 - 0.40) = 0.30$. Conditional excitation is $0.30/0.85 = 0.3529$; the unconditioned excitation is 0.30 since clicked branches are ground. For an initially coherent preparation the no-click off-diagonal term must also be updated by grouping the undetected-emission and true-no-emission Kraus outcomes. Simply replacing p by detector efficiency times p would predict the wrong physical excitation loss.

Diagnose a nonphysical time step and repair it

Apply one forward-Euler step to $\dot{p}_1 = -\Gamma p_1$ with initially excited state and $\Gamma\Delta t = 1.2$. Compute the prediction and an exact completely positive alternative. Would enforcing unit trace fix Euler's failure?

▼ **Work through the solution**

Euler gives $p_1 = 1 - 1.2 = -0.2,\ p_0 = 1.2$, already with exact trace one. The finite-time amplitude-damping channel has $p = 1 - e^{-1.2} = 0.6988$, giving $p_1 = 0.3012$. For this simple model the exact Kraus step can be applied repeatedly at arbitrary duration. In a driven model an operator-splitting scheme with unitary Hamiltonian substeps and completely positive damping substeps maintains positivity while introducing a controllable splitting error. Check convergence by halving the step and comparing with an independent observable. Renormalising trace or clipping a negative eigenvalue changes the numerical trajectory without diagnosing the original time-step error.

Further reading

- Cambridge · open systems, measurements and the Lindblad equation
- Course investigation · density matrices and dephasing
- Course investigation · driven coherent response

Quantum states, dynamics and measurement · 12. How a turning point supplies the missing semiclassical phase

Far from a turning point the wavelength follows the classical momentum. At the turning point that approximation becomes singular. Derive the Airy connection, quantise a smooth well, distinguish a hard wall and estimate a tunnelling exponent with a defensible error check.

Expand the action before interpreting a local wavelength

For a stationary one-dimensional state with smooth potential $V(x)$, write $\psi = e^{iS(x)/\hbar}$. Substitution into Schrödinger's equation gives

$$S'^2 - i\hbar S'' = 2m(E - V) = p^2(x).$$

Expand $S = S_0 + (\hbar/i)S_1 + \cdots$. The leading terms give $S_0' = \pm p$ and $S_1' = -p'/(2p)$ in an allowed region where $p > 0$. Thus travelling waves have form $Cp^{-1/2}\exp[\pm i\int^x p\,dx/\hbar]$. The amplitude factor is physically essential: for a single such branch the current to leading order is $\pm|C|^2/m$, independent of x, as required for steady flux in a real potential. A constant amplitude with variable wavenumber would predict current piling up in a stationary region.

A local condition for this expansion is $\hbar|p'|/p^2 \ll 1$ together with control over higher derivatives on the wavelength scale. Away from a barrier the phase can remain accurate even if absolute amplitude normalisation must be fixed globally. Where $E < V$, define $\kappa = \sqrt{2m(V - E)}$ with units of momentum. The local solutions become $\kappa^{-1/2}\exp[\pm\int \kappa\,dx/\hbar]$. A single real decaying branch carries zero probability current by itself. Stationary tunnelling flux through a finite barrier involves a complex combination satisfying both exterior matching conditions, even though the leading size of transmission can often be inferred from the decaying exponential.

Resolve the apparent divergence near a simple turning point

Let x_t satisfy

$$V(x_t) = E,$$

with slope $F = V'(x_t) > 0$. Linearise $V - E \simeq F(x - x_t)$. The local Schrödinger equation becomes

$$\psi'' = (2mF/\hbar^2)(x - x_t)\psi.$$

Introduce length $\ell = (\hbar^2/(2mF))^{1/3}$ and $\xi = (x - x_t)/\ell$, leaving $d^2\psi/d\xi^2 = \xi\psi$. Its two independent solutions are Ai and Bi. On the positive side $\mathrm{Ai}(\xi)$ decays, whereas $\mathrm{Bi}(\xi)$ grows. A bound state extending into a semi-infinite forbidden region selects the former. At $\xi = 0$ the Airy function is finite; the apparent divergence $p^{-1/2}$ came from extending the outer approximation beyond its range.

The asymptotic forms for $s > 0$ large are

$$\mathrm{Ai}(s) \sim \frac{e^{-2s^{3/2}/3}}{2\sqrt{\pi}\,s^{1/4}}, \qquad \mathrm{Ai}(-s) \sim \frac{\sin(2s^{3/2}/3 + \pi/4)}{\sqrt{\pi}\,s^{1/4}}.$$

For a linear slope, $2s^{3/2}/3$ equals $|\int_{x_t}^{x} p\,dx|/\hbar$ on the allowed side or $\int_{x_t}^{x} \kappa\,dx/\hbar$ on the forbidden side. Matching in an interval where $|x - x_t| \gg \ell$ but the potential remains effectively linear gives the connection from a decaying tail of coefficient C to $2C\,p^{-1/2}\sin[\hbar^{-1}\int_{x}^{x_t} p\,dx + \pi/4]$. The coefficient two and phase $\pi/4$ both come from the same local solution. A turning point with zero slope, a sharp step or two turning points closer than their Airy lengths requires a different local treatment.

Join two decaying tails to a standing wave

For a smooth well with two well-separated simple turning points x_1, x_2, the left tail matches to $\sin[\hbar^{-1}\int_{x_1}^{x} p\,dx + \pi/4]$. The right tail matches to $\sin[\hbar^{-1}\int_{x}^{x_2} p\,dx + \pi/4]$. Their zeros and logarithmic derivatives must coincide throughout the overlap, up to an overall sign. The sum of the two phase arguments is the full one-way action divided by $\hbar$ plus $\pi/2$, which must be an integer multiple of π. Taking $n \geq 0$ gives

$$\int_{x_1}^{x_2} p(x, E_n)\,dx = \pi\hbar(n + 1/2).$$

For a harmonic oscillator the turning point magnitude is $a = \sqrt{2E/(m\omega^2)}$ and $p = m\omega\sqrt{a^2 - x^2}$. The action is the area of a semicircle times $m\omega$, equal to $\pi E/\omega$. The rule reproduces $E_n = \hbar\omega(n + 1/2)$ exactly for this special spectrum, even though its outer WKB wavefunction is still inaccurate near each turning point and globally for low quantum number. Equality of a spectral result does not imply equality of spatial densities or agreement of all observables.

Different boundaries change the phase budget. Take a hard wall at $x = 0$ and a linearly rising potential Fx for $x > 0$. The solution decaying to the right is $\mathrm{Ai}[(x - E/F)/\ell]$. Its exact energies follow from $\mathrm{Ai}(-E/F\ell) = 0$. The wall imposes a node without an Airy phase at the left. Setting the right-matched standing wave to zero at the wall gives

$$\int_{0}^{E/F} p\,dx = \pi\hbar(n + 3/4).$$

Thus $E/(F\ell) \simeq [3\pi(n + 3/4)/2]^{2/3}$. For $n = 0$ this is about 2.320 versus the first exact zero 2.338; with $n = 4$ the approximation is about 7.94. Replacing the wall by a second smooth turning point would instead introduce an erroneous half offset.

Separate a credible exponential estimate from its prefactor

For a barrier with two simple turning points x_a, x_b, a large forbidden action $I = \int_{x_a}^{x_b} \kappa\,dx/\hbar$ suppresses the transmitted flux approximately as $T \simeq e^{-2I}$, up to prefactors and corrections that depend on geometry and matching. For a triangular forbidden region with height $V_0 - E = \Delta$ at $x = 0$ falling as $\Delta - Fx$ to zero at Δ/F, its action is $I = 2\sqrt{2m}\Delta^{3/2}/(3\hbar F)$. A discontinuous incident-side boundary may affect the prefactor substantially. It does not justify assigning a second smooth-turning-point phase to the discontinuity. For a rectangular barrier the exact thick-barrier prefactor is $16E(V_0 - E)/V_0^2$ times the decaying exponential at equal exterior potentials. It can exceed one while total transmission remains small.

Suppose two measurements change only the forbidden width of a flat barrier from d to $d + \Delta d$, keeping both interfaces and asymptotic energies the same. In the thick limit the ratio $T(d + \Delta d)/T(d) \simeq e^{-2\kappa\Delta d/\hbar}$

cancels the common prefactor. If $\kappa/\hbar = 2.0\,\mathrm{nm}^{-1}$ and $\Delta d = 0.50$ nm, this ratio is

$$e^{-2} = 0.1353.$$

An offset error in barrier width that leaves this ratio unchanged can still cause a large error in absolute transmission. Near the top two turning points can merge, I becomes small, and the raw exponential lacks a unitary connection to the over-barrier reflection regime. Use a suitable uniform approximation or a direct scattering solution there.

Predict a high-energy spectrum without a closed-form wavefunction

Consider a pure quartic trap $V(x) = \lambda x^4$ with $\lambda > 0$. Classical turning points lie at $x = \pm a,\ a = (E/\lambda)^{1/4}$. Unlike the harmonic oscillator its period changes with energy, so its levels cannot be equally spaced. The semiclassical action can be evaluated once in dimensionless coordinates and then rescaled to any m, λ, E:

$$A(E) = \int_{-a}^{a} \sqrt{2m(E - \lambda x^4)}\, dx = 2\sqrt{2m}\, \lambda^{-1/4} E^{3/4} I_4,$$

$$I_4 = \int_0^1 \sqrt{1 - y^4}\, dy.$$

With $s = y^4$,

$$I_4 = \frac{1}{4} B(1/4, 3/2) \simeq 0.8740.$$

Quantising $A(E_n) = \pi\hbar(n + 1/2)$ predicts $E_n \simeq [\pi\hbar(n + 1/2)\lambda^{1/4}/(2\sqrt{2m} I_4)]^{4/3}$. The dependence $E \propto n^{4/3}$ and the dimensional factor $\hbar^{4/3}\lambda^{1/3}m^{-2/3}$ can be checked independently from scaling the full Schrödinger equation. That dimensional argument does not supply the coefficient or the half-integer offset. The action and boundary phases supply the additional physics.

Choose units such that $H = -d^2/dx^2 + x^4$. Then $2m = \hbar = 1$ in the scaled equation, and the leading energies are $[\pi(n + 1/2)/(2I_4)]^{4/3}$. For $n = 0$ this gives about 0.868, below the accurately computed ground value around 1.060. An independent variational Gaussian $\psi = (\beta/\pi)^{1/4} e^{-\beta x^2/2}$ has energy $\beta/2 + 3/(4\beta^2)$, minimised at $\beta^3 = 3$ and giving 1.082. Its proximity to the true ground energy has a different status: the Gaussian energy is an upper bound by the variational principle, while leading WKB provides no universal bound for this potential. Using agreement between two approximate numbers as proof of a rigorous interval would require actual lower and upper estimates.

At larger n the period relation provides an additional prediction. Since $A \propto E^{3/4}$,

$$T_{\mathrm{cl}} = 2\, dA/dE \propto E^{-1/4};$$

the frequency increases at larger amplitude, and the local spacing should grow as $E^{1/4}$. Numerically compute a handful of high levels with a direct finite-difference or spectral Hamiltonian, then plot $E_n^{3/4}$ against $n + 1/2$. A nearly linear trend is a more sensitive diagnostic of the predicted law than comparing a single apparent power from a log plot at low n. Increase domain size and basis cutoff independently. A truncation that

confines high-energy tails can falsely steepen the spacing, while a coarse grid can flatten the kinetic-energy dispersion. Both errors become more visible for rapidly oscillating states.

Relate a barrier action to a coherent population oscillation

Take two identical minima separated by a smooth barrier. At energies below the barrier top but above each isolated-well ground, there are two classically allowed intervals and a central forbidden region. If the barrier is sufficiently opaque, construct approximate localised states $|L\rangle, |R\rangle$ with equal one-well energy E_0. In their span the simplest symmetric Hamiltonian is $H = E_0 I - J(|L\rangle\langle R| + |R\rangle\langle L|)$ with $J > 0$ in a suitable phase convention. The even and odd combinations have energies $E_0 - J, E_0 + J$. The tunnel splitting is $2J$, set to leading exponential order by the overlap of amplitudes through the central forbidden interval.

A single under-barrier wavefunction amplitude carries e^{-I},

$$I = \int \kappa dx/\hbar.$$

For an open barrier the transmission probability is proportional to e^{-2I}. For a symmetric double-well splitting the relevant quantity is a coherent off-diagonal energy amplitude, so its leading exponential is $J \propto e^{-I}$, times a prefactor involving the oscillation scales and matching in the two wells. Confusing these exponents can lead to a severe error when using a calculated spatial barrier to predict a coherent tunnel frequency. The exact prefactor depends on the potential and turning-point normalisations. A direct symmetric numerical eigenproblem giving the difference of two nearly equal energies requires enough precision to resolve an exponentially small splitting.

Prepare $|L\rangle$ at time zero. Its evolution under the two-state Hamiltonian gives

$$P_R(t) = \sin^2(Jt/\hbar);$$

the first full transfer time is $\pi\hbar/(2J)$. If the observed energy splitting expressed as ordinary frequency is $2J/h = 20$ Hz, the first transfer occurs at 25 ms. Changing only the barrier action by $+0.70$ at fixed prefactor multiplies the splitting by $e^{-0.70} = 0.497$ and approximately doubles the transfer time. An open-barrier transmission experiment over the same change instead sees factor $e^{-1.4} = 0.247$. Both tests depend exponentially on barrier shape but interrogate different quantities.

Now add a calibrated tilt $\epsilon = E_L - E_R$. The 2×2 Hamiltonian after trace subtraction is $(\epsilon/2)\sigma_z - J\sigma_x$, and the maximum transfer falls to $4J^2/(\epsilon^2 + 4J^2)$. At $\epsilon = 2J$ it is one half while the oscillation angular frequency increases from $2J/\hbar$ to $2\sqrt{2}J/\hbar$. A loss of transfer contrast need not be environmental decoherence; static tilt alone yields this combination of increased frequency and reduced amplitude. Measuring both quantities and reversing the tilt tests the two-level model. If a higher intrawell level lies within the tilt or excitation range, a single tunnelling doublet no longer spans the relevant motion and a larger basis is required.

Understand the angular correction at a radial turning point

In three dimensions the reduced radial equation on $r > 0$ has a centrifugal term $\hbar^2 l(l+1)/(2mr^2)$. Applying the flat-coordinate WKB rule directly to this equation can give incorrect phases at its singular endpoint, especially for low l. A way to see the appropriate angular correction without guessing is to set

$r = r_0 e^x$, mapping the half-line to all real x, and write $u(r) = e^{x/2} w(x)$. The arbitrary reference r_0 keeps the logarithm dimensionless. Since $d/dr = r^{-1} d/dx$, differentiating twice yields

$$\frac{d^2u}{dr^2} = \frac{e^{x/2}}{r^2}\left(w'' - \frac{1}{4}w\right).$$

The resulting equation for w has squared local action momentum in x proportional to $2mr^2(E - V) - \hbar^2(l + 1/2)^2$. Thus returning its action $\int p_x dx$ to r uses $l(l+1) \to (l+1/2)^2$ in the radial semiclassical momentum. This is the Langer correction. It does not change the exact angular eigenvalue $L^2 = \hbar^2 l(l+1)$; it accounts for the coordinate transformation and endpoint behaviour of the semiclassical representation. Confusing those statements would alter the exact radial equation to repair an approximate integration rule.

An instructive check is the isotropic oscillator. For $V = m\omega^2 r^2/2$ and effective classical angular momentum $L = (l + 1/2)\hbar$, the radial action over its two roots is

$$A_r = (\pi/2)(E/\omega - L).$$

One can obtain this without an explicit contour integral by noting that radial motion repeats at twice the Cartesian frequency, so

$$dA_r/dE = \pi/(2\omega).$$

Fix the integration constant at the circular-orbit threshold $E = \omega L$, where the two radial turning points meet and $A_r = 0$. Quantising $A_r = \pi\hbar(n_r + 1/2)$ gives $E = \hbar\omega(2n_r + l + 3/2)$, matching the exact three-dimensional oscillator spectrum. For $l = 0$ the naive uncorrected centrifugal term has no inner turning point; transferring a generic two-smooth-turning-point formula into that half-line problem without reconsidering its origin is especially unsafe.

Calculate an exponentially small transmission without subtracting nearly equal numbers

In a classically forbidden interval the local solution contains a growing and a decaying branch. In a long interval numerical transfer matrices can contain elements as large as e^I, while the quantity sought is as small as e^{-2I}. Straight propagation of a decaying boundary condition in the unstable direction can contaminate it with a tiny growing roundoff component which eventually dominates. A plot matching the expected slope over the first few decay lengths gives no guarantee that a final tiny flux is accurate. Numerically stable alternatives include propagating a logarithmic derivative with controlled handling of its poles, matching independently normalised left and right solutions, or composing scattering matrices for short slices while retaining decaying propagation factors.

For a square barrier the exact transmission at $E < V_0$ in equal exterior regions is $T = [1 + V_0^2 \sinh^2(\kappa d/\hbar)/(4E(V_0 - E))]^{-1}$. With $E = V_0/2$ this simplifies to $\operatorname{sech}^2 I$. At $I = 20$ its value is approximately $4e^{-40} = 1.70 \times 10^{-17}$. Computing $T = 1 - R$ in ordinary double precision can give exactly zero since the reflection probability rounds to one. Compute transmitted amplitude or its logarithm directly and compare $\log T \simeq \log 4 - 2I$. At a moderate I, where both currents are accurately representable,

independently verify $R + T = 1$. Moving the benchmarking point across the barrier-opacity range tests both the physical approximation and the floating-point representation of the desired observable.

Resolve a wave through a classical turning point

Vary the physical length scale in a locally linear potential. Compare the numerical Airy solution and the two semiclassical branches near and away from the turn.

Open this experiment online

Try the calculation

Turning-point length (nm) **0.8**

Turning-point position (nm) **0.4**

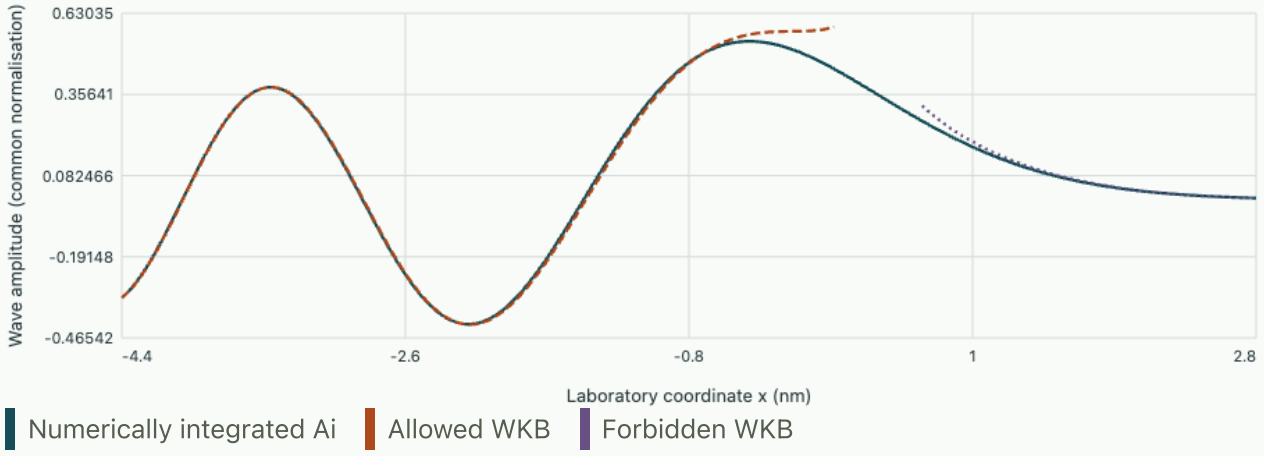

The turn at 0.4 nm has finite Ai value 0.35503. The characteristic transition length is 0.8 nm; the asymptotic curves omit the nearest 0.35 lengths.

Connect semiclassical action to counting states in phase space

The action quantisation rule can also be read as a statement about the number of energy levels below a chosen energy. In a one-dimensional smooth confining well the classically accessible region in (x, p) has area $\Omega(E) = \int dx \int_{-p(x)}^{p(x)} dp = 2A(E)$. Roughly one stationary state fits per phase-space area $2\pi\hbar$, with an endpoint correction from the turning-point phases. The smoothed density of states follows by differentiating $\bar{N}(E) \simeq \Omega(E)/(2\pi\hbar)$, giving $\bar{\rho}(E) = T_{\mathrm{cl}}/(2\pi\hbar)$. This counts levels after smoothing over individual spacing; the exact count is a staircase and its derivative a sum of delta functions. A formula for a smooth density cannot predict each state in a near-degenerate doublet without additional phase and tunnelling information.

For a particle in a two-dimensional rigid rectangular region, apply the same counting idea in four-dimensional phase space. With area $\mathcal{A} = L_x L_y$, accessible momentum disk $p_x^2 + p_y^2 < 2mE$, the scalar smoothed count is $\bar{N}(E) \simeq \mathcal{A}\pi(2mE)/(2\pi\hbar)^2 = \mathcal{A}mE/(2\pi\hbar^2)$. Its derivative is a constant, unlike the $E^{-1/2}$ free one-dimensional density and the $\sqrt{E}$ three-dimensional density. This leading area law can be checked by directly

counting positive integer lattice points satisfying $(n_x\pi/L_x)^2 + (n_y\pi/L_y)^2 < k^2$ in a Dirichlet rectangle. It works after smoothing when many modes fit in the domain and makes no assumption that the rectangular shape itself models all possible billiards.

A more sensitive check accounts for the excluded axes in that first-quadrant lattice count. For a Dirichlet rectangle subtract half of the positive-axis contributions at each boundary of the quadrant. In terms of perimeter $\mathcal{P} = 2(L_x + L_y)$, the resulting first two terms are $\bar{N}(k) \simeq \mathcal{A}k^2/(4\pi) - \mathcal{P}k/(4\pi) + O(1)$ after appropriate smoothing, with fluctuations from discrete modes beyond this local expansion. Neumann boundaries reverse the sign of the perimeter term because modes constant along one coordinate are admitted. Area alone determines the bulk growth, whereas boundary condition affects the next correction. The finite corner term and spectral fluctuations require more care than assigning one universal integer offset to every billiard.

The distinction between smoothing and exact prediction also explains why counting every computed eigenvalue only up to an arbitrarily sharp threshold can produce apparent convergence jumps as the mesh improves. If a true eigenvalue lies near the threshold, a small numerical shift changes an integer count by one. Compare a range of cumulative counts, track nearby eigenvalue residuals and study the smoothed density with a bandwidth large enough to include several levels but small compared with the energy scale of interest. Repeating for two geometries of equal area and different perimeter tests actual geometric information beyond an overall energy scale.

Measure the boundary correction in a square cavity

A scalar mode calculation uses a square of side 10.0 length units and reports all Dirichlet modes with $k < 5.00$ inverse length units. Find the leading area estimate, add the perimeter correction, and compare with a direct lattice-point count.

▼ **Derive and check the result**

The area contribution is $100(25)/(4\pi) = 198.94$; the perimeter subtraction is $40(5)/(4\pi) = 15.92$, giving a smoothed estimate 183.03 before corner and oscillatory terms. The exact count uses positive integers n_x, n_y satisfying $n_x^2 + n_y^2 < (50/\pi)^2 \simeq 253.30$. For each $n_x = 1, \ldots, 15$, count $n_y \le \lfloor \sqrt{253.30 - n_x^2} \rfloor$ and sum. The values are $15, 15, 15, 15, 15, 14, 14, 13, 13, 12, 11, 10, 9, 7, 5$, giving 183. Agreement at this threshold is close, while the leading area term alone is about sixteen high. Nearby thresholds can have different residual fluctuations, so the close integer agreement at this one wave number should not be treated as an exact two-term counting rule. The same calculation on a different aspect ratio at fixed area would change both the perimeter and the lattice pattern.

Measure a reflection phase separately from the travel action

In a one-dimensional cavity a quantisation condition can be written as a round-trip phase budget. Let an allowed region accumulate one-way phase $\theta(E)$ and let reflections at its left and right bounding regions have phase shifts $\phi_L(E), \phi_R(E)$ for unit-modulus reflection. A self-reproducing wave requires

$$2\theta(E) + \phi_L(E) + \phi_R(E) = 2\pi n,$$

with an integer depending on the chosen phase origin. A rigid Dirichlet wall reflects with amplitude -1; a smooth decaying turning point contributes a different reflection phase that the Airy matching fixed. Adding 2π to one reflection phase simply relabels the integer, but an energy-dependent boundary phase alters the spacing and cannot be removed by a constant relabelling.

A useful boundary that can be solved without approximation is a finite step $V = 0$ for $x < 0,\ V = V_0 > E$ for $x > 0$. With incident wave e^{ikx}, reflected amplitude re^{-ikx} and exterior evanescent $te^{-\kappa x}$, where k, κ here are inverse lengths, continuity of the value and derivative gives

$$1 + r = t,\ ik(1 - r) = -\kappa t.$$

Solving yields

$$r = (k - i\kappa)/(k + i\kappa).$$

Its magnitude is one and its phase can be written $-2\arctan(\kappa/k)$. As V_0 tends to infinity with E fixed, the phase tends to $-\pi$, giving the rigid-wall limit. Just below the step threshold it approaches zero. Two mirrors with equally complete reflection can therefore produce different cavity eigenfrequencies.

Near the rigid-wall limit expand $r \simeq -e^{2ik/\kappa}$. To leading order the exterior behaves like a rigid reflecting plane displaced a penetration length $1/\kappa$ beyond the interface. For a deep symmetric finite well of width L, the phase condition approximately replaces L by $L + 2/\kappa(E)$. This shift lowers levels relative to a rigid well of the same geometric width. It is only an approximation because penetration depends on energy. Matching logarithmic derivatives independently on the two sides is preferable near threshold or when the barrier shape varies over the penetration length.

Infer a boundary penetration from a measured free spectral interval

A long one-dimensional cavity with identical deep finite barriers has a measured low-mode wavenumber interval 2% smaller than π/L. Under the constant penetration approximation infer $1/\kappa$ relative to L and identify a check on the approximation.

▼ **Derive and check the result**

Set $\pi/(L + 2/\kappa) = 0.98\pi/L$. Then $2/\kappa L = 1/0.98 - 1 = 0.0204$, or a penetration distance $1/\kappa = 0.0102L$ at each wall. For $L = 100$ nm this is about 1.02 nm. Measure several successive intervals to see whether one penetration length describes them. An increasing penetration with energy changes the reflection-phase derivative and therefore the spacing; the physical group delay then contains travel time in the allowed region plus boundary delay. A geometrically longer rigid cavity would produce a similar single interval but a different energy dependence, so an isolated observation does not uniquely identify the cause.

Tutorial · determine what a semiclassical estimate can certify

Recover the period and level spacing from an action curve

For $A(E) = \int_{x_1}^{x_2} p\,dx$ in a smooth bound well, derive dA/dE and infer the local energy spacing. Apply it to a 100 kHz oscillator and express spacing in both joules and electronvolts.

▼ **Work through the solution**

The endpoint derivative terms vanish because $p = 0$ at both ends. Inside, $\partial p/\partial E = m/p = 1/v$, so dA/dE equals the classical one-way travel time $T_{\mathrm{cl}}/2$. From the quantisation rule increasing n by one gives local spacing $\Delta E \simeq 2\pi\hbar/T_{\mathrm{cl}} = h f_{\mathrm{cl}}$. At 100 kHz this is 6.626×10^{-29} J or 4.136×10^{-10} eV. In an anharmonic potential the classical period generally depends on energy; use the local slope or integrate between adjacent solutions rather than assuming a universal ladder spacing.

Check whether the Airy overlap interval exists

For a model atom with mass 1.44×10^{-25} kg let $F = 1.0 \times 10^{-21}$ N at the turning point. Find the Airy length. If $|V''| = 1.0 \times 10^{-16}\,\mathrm{N/m}$, estimate the relative quadratic correction at 5ℓ.

▼ **Work through the solution**

The length is $(\hbar^2/2mF)^{1/3} \simeq 3.38 \times 10^{-8}$ m. The ratio of quadratic potential term to linear term at 5ℓ is $|V''|5\ell/(2F) \simeq 0.00845$. Thus several Airy lengths can fit in a region with small curvature correction for these numbers. The parameter $|x - x_t|/\ell = 5$ should still be checked against the desired asymptotic precision; the result supports an overlap interval but does not guarantee a prescribed decimal accuracy. Bringing a second turning point within a few such lengths changes the matching problem qualitatively.

Compare a triangular action with a rectangular shortcut

A linearly falling barrier has initial forbidden wavenumber $\kappa_0/\hbar = 3.0\,\mathrm{nm}^{-1}$ and width $d = 2.0$ nm. Compare its exponential estimate with that of a rectangle using initial height throughout. What would a rectangle at the average potential height predict?

▼ **Work through the solution**

For the triangle, $\kappa(x) = \kappa_0\sqrt{1 - x/d}$, so $I = (2/3)(\kappa_0/\hbar)d = 4.0$, giving $e^{-8} = 3.35 \times 10^{-4}$. The initial-height rectangle gives $I = 6$ or $e^{-12} = 6.14 \times 10^{-6}$. Averaging the potential height first gives wavenumber $\kappa_0/\sqrt{2}$, $I = 4.24$ and $e^{-8.49} \simeq 2.06 \times 10^{-4}$. The action averages the square root of the local energy gap, not the gap itself. Agreement in mean barrier height is insufficient for predicting an exponential observable. Exact matching or a direct numerical probability current would be needed to determine the prefactor.

Further reading

- Fließbach · Näherungsmethoden
- Cambridge · WKB, Airy matching and semiclassical quantisation
- Course investigation · semiclassical quantisation and tunnel splitting

Quantum states, dynamics and measurement · 13. Measure geometric phase from three spin overlaps

A numerical eigensolver is free to choose a new complex phase for every eigenvector. A product of three measured overlaps removes these arbitrary choices. Relate its argument to a spherical triangle, then see why a vanishing overlap destroys phase information.

Close a loop before assigning a phase

Choose three normalised spinors $|A\rangle, |B\rangle, |C\rangle$. Their individual phases change under $|j\rangle \mapsto e^{i\gamma_j}|j\rangle$. In the cyclic product $Z = \langle A|B\rangle\langle B|C\rangle\langle C|A\rangle$ every phase occurs once with each sign. This product is a gauge-independent complex quantity although none of the three overlap phases separately is.

Take A at the north pole of the Bloch sphere, B on its equator at longitude zero and C on the equator at longitude φ with $-\pi < \varphi < \pi$. Choose $A = (1, 0)$, $B = (1, 1)/\sqrt{2}$ and $C = (1, e^{i\varphi})/\sqrt{2}$. Direct multiplication gives

$$Z = \frac{1 + e^{i\varphi}}{4} = \frac{\cos(\varphi/2)}{2} e^{i\varphi/2}.$$

The oriented triangle bounded by the two meridians and the shortest equator arc has solid angle $\Omega = \int_0^{\varphi} d\phi \int_0^{\pi/2} \sin\theta \, d\theta = \varphi$. The overlap product therefore has argument $+\Omega/2$ for this order. The geometric phase of a ket parallel-transported along the oriented shortest geodesics is the negative of this overlap-product argument. Stating which sign is reported prevents an apparent contradiction. This construction requires no arbitrary differentiation of eigenvector phases.

Separate a well-defined angle from an observable signal

At 90 degrees longitude the triangle encloses $\pi/2$ steradians. The complex product is $(1 + i)/4$, its magnitude is about 0.354 and its argument is $\pi/4$. As longitude approaches 180 degrees, the equatorial states approach orthogonality. The magnitude tends to zero, and even a small absolute uncertainty in the measured complex product yields a large angular uncertainty. Exactly at orthogonality there is no unique shortest equatorial arc or measurable phase of the zero complex number.

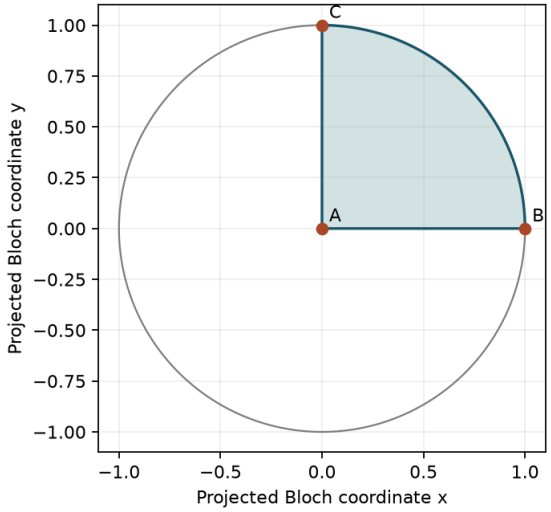

The north pole projects to the centre and the equator to the rim. The filled planar region illustrates the path, whereas its solid angle is integrated on the sphere.

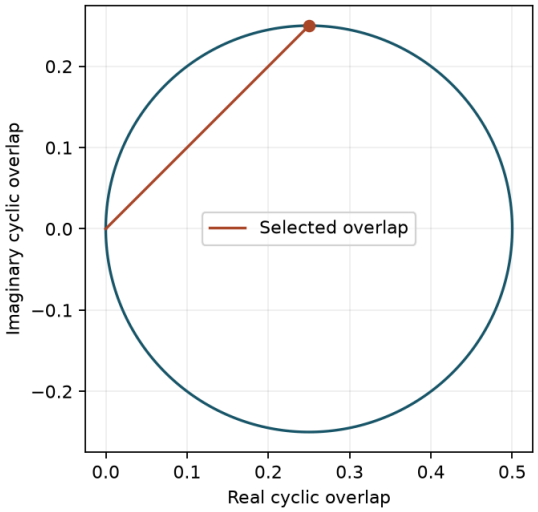

At the origin the overlap phase is undefined. Reversing the circuit reflects the complex product across the real axis.

Exercise · reverse the circuit

What happens to amplitude and phase if the experiment uses A, C, B in place of A, B, C?

▼ **Work through the solution**

Hermitian symmetry conjugates every overlap while the scalar factors can be reordered. The new product equals Z^*. Its magnitude stays fixed and both signed phase conventions reverse.

Exercise · estimate the angular uncertainty

A measured overlap has magnitude 0.02 and a small perpendicular complex uncertainty of magnitude 0.002. What is the first-order phase uncertainty? Compare with a magnitude of 0.4.

▼ **Work through the solution**

For a small transverse displacement the angular change is $\delta\theta \simeq |\delta Z_\perp|/|Z|$. The answers are 0.10 rad and 0.005 rad. Identical absolute detector noise therefore gives very different phase precision near an orthogonal pair.

Try it · Geometric phase

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Third-state longitude (degrees) **90**

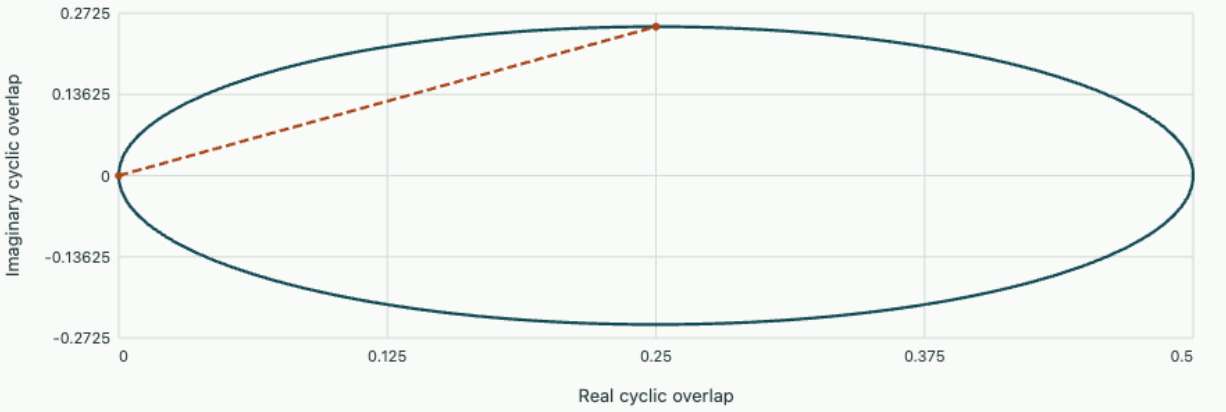

Cyclic-overlap locus Selected overlap

Overlap modulus 0.35355; overlap phase 0.7854 rad. Parallel-transport phase has the opposite sign. Arbitrary individual state phases cancel.

Download the corresponding editable notebook. The browser recalculates the reduced equations directly; the notebook exposes the full calculation and its checks.

Open editable notebook Download .ipynb

Quantum states, dynamics and measurement · 14. From electron exchange to atomic shells and molecular binding

An orbital model can explain why a closed atomic shell is special and why a shared electron can bind two protons. Both predictions need an energy as well as a rule for counting states. We construct the energies and check them against their separated-particle limits.

Count states and find what changes in an atom with many electrons

For a hydrogenic principal number n, allowed orbital quantum numbers run from 0 to $n-1$. Each l permits $2l+1$ magnetic projections. Two spin states accompany each orbital. An antisymmetric many-electron state vanishes if two occupied spin orbitals coincide, giving the shell capacity

$$N_n = 2\sum_{l=0}^{n-1}(2l+1) = 2n^2.$$

This is a counting bound. Hydrogen energies depend only on n in the nonrelativistic Coulomb problem. In a neutral many-electron atom a core screens the nucleus. An orbital that penetrates the core samples a larger net attraction. Self-consistent screening and exchange split subshells. A bare $2n^2$ table cannot predict the order of filling across all atoms or the configurations of all transition metals.

For carbon, a closed $1s^2 2s^2$ core leaves two electrons in $2p$. For distinct spatial states a, b, form normalised spatial combinations $(a(1)b(2) \pm b(1)a(2))/\sqrt{2}$. A symmetric repulsive interaction gives $E_\pm = E_0 + J \pm K$ for orthogonal orbitals, with J the direct integral and K the exchange integral. The antisymmetric spatial branch pairs with a symmetric spin triplet and often lies lower in a degenerate open subshell. This gives a mechanism for Hund's first rule. Large orbital splittings can favour double occupation despite exchange.

> **Recall · Why does an inert shell show a spectroscopic jump?**
>
> Adding an electron above a completed low shell and removing one from that shell involve orbitals with different spatial extent and binding energy. Closed-shell counting also gives zero total orbital and spin angular momentum for the filled subshell. Actual ionisation energies still include relaxation of the remaining electrons.

Construct a molecular-ion energy curve in its own nonorthogonal basis

Fix two protons at separation R and work in atomic units. The nuclei move much more slowly than the electron. The electronic Hamiltonian is $h = -\nabla^2/2 - 1/r_A - 1/r_B$. For each fixed R the total Born–Oppenheimer energy includes the proton repulsion $1/R$. Use two unit-charge hydrogenic $1s$ basis orbitals ϕ_A, ϕ_B. Their overlap matters whenever a bond might form.

$$S = e^{-R}(1 + R + R^2/3),$$
$$J = -1/R + (1 + 1/R)e^{-2R},$$
$$K = -(1 + R)e^{-R}.$$

Here $J = \langle A| -1/r_B|A\rangle$ and $K = \langle A| -1/r_A|B\rangle$. They follow by integrating the two decaying exponentials in the field of the other nucleus. For example, the spherical average of $1/r_B$ on a sphere about A is $1/\max(r_A, R)$. Splitting the J radial integral there gives the expression above. Using the isolated hydrogen equation acting on the ket now gives $H_{AA} = -1/2 + J$ and $H_{AB} = -S/2 + K$.

Vary the Rayleigh quotient with respect to the coefficients. The stationarity equation is $Hc = E_{\mathrm{el}}Sc$ where the matrix on the right has unit diagonal and off-diagonal overlap S. Reflection through the midpoint fixes the even and odd combinations and their normalisations:

$$\psi_\pm = \frac{\phi_A \pm \phi_B}{\sqrt{2(1 \pm S)}},$$
$$U_\pm(R) = \frac{H_{AA} \pm H_{AB}}{1 \pm S} + \frac{1}{R}.$$

As R grows, both curves tend to $-1/2$ hartree, the energy of a neutral H atom and a distant proton. At $R = 2$ bohr the even branch is about -0.554 hartree. Its energy below the dissociation threshold demonstrates binding even within this restricted trial space. The minimum of this curve determines a bond length only within the chosen basis and the fixed-nucleus approximation.

OneSimulation experiment · Two-centre variational energies for the hydrogen molecular ion

Try the molecular density between two nuclei

Move the nuclei and inspect how constructive and destructive interference change the electron density on their connecting line. The odd combination keeps a node at the centre. Compare the even-state energy with the separated-atom value of −0.5 hartree and locate a bound separation.

$$|\psi_\pm(x, 0, 0)|^2 = \frac{\left(e^{-|x+R/2|} \pm e^{-|x-R/2|}\right)^2}{2\pi(1 \pm S(R))}.$$

The orbitals are normalised in three dimensions and the plot is a line through that density. This is the same fixed two-orbital variational approximation as the OneSimulation notebook below. Nuclear motion and orbital relaxation are excluded. Use the notebook's eigenproblem to check the displayed total energies independently.

Open this experiment online

Try the calculation

Nuclear separation (bohr) **2.5**

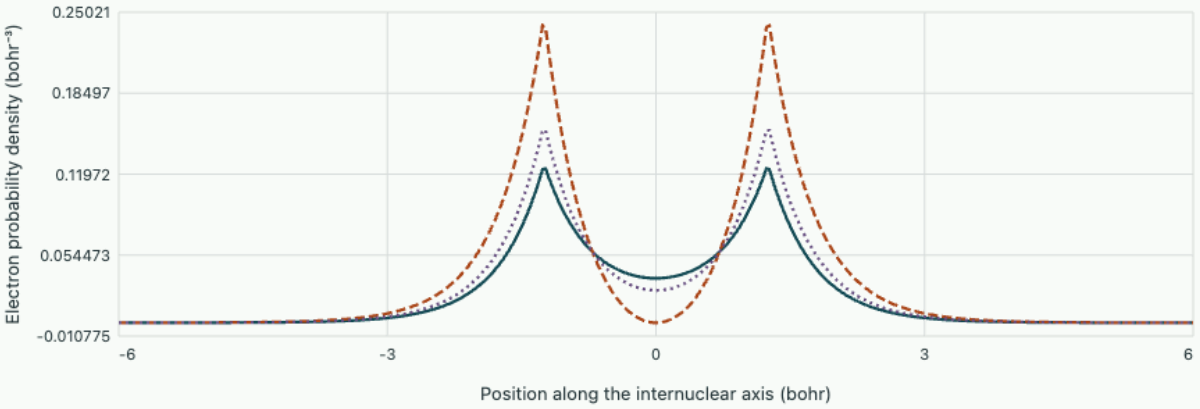

Even combination |ψ+|² Odd combination |ψ−|² Incoherent half-and-half density

Nuclei at x = ±1.25 bohr. Overlap S = 0.458308. Total even/odd energies = -0.564829 / -0.290642 hartree. Even-state binding relative to H + proton = 1.7641 eV. The plotted line samples a three-dimensional density; its one-dimensional integral need not equal one.

Reproduce with OneSimulation

Change the range or sampling of nuclear separations. The solver orthonormalises the overlap metric before diagonalising. Compare its grid minimum with the even-state formula and distinguish grid error from the error inherent in the two-orbital basis.

Run the downloaded notebook with OneSimulation, NumPy and Matplotlib installed. The output below was recorded with the displayed calculation.

```
"""Hydrogen molecular ion in a fixed pair of hydrogenic 1s orbitals; atomic units."""
import numpy as np
import matplotlib.pyplot as plt
from onesim.numerics.linear.matrices.dense_linear_algebra import symmetric_eigenvalues

def total_energies(r):
    s = np.exp(-r) * (1 + r + r*r/3)
    j = -1/r + (1 + 1/r) * np.exp(-2*r)
    k = -(1 + r) * np.exp(-r)
    h = np.array([[-0.5+j, -0.5*s+k], [-0.5*s+k, -0.5+j]])
    metric = np.array([[1.0, s], [s, 1.0]])
    chol = np.linalg.cholesky(metric)
    left = np.linalg.solve(chol, h)
    orthogonal = np.linalg.solve(chol, left.T).T
    return symmetric_eigenvalues(orthogonal) + 1/r

distance = np.linspace(0.9, 8.0, 711)
bonding, antibonding = np.array([total_energies(r) for r in distance]).T
i = np.argmin(bonding)
print(f"Grid minimum R = {distance[i]:.3f} bohr")
print(f"Total energy = {bonding[i]:.6f} hartree")
print(f"Binding relative to H + proton = {(-0.5-bonding[i])*27.2114:.3f} eV")
plt.plot(distance, bonding, label="Even combination")
plt.plot(distance, antibonding, label="Odd combination")
plt.axhline(-0.5, linestyle=":", color="grey", label="Separated atom and ion")
```

```
plt.xlabel("Nuclear separation (bohr)")
plt.ylabel("Total energy (hartree)")
plt.legend()
plt.tight_layout()
plt.savefig("physics-molecular-ion.png", dpi=160)
```

Computed output

```
Grid minimum R = 2.490 bohr
Total energy = -0.564831 hartree
Binding relative to H + proton = 1.764 eV
```

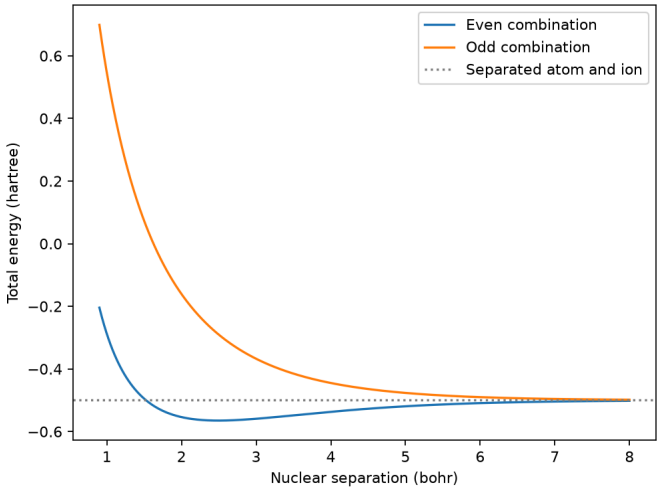

Two-centre variational energies for the hydrogen molecular ion.

Download editable notebook · Download Python script

Add the second electron without assigning an electron to a nucleus

For neutral hydrogen include $1/r_{12}$ in the electronic Hamiltonian. The separated neutral-atom energy is now -1 hartree. With $u = \phi_A(1)\phi_B(2)$ and $v = \phi_B(1)\phi_A(2)$ the two assignments have overlap S^2. Spatial symmetrisation gives the covalent Heitler–London trial states and their energies:

$$U_{\pm}(R) = \frac{\langle u|H_{\mathrm{el}}|u\rangle \pm \langle u|H_{\mathrm{el}}|v\rangle}{1 \pm S^2} + \frac{1}{R}.$$

The plus state must carry the antisymmetric spin singlet. The off-diagonal matrix element includes electronic hopping and two-electron terms. Its sign and the overlap both enter the total energy. At large R, $S \to 0$ and each product contains one atom at each site. In contrast, a doubly occupied bonding molecular orbital expanded in the site basis contains ionic configurations with both electrons on one site. Optimising the relative weights of covalent and ionic configurations improves the dissociation limit.

As a worked matrix test take $R = 2$, $S^2 = 0.25$, $H_{uu} = -1.48$ and $H_{uv} = -0.45$ hartree. These are illustrative matrix elements, not tabulated hydrogen data. The singlet energy is -1.044 hartree after nuclear repulsion. The triplet yields about -0.873 hartree. Only the first trial value lies below the separated-atom threshold. A true equilibrium requires repeating the calculation over R.

Tutorial on shell models and bonds

1 · Count the microstates of two electrons in a p subshell

▼ **Work through the solution**

Six spin orbitals allow $\binom{6}{2} = 15$ antisymmetric two-electron determinants. For two parallel spins choose two distinct spatial orbitals in $\binom{3}{2} = 3$ ways for a fixed spin orientation. This counting alone provides no energy ordering.

---

2 · Diagnose a false bond from an omitted term

▼ **Work through the solution**

If a calculation reports only E_{el}, decreasing R can lower it through attraction to both nuclei. The force on the separation comes from $-d(E_{\mathrm{el}} + 1/R)/dR$. Omitting repulsion gives a false energy minimum or a collapse at short distance.

---

3 · Estimate the vibrational frequency near a computed minimum

▼ **Work through the solution**

Suppose $U''(R_e) = 0.10$ hartree/bohr² and the nuclear reduced mass is 918 electron masses. The angular frequency in atomic units is $\sqrt{0.10/918} = 0.0104$. In this harmonic model the vibrational energy spacing is 0.0104 hartree, or 0.284 eV. Replacing both protons by twice their mass divides the spacing by $\sqrt{2}$.

For further study see Griffiths and Schroeter, Introduction to Quantum Mechanics, 3rd edition, sections 5.2 and 8.2–8.4 on atomic configurations and variational molecular energies.

Quantum states, dynamics and measurement · 15. The scattering integral equation and successive Born terms

A weak potential often produces an angular pattern resembling its Fourier transform. What controls that approximation? A particle can propagate between two interactions with the same potential. Accounting for that propagation builds an integral equation and supplies a calculable second term.

Construct a Green function with the outgoing boundary condition

For a short-range potential and incident energy $E = \hbar^2 k^2/(2m) > 0$, move the potential to the right side of the stationary Schrödinger equation. The outgoing Helmholtz kernel satisfies $(\nabla^2 + k^2)e^{ikR}/R = -4\pi\delta(\mathbf{R})$. Its outward phase is appropriate for incident states normalised to a plane wave $e^{i\mathbf{k}\cdot\mathbf{r}}$ and time dependence $e^{-iEt/\hbar}$.

$$\psi(\mathbf{r}) = e^{i\mathbf{k}\cdot\mathbf{r}} - \frac{m}{2\pi\hbar^2}\int d^3r'\,\frac{e^{ik|\mathbf{r}-\mathbf{r}'|}}{|\mathbf{r}-\mathbf{r}'|}V(\mathbf{r}')\psi(\mathbf{r}').$$

Apply the differential operator to check the source and sign. Far outside the interaction region expand the distance in the phase as $r - \hat{\mathbf{r}}\cdot\mathbf{r}'$ and in the denominator as r. For elastic scattering $\mathbf{k}' = k\hat{\mathbf{r}}$, the coefficient of e^{ikr}/r is the amplitude.

$$f(\mathbf{k}',\mathbf{k}) = -\frac{m}{2\pi\hbar^2}\int d^3r'\,e^{-i\mathbf{k}'\cdot\mathbf{r}'}V(\mathbf{r}')\psi(\mathbf{r}').$$

Interpret each application of the free propagator

In operator notation the integral equation reads $|\psi\rangle = |\mathbf{k}\rangle + G_0^{(+)}V|\psi\rangle$. Substitution of its right side for $|\psi\rangle$ gives a plane wave, one intermediate propagation, then two propagations and so forth. The corresponding transition operator obeys

$$T = V + VG_0^{(+)}T = V + VG_0^{(+)}V + VG_0^{(+)}VG_0^{(+)}V + \cdots.$$

The first term measures the transform of V at the measured momentum transfer. In the second term an intermediate momentum is integrated over with denominator $E - \hbar^2p^2/(2m) + i0$ for p a wave number. It can differ from the incident momentum. Energy is conserved between the asymptotic states. The infinitesimal imaginary part retains the outgoing-wave condition in the integral.

> **Recall · A practical failure of first order**
>
> For a real potential the first Born forward amplitude is real. The exact optical theorem reads $\sigma_{\text{tot}} = (4\pi/k)\,\mathrm{Im}\,f(0)$. The first nonzero total cross section is order V^2. Its consistent forward contribution comes from second order, so combining orders indiscriminately can produce an apparent violation of flux conservation.

Sum a solvable sequence and test a numerical truncation

Consider a separable interaction $V = \lambda|g\rangle\langle g|$ with a square-integrable form factor. Each repeated propagation contributes the scalar $I(E) = \langle g|G_0^{(+)}(E)|g\rangle$. Trying $T = |g\rangle\tau\langle g|$ gives $\tau = \lambda + \lambda I\tau$ and hence

$$\tau = \frac{\lambda}{1 - \lambda I},$$

$$\tau_N = \lambda \sum_{j=0}^{N-1} (\lambda I)^j.$$

For $q = \lambda I = 0.20 - 0.10i$ the first-order relative amplitude error is $|q| = 0.224$. Truncating after two and three terms reduces it to $|q|^2 = 0.050$ and $|q|^3 = 0.0112$. A value approaching $q = 1$ signals strong enhancement. Exact poles follow by continuation to the correct energy domain. A small cross section in just one detector can arise from a form-factor zero even when the repeated-scattering factor is large.

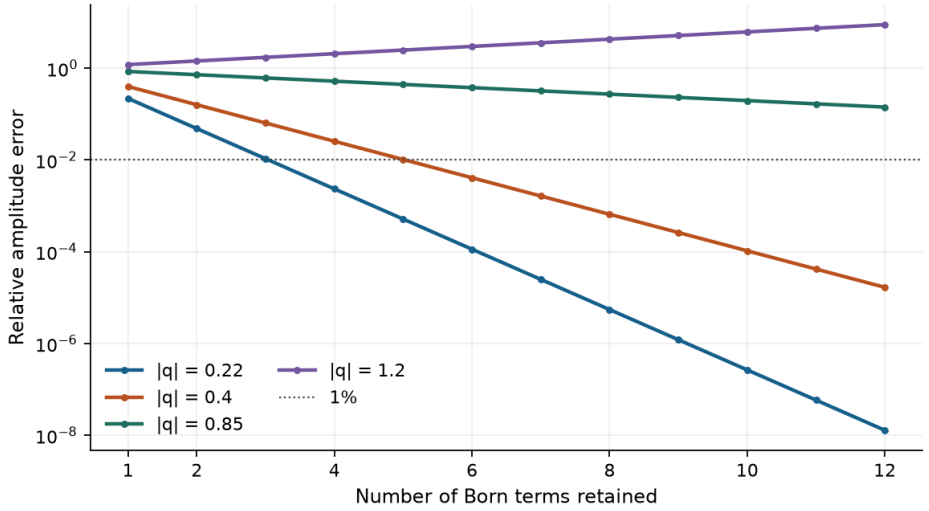

Exact relative errors for the separable model. A small change in the magnitude of the repeated-scattering factor can greatly change how many Born terms are useful.

Try it · Born

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Repeated-scattering magnitude |λI| **0.7**

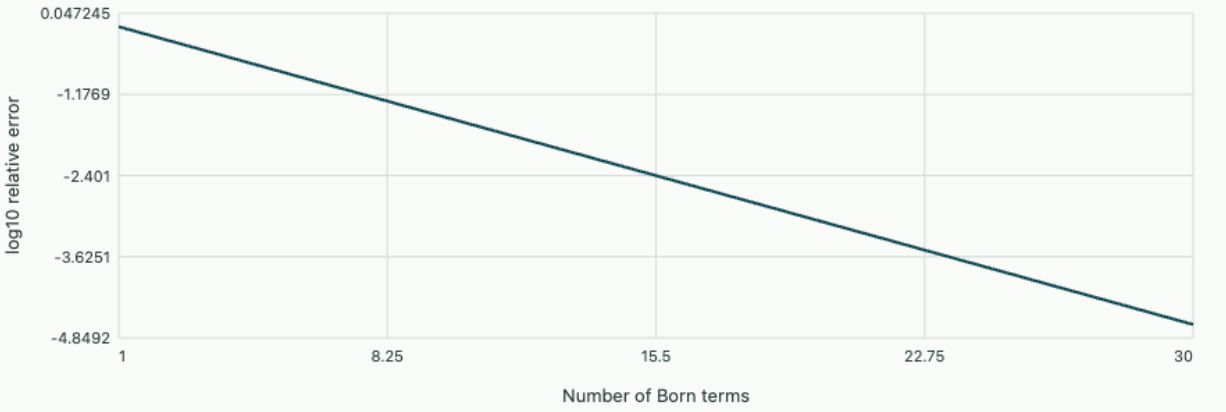

Exact separable-model truncation error

Relative error is |λI|^N. The geometric series converges.

Tutorial on recognising repeated scattering

1 · Check a Born-amplitude dimension

▼ Work through the solution

$m/\hbar^2$ has units of inverse energy times inverse length squared. Multiplying by an energy and a three-dimensional volume produces length. Squaring the amplitude gives the area per unit solid angle.

2 · Choose the number of terms

▼ Work through the solution

For a scalar $|q| = 0.40$, the relative truncation error is $|q|^N$. Five terms give 0.01024, slightly above 1%. Six give 0.00410. In a general potential no single scalar controls every channel, so repeat convergence tests with multiple incident energies and angles.

---

3 · Distinguish invertibility from convergence

▼ **Work through the solution**

At $q = -1.2$ the exact expression $\lambda/(1 - q)$ is finite. Individual Born terms grow and alternate. Failure of the expansion need not signal an infinite physical scattering amplitude. Solving the full integral equation and checking its residual are different tests from summing the series.

For further study see Griffiths and Schroeter, Introduction to Quantum Mechanics, 3rd edition, sections 10.3–10.4 on phase shifts, the integral equation and the Born series.

Quantum states, dynamics and measurement · 16. Spontaneous radiation, stimulated transitions and thermal detailed balance

A coherent laser can drive Rabi oscillations in two atomic levels. A large cavity filled with thermal radiation presents many modes with random phases. In this limit one can infer the ratios of absorption, stimulated emission and spontaneous emission coefficients by requiring the observed equilibrium to be independent of the atoms used to probe it.

Specify the convention for an Einstein coefficient

Let $E_2 - E_1 = h\nu$ and let g_i count the unresolved magnetic substates of level i. Define ρ_ν as energy per unit volume per hertz including both polarisations, and use coefficients averaged over equally populated lower or upper substates as appropriate. For an isolated narrow line in isotropic incoherent radiation the upward rate per atom is $B_{12}\rho_\nu$. The downward rate is $A_{21} + B_{21}\rho_\nu$. For a broad line average ρ_ν against a normalised line shape.

At equilibrium $N_2/N_1 = (g_2/g_1)\exp(-h\nu/k_BT)$ while stationarity requires $N_1 B_{12}\rho_\nu = N_2(A_{21} + B_{21}\rho_\nu)$. Isolate the unknown radiation density:

$$\rho_\nu = \frac{A_{21}/B_{21}}{(g_1 B_{12}/g_2 B_{21})e^{h\nu/k_BT} - 1}.$$

Count the cavity modes and identify the coefficients

In a large periodic volume V the count of wave vectors in a shell is $V\,4\pi k^2 dk/(2\pi)^3$. Two transverse polarisations double it. With $k = 2\pi\nu/c$ this gives $8\pi\nu^2/c^3$ modes per volume per hertz. Each thermally occupied mode has mean energy $h\nu/(e^{h\nu/k_BT} - 1)$ above vacuum. Equality at every temperature fixes both independent factors in the rate expression.

$$g_1 B_{12} = g_2 B_{21},$$

$$\frac{A_{21}}{B_{21}} = \frac{8\pi h\nu^3}{c^3}.$$

The spontaneous-to-stimulated ratio in a blackbody is $e^{h\nu/k_BT} - 1$. Knowing the line frequency fixes a ratio of coefficients, not the dipole matrix element or the absolute lifetime. An isolated upper level with several radiative channels has total decay rate $\sum_j A_{2j}$ and emits into a specific channel with probability $A_{2j}/\sum_j A_{2j}$.

> **Recall · Frequency and angular frequency conventions**
>
> Spectral energies obey $\rho_\nu\,d\nu = \rho_\omega\,d\omega$. With $\omega = 2\pi\nu$, $\rho_\omega = \rho_\nu/(2\pi)$. A coefficient multiplying ρ_ω therefore changes by 2π. Comparing a printed A/B ratio requires checking this convention and whether intensity or energy density was used.

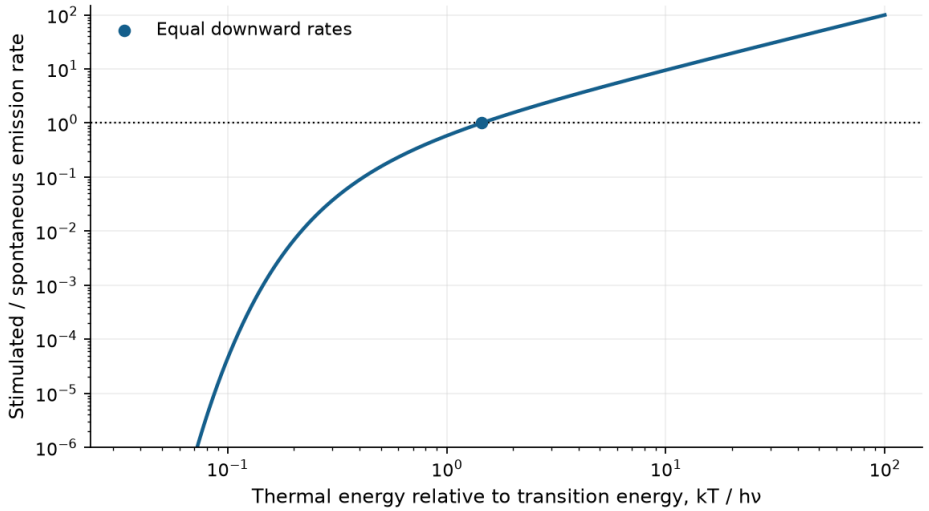

The relative importance of stimulation can be read from the ratio of bath energy to transition energy. The crossing fixes the bath temperature where the two downward rates agree.

Try it · Thermal emission

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Transition frequency (THz) **20**

Bath temperature (K) **300**

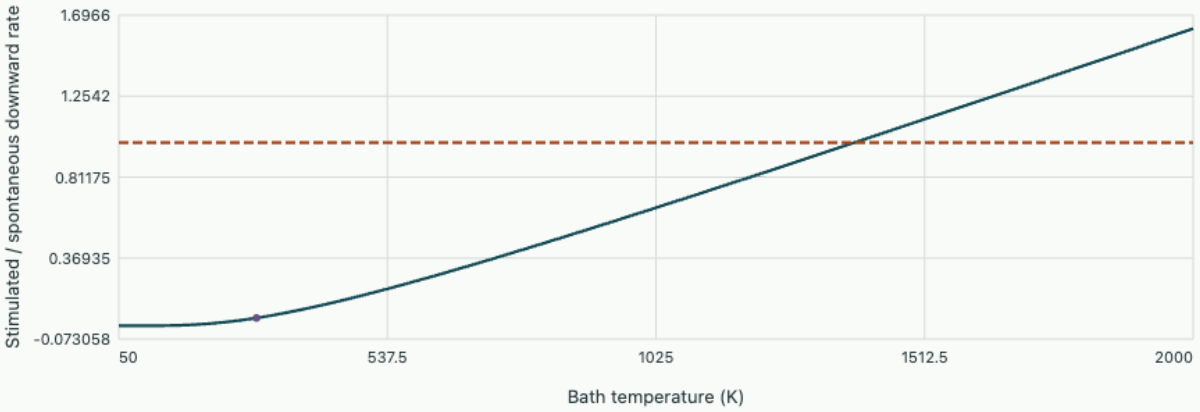

Thermal occupation · Equal downward rates · Selected bath

At selected bath stimulated/spontaneous = 0.042517. Equality occurs at 1384.8 K; the transition energy is hν.

Predict a lifetime and an equilibrium occupation

Consider a 10 GHz transition. Its characteristic temperature $h\nu/k_B$ is about 0.480 K. At 4 K the mean photon occupation is $1/(e^{0.120}-1) = 7.84$, so stimulated downward transitions dominate spontaneous ones. At 0.1 K the occupation is only 0.00830. A 600 THz optical transition at room temperature has $h\nu/k_BT \simeq 96$, leaving essentially no thermal photons at that frequency.

If an upper state has two decay channels of 4×10^7 and 1×10^7 s⁻¹, its vacuum population after a short excitation decays as $e^{-5\times10^7 t}$. The lifetime is 20 ns. The photon yield in the stronger branch is 0.80 in the absence of quenching and losses. Rates inferred from measured photon counts also need a detector-efficiency correction.

Tutorial on reading a line measurement

1 · Correct for unresolved degeneracies

▼ **Work through the solution**

If $g_1 = 2$, $g_2 = 6$, then $B_{12} = 3B_{21}$ in the convention of this chapter. Population inversion for net stimulated gain requires $N_2B_{21} > N_1B_{12}$, or $N_2/g_2 > N_1/g_1$.

---

2 · Find the bath temperature for equal downward contributions

▼ **Work through the solution**

Equality requires mean occupation one, hence $e^{h\nu/k_BT} = 2$. For 10 GHz, $T = 0.480/\ln 2 = 0.693$ K. Upward absorption may still occur at this temperature.

---

3 · Infer a missing nonradiative loss

▼ **Work through the solution**

An isolated excited population lasts 16 ns instead of 20 ns while calibrated radiative rates still sum to 5.0×10^7 s⁻¹. An independent nonradiative channel adds $1/(16\,\mathrm{ns}) - 5.0 \times 10^7 = 1.25 \times 10^7$ s⁻¹. The radiative quantum yield is 0.80.

For further study see Griffiths and Schroeter, Introduction to Quantum Mechanics, 3rd edition, section 11.3 on spontaneous emission and atomic lifetimes.

Quantum states, dynamics and measurement · 17. Which properties of an unknown spin can an instrument copy?

A device calibrated with spin up and down can write two perfectly distinguishable records. Asking what happens to a superposition reveals a limit on information extraction and explains why state tomography needs an ensemble.

Apply linearity to a calibrated copying apparatus

Let a blank qubit be prepared in $|0\rangle$. A unitary device might obey $U|0,0\rangle = |0,0\rangle$ and $U|1,0\rangle = |1,1\rangle$. Both rules are realised by a controlled NOT on the blank. Linearity determines the device on any superposition of these inputs. For $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ the output is

$$U|+,0\rangle = (|0,0\rangle + |1,1\rangle)/\sqrt{2}.$$

The desired two-copy state would be $|+,+\rangle = (|0,0\rangle + |0,1\rangle + |1,0\rangle + |1,1\rangle)/2$. Measure both qubits along z. The actual output has zero probability of unequal results. The desired output has probability one half of unequal results. These are experimentally distinguishable preparations.

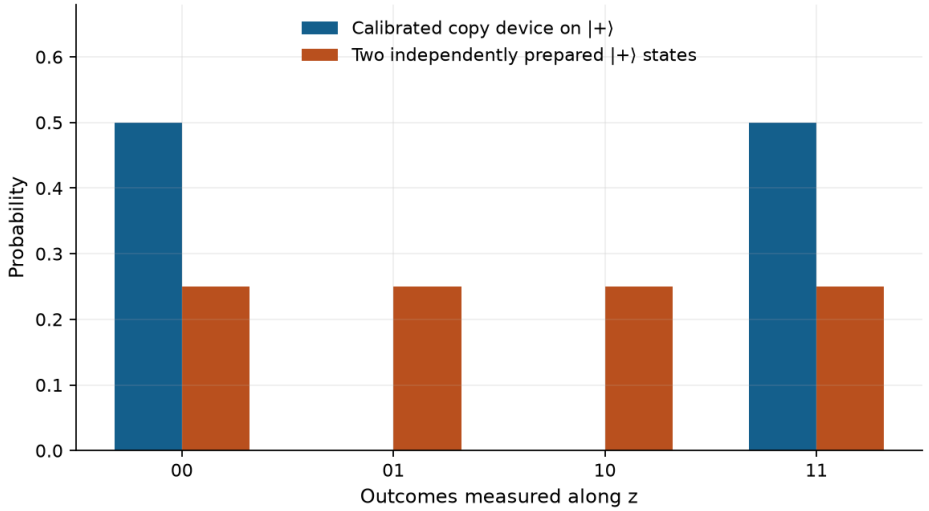

The calibrated copying device gives only matching spin records along the calibration axis. Two independently prepared superpositions also give the two unequal pairs.

Try it · Spin copy

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Input polar angle (degrees) 90

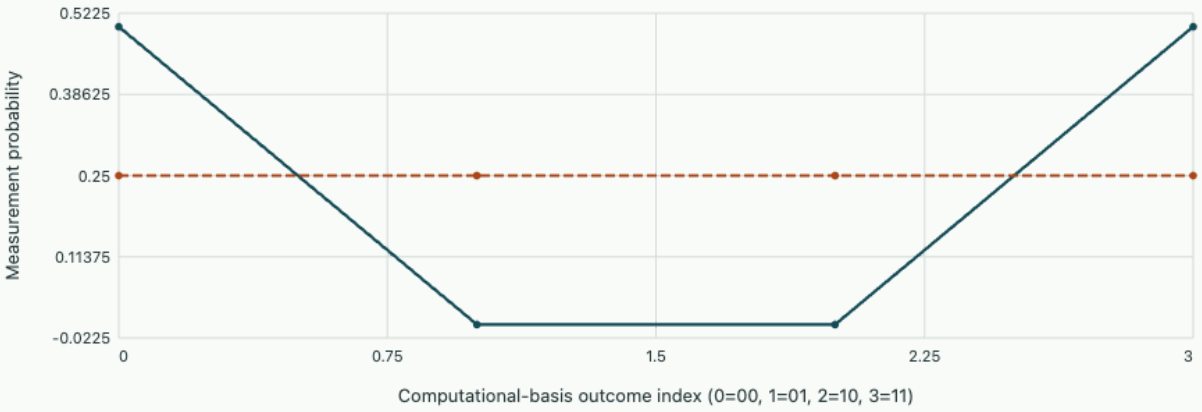

Calibrated basis copier Two independently prepared copies

The copier creates correlated records. At a superposition input they differ from two independent copies; at either calibrated basis state they agree.

Derive the obstruction without assuming a particular basis

Suppose an apparatus with fixed blank and fixed prepared environment clones each of two pure states $|\psi\rangle$ and $|\phi\rangle$. After the operation any residual machine states may depend on the input. Because the two output copies are pure, they factor from the machine in each branch. Taking the magnitude of the inner product before and after unitary evolution gives

$$s = s^2\,|\langle M_\psi | M_\phi \rangle|,$$
$$s = |\langle \psi | \phi \rangle|.$$

The machine overlap is at most one, so inputs with $0 < s < 1$ cannot satisfy this equality. Orthogonal states form the special class for which exact, deterministic copying is possible. For $|0\rangle$ and $|+\rangle$, $s = 1/\sqrt{2}$ but two copies have overlap $1/2$ even before multiplying by the machine overlap. Additional hidden apparatus states cannot fix this discrepancy.

Recall · Pure, correlated and mixed outputs

Partial trace of the actual two-qubit entangled output gives each single qubit $\rho = I/2$. That result matches two fair z histograms, yet the pair has perfect z correlation. Single-axis marginals never certify that a pair consists of independent copies.

Relate disturbance to an observable record

A measurement trying to leave each state intact while recording its identity would have outputs $|\psi\rangle|m_\psi\rangle$ and $|\phi\rangle|m_\phi\rangle$. Inner-product preservation gives $\langle\psi|\phi\rangle = \langle\psi|\phi\rangle\langle m_\psi|m_\phi\rangle$. If the inputs overlap, a nondisturbing measurement of both must leave the same record up to phase. This argument assumes exact preservation. Weak or approximate measurements can provide partial information at the cost of a disturbance.

With independently prepared copies, spin tomography estimates different components on different trials. If N separate x measurements have expectation r_x, then $\hat{r}_x = (n_+ - n_-)/N$ has variance $(1 - r_x^2)/N$. At $r_x = 0$ one hundred trials give standard error 0.10. Copying a single unknown preparation could not produce these independent statistics.

Tutorial on spin records

1 · Predict a transverse measurement

▼ **Work through the solution**

For the Bell state above each individual x result is fair, but the product of the two x results is always $+1$. For the desired $|+,+\rangle$ both individual values are always $+1$. This separates the two preparations with repeated measurements.

---

2 · Distinguish a classical record from an unknown state

▼ **Work through the solution**

A classical bit is encoded in two distinguished orthogonal states. Controlled NOT copies that alphabet. Rotating the incoming state to an unknown basis introduces nonorthogonal alternatives relative to that alphabet, for which the same guarantee fails.

---

3 · How many new preparations are needed?

▼ **Work through the solution**

To estimate a spin component near zero to one standard error of 0.02 requires $N \simeq 1/0.02^2 = 2500$ independent preparations for that axis. Splitting the same budget among three axes increases each

component's error. Repeating a projective measurement on one qubit measures the prepared post-measurement state.

For further study see Griffiths and Schroeter, Introduction to Quantum Mechanics, 3rd edition, section 12.4 on the no-clone theorem.

Computational quantum experiments

Follow a quantum state from its preparation to observable frequencies, correlations and counting statistics. Test how basis size, finite sampling and detector settings change each inference.

1. ## Measure tunnelling from a four-site spectrum

   Four localised orbitals form a minimal double well. Raise the inner two site energies and ask why two states concentrated in the outer wells still acquire different energies. Measure the splitting in the spectrum, predict an oscillation in time and reproduce the levels independently with OneSimulation.

2. ## Which part of a Bell correlation survives imperfect state preparation?

   A source emits two spin-half systems in the singlet state. Two observers select measurement directions separately and record outcomes $+1$ or -1. Construct a Bell–CHSH statistic from four measurable correlations and determine how state preparation, angular alignment and finite sample size affect the prediction.

3. ## Recover a transition frequency from two separated pulses

   A two-level system is driven by two short pulses with a free evolution interval between them. Determine why a population measurement carries phase information. Use a scan at several delays to separate a constant detuning, intrinsic coherence loss and trial-to-trial frequency variation.

4. ## When three perfectly legal spin measurements imply an impossible state

   Prepare nominally identical spin-half systems and measure along three orthogonal axes, using a different subensemble for each axis. Reconstruct a density matrix, diagnose an unphysical linear estimate and see why the choice of estimator matters in finite data.

5. ## Read wave-packet overlap from a coincidence dip

   Send one photon into each of two input ports of a balanced lossless beam splitter. Two alternatives lead to one click at each output. Derive the coincidence probability when their wave packets overlap only partially, then determine how a timing scan separates pulse duration from distinguishability in another degree of freedom.

6. ## Can an atom leave its initial state and still pass the final test?

   A driven two-level atom undergoes repeated ideal measurements. Compare two laboratory records, one retaining only runs that passed every test and another retaining every run. The distinction is visible even when every measurement is made with the same apparatus.

Computational quantum experiments · 1. Measure tunnelling from a four-site spectrum

Four localised orbitals form a minimal double well. Raise the inner two site energies and ask why two states concentrated in the outer wells still acquire different energies. Measure the splitting in the spectrum, predict an oscillation in time and reproduce the levels independently with OneSimulation.

Construct an operator before drawing an energy diagram

Label the sites $1, 2, 3, 4$ in order. The end orbitals have on-site energy zero, the middle two have energy $V > 0$ and every neighbouring pair has real hopping matrix element $-t$ with $t > 0$. The four occupation states form an orthonormal basis for one particle. Equal diagonal energies of the ends do not make their symmetric and antisymmetric combinations degenerate when there is a virtual path through the interior. There are no direct links between sites one and three or one and four. Keeping track of every allowed step will fix the power of t in the final answer.

$$H = \begin{pmatrix} 0 & -t & 0 & 0 \\ -t & V & -t & 0 \\ 0 & -t & V & -t \\ 0 & 0 & -t & 0 \end{pmatrix}.$$

Reflection maps $1 \leftrightarrow 4$ and $2 \leftrightarrow 3$, and commutes with H. In each parity sector use the normalised outer and inner basis vectors $(|1\rangle \pm |4\rangle)/\sqrt{2}$ and $(|2\rangle \pm |3\rangle)/\sqrt{2}$. The diagonal outer element vanishes in both sectors; the hopping inside the middle pair shifts its even and odd energies to $D_e = V - t$ and $D_o = V + t$. The outer–inner matrix element remains $-t$. This reduces the four-dimensional problem to two symmetric two-by-two matrices without neglecting any state.

Eliminate the barrier amplitudes and estimate the splitting

In a sector with interior energy D, the secular equation is

$$E(E - D) = t^2.$$

The lower root is $E_-(D) = [D - \sqrt{D^2 + 4t^2}]/2$, or the numerically better $-2t^2/(D + \sqrt{D^2 + 4t^2})$. For $D \gg t$ the shift is approximately $-t^2/D$. It is more negative for the even state because $D_e < D_o$. Subtract the two shifts before approximating their difference.

$$\Delta E = E_-(D_o) - E_-(D_e) \simeq t^2 \left(\frac{1}{V - t} - \frac{1}{V + t} \right) \simeq \frac{2t^3}{V^2}.$$

A single outward excursion produces an order-t^2/V shift in both end states. Only a three-hop path connects the different ends, so their splitting is smaller. Measuring an overall displacement of two spectral lines and measuring the distance between them test different physical mechanisms.

Predict a time trace and define what is measured

Denote the two low-energy eigenstates by $|e\rangle, |o\rangle$. Prepare $(|e\rangle + |o\rangle)/\sqrt{2}$. In the corresponding opposite superposition $(|e\rangle - |o\rangle)/\sqrt{2}$, the occupation at time τ equals $\sin^2(\Delta E\,\tau/2\hbar)$. Its first maximum is at $\pi\hbar/\Delta E$. These superpositions include interior weight, especially at low barrier. A detector on site four alone measures a different projector. Preparing bare site one excites both high and low doublets and can generate additional fast oscillations.

For $V/t = 4$, the two lower eigenvalues in units of t are $(3 - \sqrt{13})/2$ and $(5 - \sqrt{29})/2$, about –0.30278 and –0.19258. Their separation is about $0.11019t$. The high-barrier approximation gives $0.125t$. With $t = 1\,\mathrm{meV}$, $\hbar = 0.65821\,\mathrm{meV\,ps}$ and the first complete transfer in the dressed basis takes about 18.8 ps. If a measurement integrates over many such oscillations, it may return one half despite perfectly coherent evolution.

Compare spectra from three independent constructions

Use the parameter controls to vary the middle barrier and hopping. The plot shows the exact splitting and its leading virtual-hopping prediction over a range of V/t. The printed result gives an actual transfer time. The downloadable OneSimulation model builds the original four-site Hamiltonian and solves its eigenvalue problem by contour integration. Compare all four eigenvalues with the two parity blocks before interpreting their small difference. A wrong sign on one edge of this open chain can be removed by a local change of basis and may leave its spectrum unchanged, so also check any phase-sensitive observable used for preparation.

Keep the hopping fixed at 1 meV and compare two barriers. Start with a 4 meV barrier and raise it to 8 meV. Record both splittings and estimate the change in the first transfer time before reading off the prediction.

Open this experiment online

Try the calculation

Middle site energy V (meV) 4

Hopping t (meV) 1

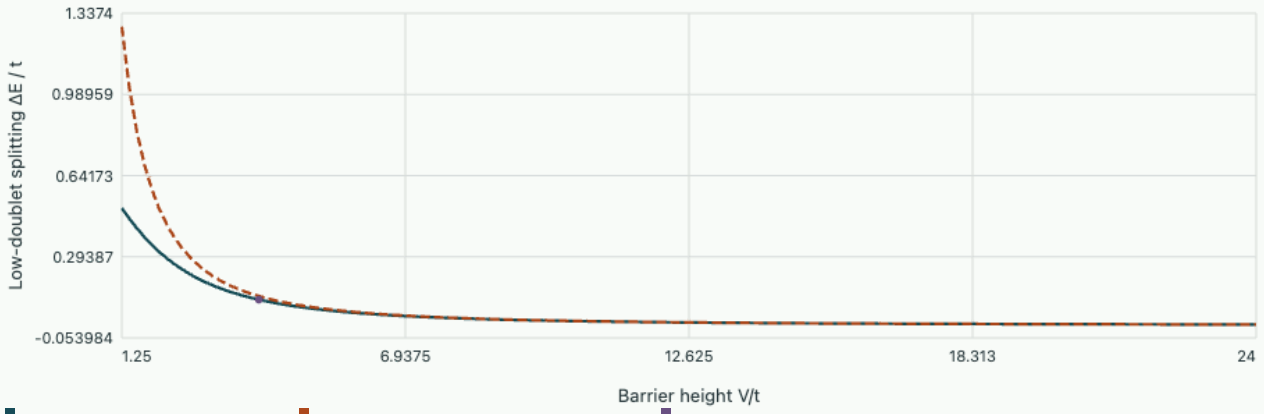

Exact parity sectors | Three-hop approximation | Selected barrier

Splitting 0.11019 meV. First transfer between dressed end states at 18.766 ps.

Download the OneSimulation model

In your OneSimulation environment, run the model with these inputs.

```
python spectra-and-orbits.py wells --barrier 4 --hopping 1
```

Tutorial

Hold V fixed and halve t in the high-barrier limit. How do the common shift, splitting and transfer time change?

▼ **Work through the solution**

The leading common shift scales as t^2/V and falls by four. The splitting scales as t^3/V^2 and falls by eight. The transfer time grows eightfold. Halving both V and t instead keeps their ratio fixed and halves every eigenvalue.

Explain why a small static asymmetry between the end wells can suppress almost complete transfer while increasing their measured frequency splitting.

▼ **Work through the solution**

In a two-state reduction, detuning δ and coupling J give separation $\sqrt{\delta^2 + 4J^2}$. Starting in one localised state, the maximum transfer probability is $4J^2/(\delta^2 + 4J^2)$. Both the frequency and the contrast are needed to infer the coupling.

How could one check that a contour solver has found both members of a closely spaced doublet rather than printing one twice?

▼ Work through the solution

Check the rank of the independent projected vectors and the residual norm for each vector. Then exploit reflection parity or solve in each symmetry sector independently. Repeating a low residual for two nearly parallel vectors does not establish a two-dimensional eigenspace.

The continuous barrier calculation provides a different model in which barrier width eventually produces an exponential dependence.

Computational quantum experiments · 2. Which part of a Bell correlation survives imperfect state preparation?

A source emits two spin-half systems in the singlet state. Two observers select measurement directions separately and record outcomes $+1$ or -1. Construct a Bell–CHSH statistic from four measurable correlations and determine how state preparation, angular alignment and finite sample size affect the prediction.

Relate a counted coincidence to an operator

Denote unit measurement directions by $\mathbf{a}, \mathbf{b}$. Outcome projectors are $(1 + s\,\boldsymbol{\sigma}\cdot\mathbf{a})/2$ and $(1 + t\,\boldsymbol{\sigma}\cdot\mathbf{b})/2$, for $s, t = \pm 1$. The singlet has zero individual spin expectation and tensor $\langle \sigma_i \otimes \sigma_j \rangle = -\delta_{ij}$. The diagonal entries follow because measurements along any common axis are exactly opposite and the state is rotationally invariant. Its joint probability is

$$p(s,t|\mathbf{a},\mathbf{b}) = \tfrac{1}{4}(1 - st\,\mathbf{a}\cdot\mathbf{b}),$$

$$E(\mathbf{a},\mathbf{b}) = -\mathbf{a}\cdot\mathbf{b}.$$

Summing over either outcome leaves one half for the other observer, independent of the remote angle. Covariance can change while the individual counting rates remain constant. At an angle of 60 degrees the probabilities of the two equal-outcome combinations sum to one quarter, and of opposite outcomes to three quarters. Their signed difference gives –0.5. The lab models spin-axis angles; linear polarisation analysers give a doubled angle in their effective two-state representation.

Derive a bound and choose a set of directions

If a trial had four predetermined local binary responses A, A', B, B', its combination $AB + AB' + A'B - A'B' = A(B + B') + A'(B - B')$ would have magnitude two, since exactly one of the parentheses vanishes. Averaging over a common distribution of such assignments preserves the bound. The quantum prediction combines expectation values measured on different subensembles with independently chosen setting pairs. Use coplanar angles $a = 0,\ a' = \pi/2,\ b = \theta,\ b' = -\theta$.

$$S = E(a,b) + E(a,b') + E(a',b) - E(a',b') = -2(\cos\theta + \sin\theta).$$

At $\theta = 45^\circ$ its magnitude is $2\sqrt{2}$. The derivative vanishes there, making small symmetric angular errors second order. If a fraction $1 - v$ of pairs is completely unpolarised, the density operator is $v|\Psi^-\rangle\langle\Psi^-| + (1 - v)I/4$. Every correlator gains a factor v. The best attainable magnitude exceeds two only when $v > 1/\sqrt{2}$. This is a statement about this measurement family and this noise model. A pair source and a scalar contrast alone do not define an experimental protocol.

Assess one set of finite counts

At $v = 0.80$, $\theta = 45^\circ$ the expected magnitude is about 2.263. Each trial product is a binary variable with variance $1 - E^2$. For N independent pairs per setting, the sample correlation variance is $(1 - E^2)/N$. Assuming independent subensembles, variance of the signed sum is the sum of the four variances. Here

every squared correlation is 0.32. With 1000 trials per setting the approximate standard error of the statistic is

$$\sqrt{4(1-0.32)/1000} = 0.0522.$$

The excess 0.263 is roughly five standard errors under this idealised independent-trial model. It is not a substitute for a finite-sample hypothesis test that accounts for an actual protocol's setting choices, losses, timing and possible memory.

Change settings and preparation independently

Inspect the predicted absolute statistic and the local assignment bound. Find the angles for which the chosen source can exceed two. Then vary visibility at a fixed angle. Small individual spin polarisation drifts could alter the marginal counts; unequal channel efficiencies could change a coincidence-selected sample. Monitoring all four joint count tables and both single-observer marginal distributions allows more checks than retaining only their final sum.

Open this experiment online

Try the calculation

Preparation visibility **0.8**

Symmetric analyser angle (degrees) **45**

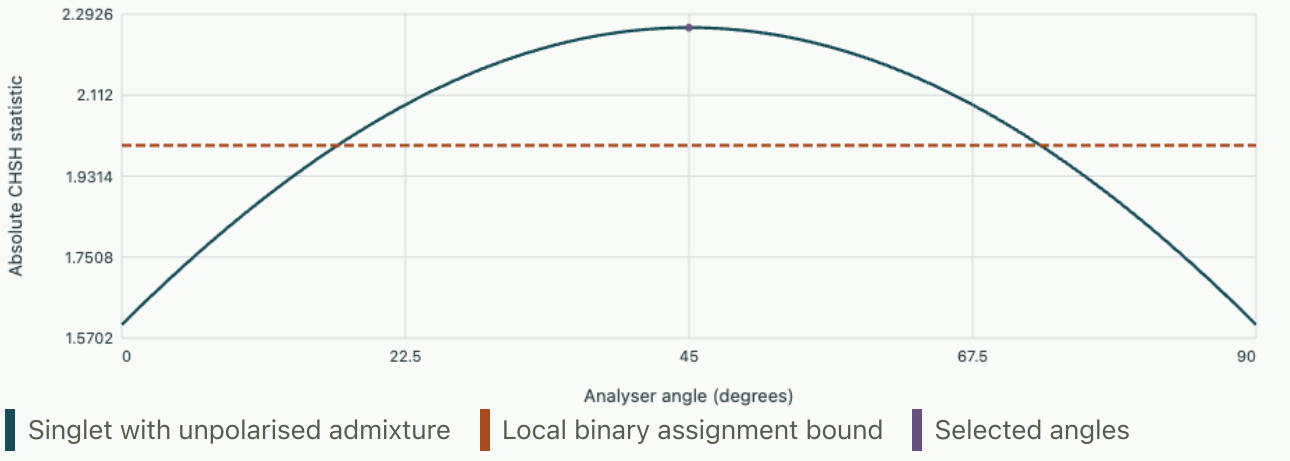

Predicted |S| = 2.2627. With these angles the visibility required to reach two is 0.70711.

Tutorial

At visibility 0.70, can increasing the number of trials eventually make the expected maximum exceed two?

▼ **Work through the solution**

No. The largest mean value is $2\sqrt{2}(0.70) = 1.980$. More trials narrow the sampling distribution around that value. Changing the preparation or measurement family addresses a different question.

What is the predicted fractional loss at an angle δ away from the optimum, and how large is it at five degrees?

▼ **Work through the solution**

Write $\cos\theta + \sin\theta = \sqrt{2}\cos(\theta - \pi/4)$. The fractional loss is $1 - \cos\delta$, about 0.00381 at five degrees. Asymmetric offsets on different analyser axes need not enter the same way.

A paper gives one large value of $|S|$ but discards all pairs for which one detector was silent. Can the ideal bound derived above automatically be applied to these retained counts?

▼ **Work through the solution**

The derivation averaged all setting pairs over a common underlying ensemble. Setting- or state-dependent retention can alter those ensembles. One needs the trial definition, data handling and efficiency assumptions to interpret the reported statistic.

For state preparation and independent measurement on tensor factors, revisit the composite-state chapter.

Computational quantum experiments · 3. Recover a transition frequency from two separated pulses

A two-level system is driven by two short pulses with a free evolution interval between them. Determine why a population measurement carries phase information. Use a scan at several delays to separate a constant detuning, intrinsic coherence loss and trial-to-trial frequency variation.

Follow the amplitudes through both pulses

Use basis $|g\rangle, |e\rangle$. Let the first pulse prepare $(|g\rangle + |e\rangle)/\sqrt{2}$ in the frame of the applied oscillator. A detuning Δ in radians per second accumulates a relative phase $\Delta\tau$ in free time τ. Immediately before the second pulse each state has population one half; measuring population here cannot reveal detuning. Choose a second pulse equivalent to projecting onto $(|g\rangle + e^{-i\phi}|e\rangle)/\sqrt{2}$. Then one output population is

$$P(\tau, \phi) = \tfrac{1}{2}[1 + \cos(\Delta\tau - \phi)].$$

The sign of the phase follows our definitions of detuning and the reference pulse. Shifting the reference phase makes it possible to check that sign. A measurement using only $\phi = 0$ cannot distinguish Δ from $-\Delta$. Coherence is the off-diagonal density-matrix element that the second pulse converts into an observable population difference.

Derive two envelopes with different origins

A homogeneous decay time T_2 multiplies the off-diagonal element by $e^{-\tau/T_2}$. If instead the detuning of each experimental shot is constant during that shot and drawn from a Gaussian distribution of mean $\bar{\Delta}$ and standard deviation σ_Δ, average the complex phase, not its absolute value. Completing the square in the Gaussian integral gives

$$\langle e^{i\Delta\tau} \rangle = e^{i\bar{\Delta}\tau} e^{-\sigma_\Delta^2 \tau^2/2}.$$

With both independent mechanisms the measured contrast is $C(\tau) = \exp[-\tau/T_2 - \sigma_\Delta^2\tau^2/2]$. One mechanism gives an initial logarithmic slope, the other adds curvature. A single low contrast measured at just one delay cannot establish which one caused it.

Infer parameters from quadratures and a frequency scan

Record the $\phi = 0, \pi/2$ outputs, then form $X = 2P_0 - 1 = C\cos(\bar{\Delta}\tau)$ and $Y = 2P_{\pi/2} - 1 = C\sin(\bar{\Delta}\tau)$. Their length estimates contrast and their angle estimates accumulated phase modulo 2π. Suppose at $0.50\,\mathrm{ms}$ the readings are 0.500 and 0.850. Then $X = 0,\ Y = 0.700$, and the phase is $\pi/2$ modulo 2π. With an independent bound $|\bar{\Delta}/2\pi| < 1\,\mathrm{kHz}$, the implied cyclic detuning is $+500\,\mathrm{Hz}$. Longer delays resolve fine frequency differences at the cost of phase ambiguity and declining contrast. Use a shorter-delay record to unwrap a longer one.

At a contrast of 0.700 the relation $-\ln(0.700) = \tau/T_2 + \sigma_\Delta^2\tau^2/2$ supplies one equation for two unknowns. Several delays, with independent error bars, are needed to identify the linear and quadratic terms. At very low

signal the length $\sqrt{X^2 + Y^2}$ has a positive noise bias. Fitting the two directly measured quadratures can be preferable to fitting the length.

Explore two mechanisms

The graph holds the homogeneous T_2 at 2 ms. It lets you set signed cyclic detuning and shot-to-shot cyclic frequency spread. Compare both quadrature populations with the upper contrast envelope. First reverse the sign of the detuning, then increase the spread without changing the mean. Which effect interchanges peaks of one quadrature and which damps both?

Open this experiment online

Try the calculation

Mean cyclic detuning (Hz) **500**

Shot-to-shot rms spread (Hz) **80**

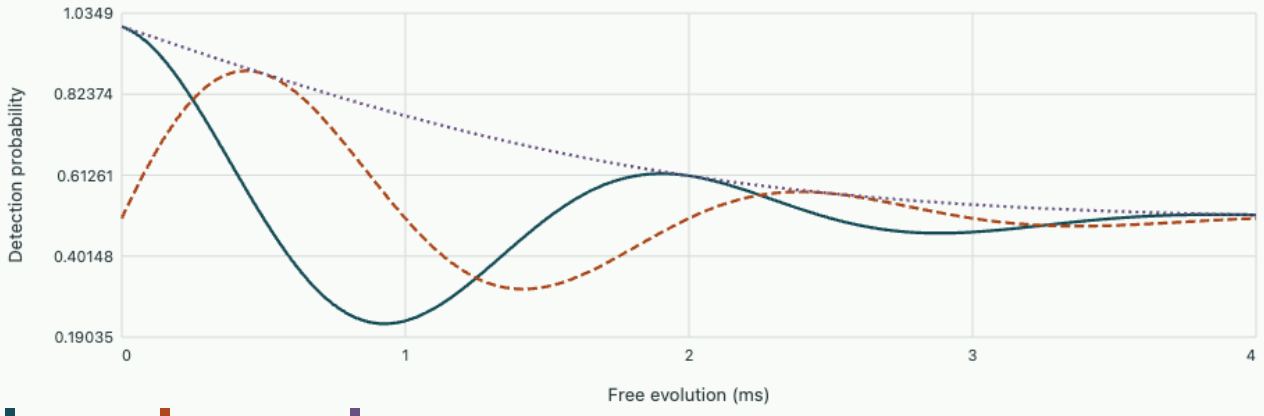

Contrast after 1 ms = 0.53455. The two probabilities at 0.50 ms are 0.5 and 0.87729.

Tutorial

What is the fringe spacing in a scan of cyclic detuning at fixed delay of 0.80 ms?

▼ **Work through the solution**

One full phase cycle needs $\delta f \, \tau = 1$. The spacing is

$$1/0.00080 = 1250\,\mathrm{Hz}.$$

The angular-frequency spacing is 2π times larger.

A fixed delay gives

$$P_0 = P_{\pi/2} = 0.5.$$

Does that show a phase of 45 degrees?

▼ **Work through the solution**

Both centred quadratures are zero, so the contrast is unresolved and the angle is undefined. A phase of 45 degrees with finite contrast gives equal, nonzero centred quadratures.

Explain how a refocusing pulse halfway through the evolution helps identify quasistatic detuning noise.

▼ **Work through the solution**

The pulse exchanges the phase-accumulating states, reversing phase winding for a detuning that stays constant through both halves. Random variations during the interval need not cancel. Recovering contrast in the refocused sequence supports a slowly varying contribution to the original envelope.

Compare a continuous drive in the transition chapter with the role of the second pulse here.

Computational quantum experiments · 4. When three perfectly legal spin measurements imply an impossible state

Prepare nominally identical spin-half systems and measure along three orthogonal axes, using a different subensemble for each axis. Reconstruct a density matrix, diagnose an unphysical linear estimate and see why the choice of estimator matters in finite data.

Translate detector counts into three components

Write $\rho = (I + \mathbf{r}\cdot\boldsymbol{\sigma})/2$. For measurement along a unit direction $\mathbf{n}$, the probability of spin up is

$$p_+ = (1 + \mathbf{n}\cdot\mathbf{r})/2.$$

With axes x, y, z, count k_j positive outcomes in N_j trials. Linear inversion gives

$$r_j = 2k_j/N_j - 1.$$

The trace condition is already built into this parametrisation. Positivity remains a separate physical requirement. Using the Pauli product identity, $(\mathbf{r}\cdot\boldsymbol{\sigma})^2 = |\mathbf{r}|^2 I$, the eigenvalues are $(1 \pm |\mathbf{r}|)/2$. They form legitimate probabilities only for $|\mathbf{r}| \le 1$.

Each axis can report a frequency in $[0, 1]$ even when the three reconstructed components have length greater than one. This need not imply a broken detector. The axes are measured on different sets of particles and can fluctuate in mutually inconsistent directions. A state that slowly drifts between the measurements, however, can produce deviations that persist as more particles are counted. Interleaving the settings can test that explanation.

Reconstruct a specific estimate

In 25 shots per axis an experiment obtains positive counts $23, 19, 19$. The reconstructed components are $0.84, 0.52, 0.52$, with norm about 1.1165. The smaller eigenvalue is approximately −0.058. A fourth analyser set opposite the inferred Bloch direction would therefore receive a predicted negative probability from this linear model. That prediction exposes the issue independently of any particular plotting convention.

A simple constrained estimate minimises Euclidean distance from the measured vector to the closed unit ball. If the vector lies inside, the minimiser is the measured vector. Outside, a Lagrange multiplier or elementary projection shows the minimiser is the vector divided by its length. One can now predict valid probabilities for every analyser direction. The result minimises a specified quadratic loss; it need not maximise the likelihood of the original counts.

Why a likelihood fit weights measurements differently

For independent binomial trials, the log likelihood up to constants is $\sum_j [k_j \ln(1 + r_j) + (N_j - k_j)\ln(1 - r_j)]$. The positive physical state constraint couples the components. Differentiating each term gives a score $[2k_j - N_j - N_j r_j]/(1 - r_j^2)$. Thus the cost of moving a component away from its unconstrained estimate depends on both its count and its location. A radial Euclidean projection ignores these different weights. Doubling only N_z while leaving its relative count unchanged preserves the

inversion vector but changes the evidence for retaining its z component. A fit using likelihood can account for that evidence.

The shape of the confidence region is a further question. An unconstrained Gaussian ellipse may extend beyond physical states near the surface. A physical point estimate does not by itself specify an uncertainty interval; a measurement report should include both the estimator and the count data or a stated statistical model.

Evaluate a new measurement using the two estimates

Fix the first count at 23 of 25. Vary the positive counts on y, z, then scan a test analyser in the plane through the estimated direction and an orthogonal axis. The graph plots probability versus angle from the raw vector. Radial projection preserves that direction. Predict which orientation first gives an impossible probability when the estimate moves outside the ball. A truly independent fourth measurement near that direction is more informative about the failure than a repeated plot of the three fitted counts.

Open this experiment online

Try the calculation

Up counts along y (out of 25) **19**

Up counts along z (out of 25) **19**

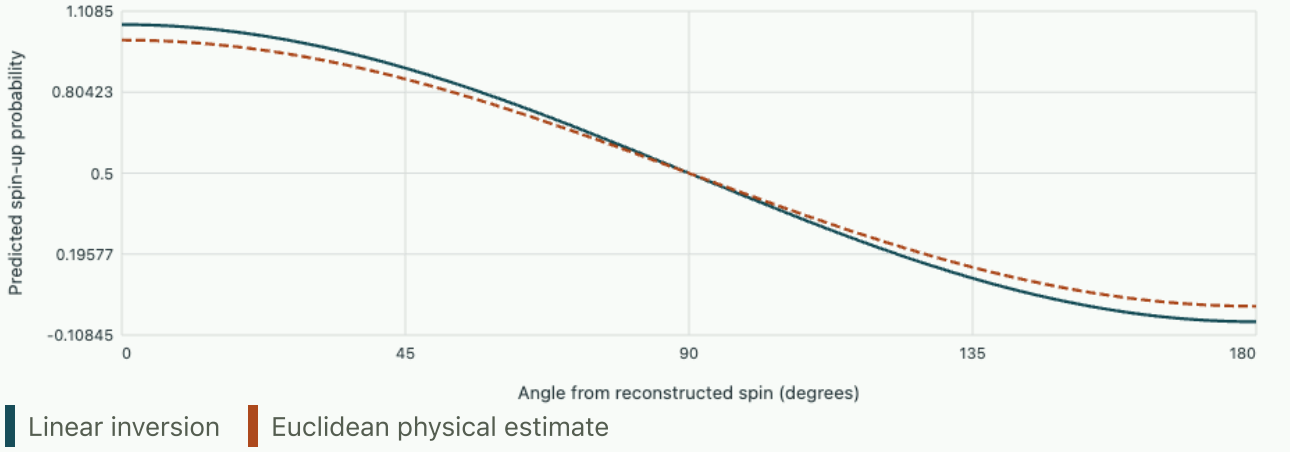

Reconstructed Bloch radius 1.1164. Eigenvalues 1.0582 and -0.058211.

Tutorial

All three directions are measured on 100 independent systems and the counts are 50, 50 and 50. Does this prove that every specimen was prepared in a mixed state?

▼ **Work through the solution**

The estimate is $I/2$ as a description of the measured ensemble. Many preparations have the same density operator, including equally likely pure states pointing along opposite directions. The counts do not identify one particular preparation history.

A data analysis repairs a negative eigenvalue by setting it to zero while retaining the original positive eigenvalue unchanged. Which density-matrix condition now fails?

▼ **Work through the solution**

The trace was initially one. Removing a negative number increases that trace above one. A valid construction must enforce positivity and normalisation together. Projection onto the state ball automatically gives both for a two-level system.

Why is it incorrect to send a single qubit successively through three projective analysers and use the resulting counts in this inversion?

▼ **Work through the solution**

A projective measurement changes the state entering the next analyser. The likelihood above assumes independent preparations of a common input state. Sequential data instead require the conditional measurement operations and their ordering.

Reconstruct how a partial trace produces a mixed state in the entanglement and measurement chapter.

Computational quantum experiments · 5. Read wave-packet overlap from a coincidence dip

Send one photon into each of two input ports of a balanced lossless beam splitter. Two alternatives lead to one click at each output. Derive the coincidence probability when their wave packets overlap only partially, then determine how a timing scan separates pulse duration from distinguishability in another degree of freedom.

Track the two indistinguishable alternatives

Choose input creation operators $a^\dagger, b^\dagger$ and outputs $c^\dagger, d^\dagger$. A possible symmetric unitary convention maps $a^\dagger$ to $(c^\dagger + id^\dagger)/\sqrt{2}$ and $b^\dagger$ to $(ic^\dagger + d^\dagger)/\sqrt{2}$. For input $a^\dagger b^\dagger|0\rangle$, expand the product. The two one-in-each-output terms have coefficients $1/2$ and $i^2/2 = -1/2$. They cancel for otherwise identical photons. The same-port outputs remain and their squared probabilities add to one after using the $\sqrt{2}$ normalisation of a two-photon mode. The photons do not interact with each other; bosonic amplitudes for indistinguishable output records interfere.

Allow the internal wave packets to differ

Let the photons carry normalised internal states f, g, which may include arrival time, polarisation and spectrum. In the coincidence sector the two possible output states are proportional to $c_f^\dagger d_g^\dagger|0\rangle$ and $c_g^\dagger d_f^\dagger|0\rangle$. Their overlap is $\langle f|g\rangle\langle g|f\rangle = |\langle f|g\rangle|^2$. Squaring the difference of their amplitudes gives

$$P_{\text{coinc}} = \tfrac{1}{2}(1 - |\langle f|g\rangle|^2).$$

Orthogonal internal states give one half, the value obtained by adding alternative paths as independent probabilities. A common overall optical phase of either photon leaves the modulus unchanged. Delay matters because it changes the overlap in time, even without a locked optical phase between the two sources.

Background · Can a polariser restore a missing coincidence dip?

Prepare one horizontally polarised and one vertically polarised photon in otherwise identical temporal and spatial modes. Their internal overlap is zero, giving the distinguishable coincidence probability of one half at a balanced splitter. Put identical 45° polarisers in the two input arms. Each photon passes with probability one half. Conditional on both passing, their polarisation states now agree.

$$P(\text{two pass}) = \frac{1}{4},$$

$$P(\text{coincidence} \mid \text{two pass, zero delay}) = 0.$$

At large delay the conditional coincidence probability returns to one half, so the coincidence probability per incident pair after the polarisers is one eighth. The dip can regain ideal visibility while the available event rate falls. Reporting the transmission or raw count rate along with visibility distinguishes filtering from an improvement in the source brightness.

Compute a dip with a specified width convention

Take a transform-limited Gaussian amplitude proportional to $e^{-t^2/(4\sigma^2)}$. Its detection-time probability has standard deviation σ. Delay one photon by τ. Completing the square in the overlap integral gives amplitude overlap $e^{-\tau^2/(8\sigma^2)}$ and squared overlap $e^{-\tau^2/(4\sigma^2)}$. If both photons also have linear polarisations separated by α, multiply the squared overlap by $\cos^2\alpha$. Thus the dip has full width at half depth $4\sigma\sqrt{\ln 2}$, independent of α as long as a dip remains measurable.

For $\sigma = 80$ fs and $\alpha = 30^\circ$, the coincidence probability is 0.125 at zero delay and approaches 0.5 at long delay. The half-depth level is 0.3125 at $|\tau| \simeq 133$ fs, giving full width about 266 fs. A width measured at probability 0.25 instead would not in general represent half depth. Quote the baseline, minimum and width convention if comparing different measurements.

Decide what a coincidence counter can identify

Adjust pulse duration and relative polarisation. Study peak-to-baseline depth before using the width to infer duration. A reduced visibility might equally arise from unobserved spectral distinguishability, multi-photon contamination or uncorrelated background counts. Blocking each input arm provides an independent accidental-count estimate. The ideal probability plotted assumes exactly one photon in each input and observation of both outputs irrespective of arrival time within the accepted window. A narrow coincidence window makes detection conditioning part of the model.

Open this experiment online

Try the calculation

Single-photon rms intensity width σ (fs) **80**

Relative polarisation angle (degrees) **30**

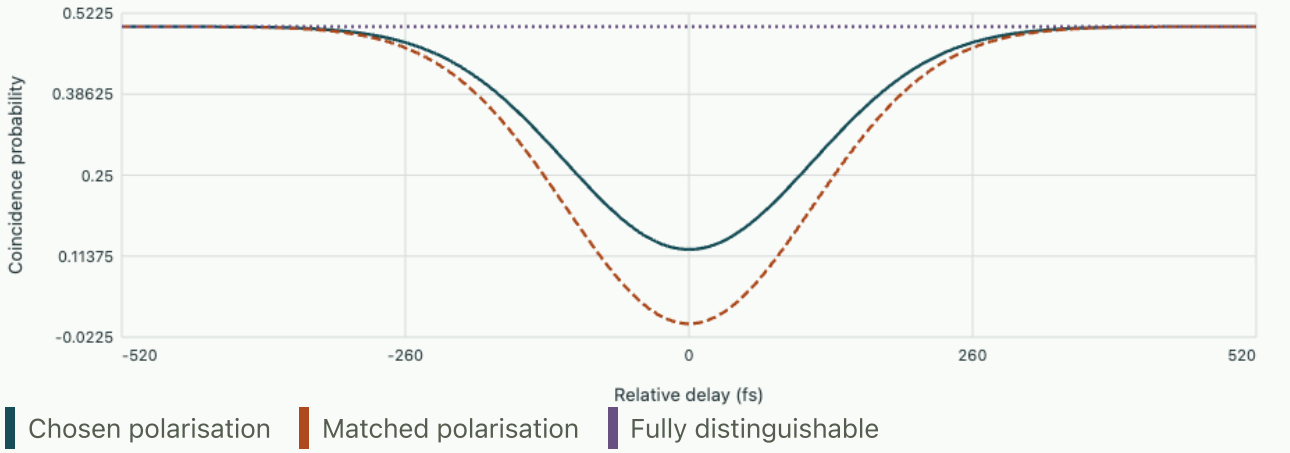

Zero-delay coincidence 0.125. Full width at half depth 266.42 fs.

Tutorial

Orthogonal linear polarisations yield no dip even at zero delay. Does a slower detector that cannot resolve polarisation automatically restore it?

▼ Work through the solution

No. The outgoing internal states remain orthogonal. Summing over an unobserved label is equivalent to tracing it out and removes the interference term. Deliberate projection of both polarisations onto a common state is another experiment with a different success probability.

For a lossless beam splitter of intensity reflectance R and transmittance $T = 1 - R$, predict the coincidence probabilities for perfectly distinguishable and identical photons.

▼ Work through the solution

The distinguishable alternatives give $R^2 + T^2$. For identical packets the amplitudes subtract, giving $(R - T)^2$. A perfectly indistinguishable pair need not give a zero minimum at an unbalanced splitter.

What delay has squared temporal overlap one half if the individual photon intensity has 100 fs rms width?

▼ Work through the solution

Set $\tau^2/(4\sigma^2) = \ln 2$, giving $|\tau| = 2(100)\sqrt{\ln 2} = 166.5$ fs. The measured full width spans both signs, about 333 fs.

Follow the single-photon and many-particle descriptions in the exchange chapter.

Computational quantum experiments · 6. Can an atom leave its initial state and still pass the final test?

A driven two-level atom undergoes repeated ideal measurements. Compare two laboratory records, one retaining only runs that passed every test and another retaining every run. The distinction is visible even when every measurement is made with the same apparatus.

Count histories or evolve a population

Write the resonant Hamiltonian in the rotating frame as $H = \hbar\Omega\sigma_x/2$. A pulse of duration T/N rotates the initial ground state into $\cos(\alpha/2)|0\rangle - i\sin(\alpha/2)|1\rangle$ with $\alpha = \Omega T/N$. The probability of passing one ground-state test is $c = \cos^2(\alpha/2)$. Each successful projection prepares the same ground state for the next interval. The probability that a complete record contains only successes is consequently $P_{\text{all}} = c^N$.

For an unread measurement, keep the atoms detected in either state. The projection removes the off-diagonal density-matrix elements, leaving a classical mixture just before the next pulse. If p_j denotes the ground fraction after measurement j, the next pulse maps it to $p_{j+1} = cp_j + (1-c)(1-p_j)$. The excess population $z_j = 2p_j - 1$ then obeys $z_{j+1} = (2c-1)z_j$. Starting in the ground state gives

$$P_{\text{final}} = \frac{1 + \cos^N(\Omega T/N)}{2}.$$

Leaving the ground state in one interval and returning in a later one contributes to the final fraction. It removes that run from the all-success record. The two probabilities therefore refer to different experimental questions, even though they use exactly the same sequence of pulses and measurements.

Follow a half-turn in two steps

Choose $\Omega T = \pi$ and make two evenly spaced measurements, including the one at the end. Each pulse is a quarter-turn of the Bloch vector and gives either result with equal probability following either previous outcome. The four two-result strings have probability one quarter each. Only 00 passes every test, whereas 00 and 10 both finish in the ground state. The fractions are therefore 0.25 and 0.50. With a single final measurement the undisturbed half-turn ends entirely in the excited state.

At fixed total duration and very large N, Taylor expansion gives

$$\log P_{\text{all}} \simeq -(\Omega T)^2/(4N).$$

The survival probability tends to one because the transition probability per interval is quadratic in its short duration. This prediction assumes ideal instantaneous projective measurements and coherent rotations between them. A measurement with a fixed finite duration cannot be packed arbitrarily densely into a fixed time.

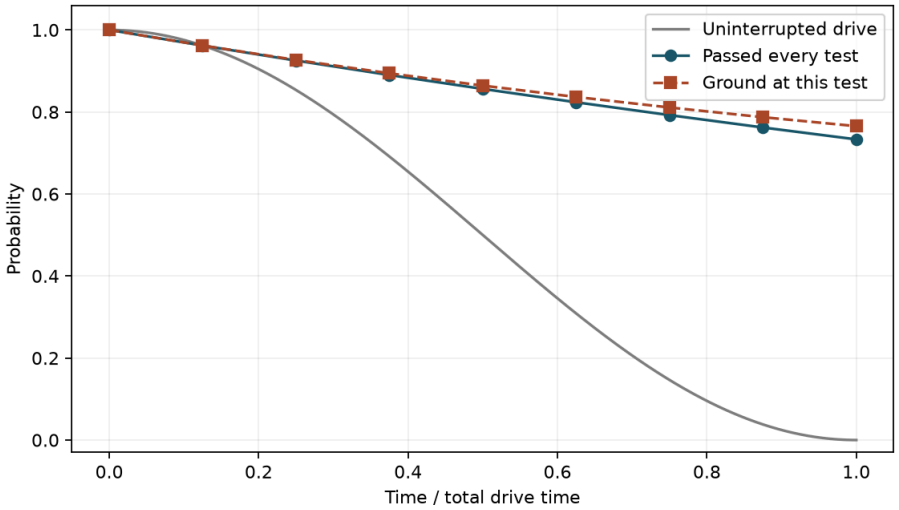

Markers represent measurements. The continuous reference curve is the ground population of an atom subjected to the same drive with every intermediate test omitted.

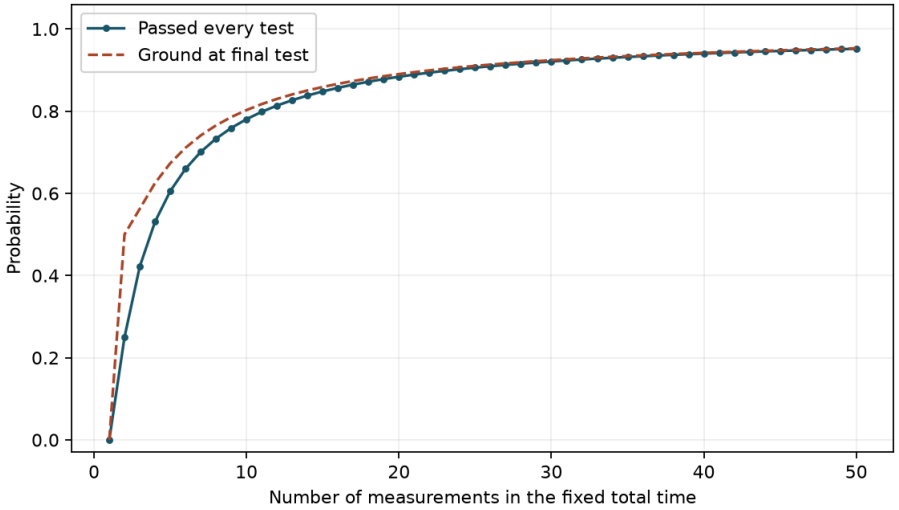

Keep the pulse area and total time fixed whilst shortening each free interval. An increasingly dense train of ideal observations preserves the initial state.

Exercise · recover a mixed population

If the atom initially occupies the ground state with probability 0.7 and has no coherences, find the final ground probability for four measurements and a total pulse area of π.

▼ **Work through the solution**

The initial contrast is 0.4. Each interval multiplies it by $\cos(\pi/4) = 1/\sqrt{2}$. Four steps give contrast 0.1, so the final ground fraction is $(1 + 0.1)/2 = 0.55$. The all-pass probability would instead contain a first measurement conditioned on the initially mixed preparation.

Exercise · budget the number of tests

Estimate the number of ideal tests needed for 99 percent all-pass survival during a coherent half-turn. Check the estimate with the editable input.

▼ Work through the solution

Equating the logarithm of 0.99 to $-\pi^2/(4N)$ gives $N \simeq 245.5$. Rounding up to 246 and evaluating the exact cosine power gives about 0.99002. This estimate concerns survival over an entire record. Filtering out only the atoms ending in the excited state would estimate a different probability.

Try it · Zeno

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Total pulse area / π 1

Intermediate measurements 8

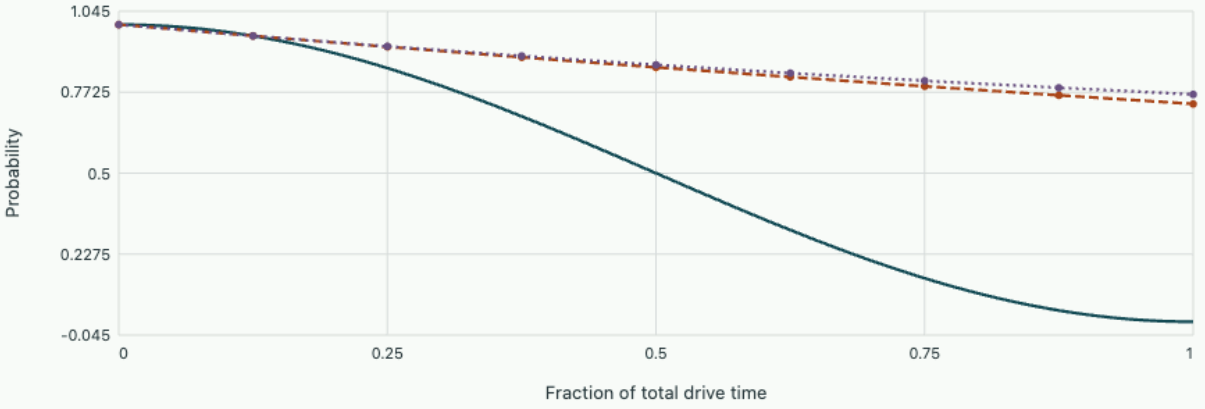

Selective survival 0.73313; unread final ground population 0.7654. Ideal instantaneous projective measurements are assumed.

Download the corresponding editable notebook. The browser recalculates the reduced equations directly; the notebook exposes the full calculation and its checks.

Open editable notebook Download .ipynb

Translational symmetry, bands and edges

Apply a translation operator, reduce a many-site Hamiltonian to a small matrix and use the resulting eigenvectors to predict an edge state. Compare a band touching with a gapped system.

1. Translational symmetry and diffraction
2. Bloch waves, band edges and effective mass
3. Predicting a gap and testing it in a periodic medium
4. How alternating bonds split a one-dimensional band
5. Finding an end state one bond at a time
6. Why a honeycomb lattice has Dirac cones
7. When a periodic dielectric has no propagating Bloch wave
8. Action density, stress waves and the boundary that reflects them

Translational symmetry and diffraction

Condensed matter and band theory

Work backwards from a scattering pattern to the repeating translations and the atoms within one cell. Learn which structural information an intensity measurement loses.

FROM FIRST PRINCIPLES

Recover a lattice spacing from phase addition

Illuminate a row of N equally spaced atoms with a coherent plane wave and observe weak elastic scattering far away. Each atom sees an incident phase and contributes an outgoing phase. Their difference depends on $\mathbf{q} = \mathbf{k}_{\text{out}} - \mathbf{k}_{\text{in}}$. With the origin at the first atom, let $\phi = \mathbf{q} \cdot \mathbf{a}$ be the phase difference between neighbours displaced by $\mathbf{a}$. Up to an overall sign convention the total amplitude is a geometric sum:

$$A = f \sum_{j=0}^{N-1} e^{ij\phi} = f \, \frac{1 - e^{iN\phi}}{1 - e^{i\phi}}.$$

Here f is the amplitude from one atom. Use $|1 - e^{ix}| = 2|\sin(x/2)|$ and take the squared magnitude. When $\phi = 2\pi m$, evaluate the original sum directly:

$$I \propto |f|^2 \frac{\sin^2(N\phi/2)}{\sin^2(\phi/2)},$$

$$I_{\text{peak}} \propto N^2 |f|^2.$$

For a three-dimensional lattice every primitive translation must meet this phase condition. The resulting peak positions identify reciprocal-lattice vectors. For several atoms in each cell the single-cell factor becomes $\sum_s f_s e^{i\mathbf{q}\cdot\mathbf{r}_s}$. Its squared magnitude can vanish even at a permitted reciprocal vector. An absent peak can therefore reveal interference inside the cell.

Background · Reciprocal lattice

$$\mathbf{G} \cdot \mathbf{R} = 2\pi n$$

A reciprocal-lattice vector leaves the phase unchanged under every direct-lattice translation. A basis of different atoms can still suppress a reflection.

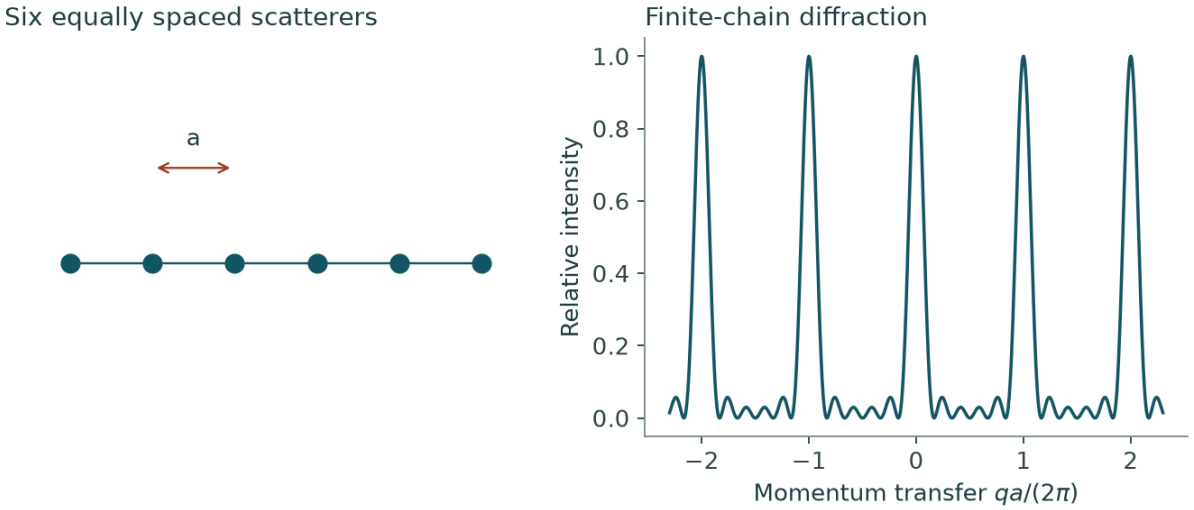

Why a finite crystal gives diffraction peaks of nonzero width

Successive sites acquire a relative phase qa. At qa = 2π every contribution returns in phase. The first zero on either side is offset by 2π/6 because six equally spaced phase vectors sum to zero. Doubling the number of sites halves this offset while leaving the peak spacing fixed.

Explore the related parameter study beside the computed figure.

DERIVATION AND APPLICATIONS

Extract atomic displacement from the loss of Bragg intensity

A lattice can keep the same mean spacing while its Bragg peaks lose intensity on heating. Consider a crystal of identical atoms whose displacements about their mean positions fluctuate during the exposure. Treat scattering in the single-scattering approximation. At a reciprocal vector of magnitude Q, every mean lattice position supplies the same phase. The instantaneous displacement projected along the scattering vector, u, adds a phase Qu. If this component is Gaussian with mean zero and variance s^2, averaging the complex amplitude gives

$$\left\langle e^{iQu} \right\rangle = \frac{1}{\sqrt{2\pi s^2}} \int_{-\infty}^{\infty} e^{-u^2/(2s^2)} e^{iQu}\, du = e^{-Q^2 s^2/2}.$$

The last equality follows by Fourier transformation of the normalized Gaussian. The coherent peak measures the square of this average amplitude. Squaring the instantaneous phase factor first would give unity and erase the effect of phase differences between sites. With $I_0(Q)$ the coherent intensity of the perfectly stationary array, the Debye–Waller attenuation is

$$I_{\mathrm{Bragg}}(Q) = I_0(Q) e^{-Q^2 s^2}.$$

The formula describes the sharp coherent contribution. The difference between the total averaged intensity and the coherent intensity appears in diffuse scattering. Stronger attenuation at large Q follows because an identical displacement spans a larger fraction of a scattering phase cycle. For corrected integrated intensities of the same reflection at two temperatures, taken at essentially the same Q, the unknown stationary intensity cancels,

$$s_2^2 - s_1^2 = -\frac{\ln[I_2(Q)/I_1(Q)]}{Q^2}.$$

This comparison assumes unchanged occupancies, instrument normalization and negligible changes in the mean structure. For isotropic displacement s^2 is one third of the mean squared three-dimensional displacement.

WORKED EXAMPLE

Recover a change in mean-square displacement

Change the inputs and investigate this derivation.

A reflection at $Q = 50.0\,\mathrm{nm^{-1}}$ falls from 10,000 to 8,000 corrected counts. Its projected variance increases by $-\ln(0.800)/2500 = 8.93 \times 10^{-5}\,\mathrm{nm^2}$. If its original projected root-mean-square displacement was $0.00600\,\mathrm{nm}$, the new value is

$$\sqrt{0.00600^2 + 0.0000893} = 0.0112\,\mathrm{nm}.$$

Adding the displacements directly would use the wrong measured quantity.

SECOND APPLICATION

Test the angular dependence with another reflection

Change the inputs and investigate this derivation.

A second specimen retains 90.0 percent of its low-temperature peak intensity at $Q = 30.0\,\mathrm{nm^{-1}}$. For the same displacement direction at twice that wavevector, the attenuation exponent quadruples. The predicted retention is

$$0.900^4 = 0.656.$$

A corrected retention near 0.90 for both reflections would instead suggest a common change in illuminated volume or flux that this displacement model cannot explain.

Try it · Debye waller

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Projected variance increase (nm²) **9.00e-5**

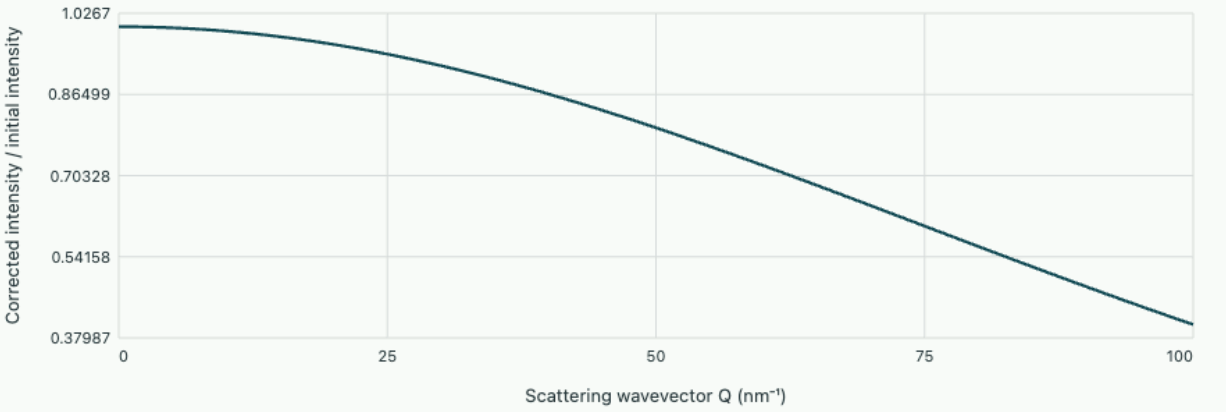

Gaussian displacement attenuation

Intensity ratio is exp(−Q²Δ⟨u²⟩). Doubling Q quadruples the attenuation exponent for the same projected displacement distribution.

What repeats in a crystal?

A one-dimensional chain with identical atoms every a has translations $R = na$. Replacing every other atom by a different species doubles the primitive translation, even if the nearest-neighbour separation stays fixed. The set of equivalent cell origins is the Bravais lattice; the atoms placed relative to each origin form the basis. Choosing a larger unit cell is allowed but folds distinct waves onto extra branches. Count independent degrees of freedom per unit length as a check on any cell choice.

In three dimensions write $\mathbf{R} = \sum_i n_i \mathbf{a}_i$. Define reciprocal vectors by $\mathbf{a}_i \cdot \mathbf{b}_j = 2\pi\delta_{ij}$. Then $\mathbf{G} = \sum_i h_i \mathbf{b}_i$ has $e^{i\mathbf{G}\cdot\mathbf{R}} = 1$ at every equivalent origin. These vectors have units of inverse length. For a rectangular lattice doubling a real-space period halves the corresponding reciprocal spacing. The first Brillouin zone consists of wavevectors closer to the reciprocal origin than to any other reciprocal point. Later we will label collective vibrations and electronic states inside this region.

Construct reciprocal space from the primitive cell

A diffraction experiment measures momentum transfer. To predict its possible values, ask which plane waves acquire the same phase at every translated lattice point. For a Bravais lattice with translations $\mathbf{R} = \sum_i n_i \mathbf{a}_i$, the condition is $\exp(i\mathbf{G}\cdot\mathbf{R}) = 1$ for every integer choice of n_i. It suffices to require

$$\mathbf{G}\cdot\mathbf{a}_i = 2\pi h_i.$$

Collect the primitive direct-lattice vectors as columns of A and their reciprocal vectors as columns of B. The defining dot products become one matrix equation. Solving it gives

$$\begin{aligned} A^{\mathsf{T}}B &= 2\pi I, \\ B &= 2\pi (A^{-1})^{\mathsf{T}}. \end{aligned}$$

Each column of B has units of inverse length. Its dot product with its own direct-lattice vector is 2π and with every other basis vector is zero. Integer combinations generate the allowed reciprocal vectors. The transpose matters whenever the direct axes are oblique.

For a triangular layer take $\mathbf{a}_1 = (a, 0)$ and $\mathbf{a}_2 = (a/2, \sqrt{3}a/2)$. The constraints on $\mathbf{b}_1$ first give $b_{1x} = 2\pi/a$ and then $b_{1y} = -2\pi/(\sqrt{3}a)$. Similarly $\mathbf{b}_2 = (0, 4\pi/(\sqrt{3}a))$. Verify that translating along $\mathbf{a}_2$ leaves the phase of $\mathbf{b}_1$ unchanged. Reciprocal vectors have units of inverse length.

How far can a wave vector move from the origin before an equivalent representative is nearer? The boundary shared with $\mathbf{G}$ solves $|\mathbf{k}| = |\mathbf{k} - \mathbf{G}|$, or $\mathbf{k}\cdot\mathbf{G} = |\mathbf{G}|^2/2$. The half-planes around the origin define the first Brillouin zone. Here the six shortest reciprocal vectors give a hexagon. With $a = 0.300\,\mathrm{nm}$ the shortest $|\mathbf{G}|$ is $4\pi/(\sqrt{3}a) = 24.18\,\mathrm{nm}^{-1}$, putting a face midpoint $12.09\,\mathrm{nm}^{-1}$ from the centre. The zone area is $(2\pi)^2/(\sqrt{3}a^2/2) = 506.5\,\mathrm{nm}^{-2}$, equal to one reciprocal primitive-cell area. Doubling all lattice lengths halves the distances in the diffraction pattern and quarters its zone area. These reciprocal positions follow from translation symmetry. The measured intensities also depend on what occupies each cell.

Add amplitudes before taking intensities

An elastic plane wave scattered from positions $na + d_j$ accumulates a phase equal to the momentum-transfer wavevector Q times position. For N identical cells the amplitude factorises:

$$A(Q) = \underbrace{\sum_j f_j(Q) e^{iQd_j}}_{F(Q)} \sum_{n=0}^{N-1} e^{iQna},$$

$$I(Q) \propto |A(Q)|^2.$$

The last sum is geometric. Its magnitude is $|\sin(NQa/2)/\sin(Qa/2)|$, with limiting value N at $Qa = 2\pi h$. Large finite crystals have narrow peaks of height proportional to N^2. The separation to the adjacent first zero is $2\pi/(Na)$; doubling coherent length halves this width. The basis factor F sets the relative integrated strengths for a given lattice. A finite instrument resolution also broadens measured peaks and should be characterised before inferring a sample size.

Background · Why are diffraction peaks from small crystals broad?

Along one direction take N equally spaced scattering planes at $x_n = na$. Near a reciprocal-lattice vector G, an additional wavevector mismatch Δq accumulates phase $na\Delta q$ at the nth plane. Summing the complex contributions gives a geometric series.

$$\frac{I(G+\Delta q)}{I(G)} = \left[\frac{\sin(Na\Delta q/2)}{N\sin(a\Delta q/2)}\right]^2.$$

The nearest zero lies at $2\pi/(Na)$. A coherently scattering length of 100 nm therefore produces a first-zero offset near 0.063 nm^{-1}. A finite width alone does not determine a lattice strain. Finite size sets a scale in absolute wavevector, whereas a distribution of fractional lattice spacings tends to give a width proportional to the reflection wavevector. Comparing several orders helps separate the two mechanisms.

An extinction is evidence about the basis

Place two sites at $d = 0, a/2$, with real scattering amplitudes $f_A = 3$ and $f_B = 1$ held constant over the angular range. At $G_h = 2\pi h/a$, $F_h = 3 + (-1)^h$. Even and odd orders have intensities proportional to 16 and 4 respectively, before angular or instrumental corrections. If both amplitudes become equal, the odd peaks vanish. If the two sites have identical chemistry as well as identical positions up to translation, a cell of length $a/2$ would describe the structure more economically. Equal scattering amplitudes in a particular probe alone cannot establish chemical identity.

For incident wavelength λ, with the total scattering angle 2θ,

$$|Q| = (4\pi/\lambda)\sin\theta.$$

At $\lambda = 0.154$ nm and spacing $d = 0.200$ nm the first Bragg peak therefore occurs at $\theta = 22.6^\circ$, or detector angle 45.3°. A factor-of-two error in reading the detector would corrupt any inferred lattice constant.

What survives a change of origin?

Translating the entire basis by d_0 multiplies $A(Q)$ by e^{iQd_0}. Intensity is unchanged, so diffraction cannot locate the absolute origin. More generally the phase of the amplitude is missing from a measurement of $|A|^2$. Peak positions suggest candidate lattices, corrected intensities test a proposed basis and peak widths probe finite coherence or strain. These are different questions even when addressed with the same scan.

Derive the selection rules of a cubic powder pattern

A powdered crystal supplies many crystallite orientations, so a fixed detector angle can intercept every family of reciprocal planes with the appropriate spacing. For a cubic lattice with conventional cell parameter a, reciprocal vectors have magnitude $2\pi\sqrt{h^2+k^2+l^2}/a$. Bragg's condition then gives

$$\sin^2\theta = \lambda^2(h^2+k^2+l^2)/(4a^2).$$

The sequence of squared sines can be used to index the pattern, but only reciprocal points with nonzero structure amplitude produce peaks.

For a monatomic face-centred cubic crystal, describe the lattice with conventional cubic coordinates and basis positions $(0,0,0)$, $(0,\tfrac12,\tfrac12)$, $(\tfrac12,0,\tfrac12)$ and $(\tfrac12,\tfrac12,0)$. For identical atomic amplitude f their phase sum is

$$F_{hkl} = f[1 + (-1)^{k+l} + (-1)^{h+l} + (-1)^{h+k}].$$

If all indices are even or all are odd, every phase factor is positive and $F = 4f$. Mixed parity makes the sum zero. Thus the first allowed squared-index sums are 3 for (111), 4 for (200), 8 for (220), 11 for (311) and 12 for (222). These extinctions encode the centring translations; adding an internal basis to each lattice site can introduce additional missing peaks. Powder intensity also depends on multiplicity and on angle-dependent scattering factors, so the strongest peak need not have the smallest indices.

Infer a lattice parameter and a strain

Suppose powder peaks under wavelength 0.1540 nm occur at detector angles $2\theta = 38.48^\circ$, 44.74° and 65.10°. Their squared half-angle sines are about 0.1086, 0.1448 and 0.2896. After division by the first they give $1 : 1.33 : 2.67$, consistent with $3 : 4 : 8$. Indexing the middle peak as (200) yields

$$a = \lambda / \sin 22.37^\circ = 0.405\,\mathrm{nm}.$$

The other two peaks independently give the same value to the displayed precision, unlike an arbitrary assignment to successive integers.

Now let the (200) detector angle fall by 0.100° under load. Differentiating $2d\sin\theta = \lambda$ gives

$$\Delta d/d \simeq -\cot\theta\,\Delta\theta.$$

Here $\Delta\theta = -0.0500^\circ = -8.73\times10^{-4}$ rad, hence

$$\Delta d/d = +2.12\times10^{-3}.$$

That is tensile strain along the scattering vector. A single peak shift cannot establish an isotropic volume strain without further orientations. Multiple peaks can also separate a uniform scale error in wavelength from a genuine directional strain. In the unstrained limit the lattice parameter extracted independently from every correctly indexed family should agree within instrumental uncertainty.

Peak position and width respond differently to imperfections. If all sampled crystallites undergo the same fractional change of their plane spacing, their peaks shift together by the differential Bragg relation. A distribution of tensile and compressive strains can instead broaden the peak around an almost unchanged mean angle. A finite crystallite size can also broaden it, but its angular dependence across several reflection families differs from the leading strain broadening. Comparing several indexed peaks therefore gives more information than treating a wide single line as direct evidence for a particular mean lattice parameter.

Instrumental angle calibration should be checked with a reference powder so that a detector zero shift does not masquerade as a specimen strain.

Tell a finite-size minimum from a basis extinction

Eight identical cells contain sites at 0 and half a lattice period. Sum the eight cell amplitudes coherently, then multiply by the basis amplitude $f_A + f_B e^{iQa/2}$. Here the form factors are real, positive and independent of angle.

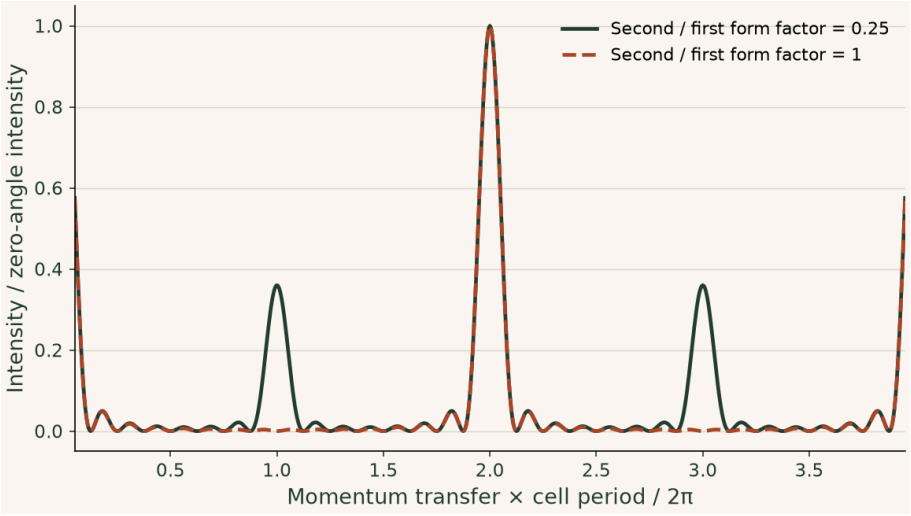

For equal form factors the odd orders vanish. With fB/fA = 1/4 an odd order has 36% of the corresponding even-order peak intensity. The subsidiary narrow fringes arise from the finite number of cells and would narrow in a longer crystal.

Try it · Diffraction chain

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Number of unit cells **8**

Second-site / first-site form factor **0.25**

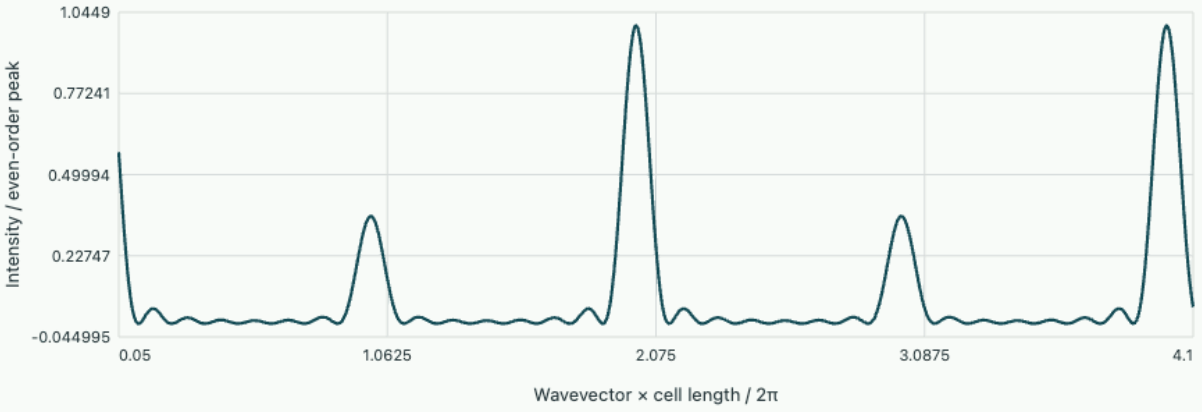

Finite two-site crystal

Odd/even principal peak ratio = 0.36. Increasing cell count narrows the peaks; the basis determines systematic absences.

▶ **Read values from the plot**

Bloch waves, band edges and effective mass

Condensed matter and band theory

Derive an electronic band from both weak and strong binding, then use its slope and curvature to predict motion.

FROM FIRST PRINCIPLES

Why translating a lattice state produces an energy band

In a perfect periodic chain the Hamiltonian is unchanged when the site labels are shifted by one. A wave may therefore have site amplitudes of the form $c_j = Ce^{ikja}$, with k defined up to $2\pi/a$. To give that phase a physical energy, choose one localized orbital per cell of spacing a. Assume an orthonormal basis, on-site energy ϵ_0, and a real hopping matrix element $-t$ to either nearest neighbour. Projecting the stationary Schrödinger equation on orbital j gives

$$Ec_j = \epsilon_0 c_j - t(c_{j-1} + c_{j+1}).$$

Successive amplitudes differ only by $e^{\pm ika}$. Cancel c_j, then combine these phases into a cosine:

$$E(k) = \epsilon_0 - t(e^{ika} + e^{-ika}) = \epsilon_0 - 2t\cos(ka).$$

For $t > 0$ the in-phase combination has the lowest energy. Comparing $k = 0$ and $k = \pi/a$ gives bandwidth $4t$. Form a narrow wave packet from neighbouring k's. Its envelope speed is

$$d(E/\hbar)/dk = (2ta/\hbar)\sin(ka).$$

Translation symmetry determines the allowed phase labels, while the coupling strength determines the numerical dispersion. A real crystal can require several orbitals and longer-range couplings.

Background · Band velocity

$$\mathbf{v}_n(\mathbf{k}) = \hbar^{-1}\nabla_{\mathbf{k}} E_n$$

This is the group velocity of a band wave packet in the ordinary semiclassical approximation. The Fermi surface determines which states can respond at low temperature.

DERIVATION AND APPLICATIONS

Use band curvature to predict acceleration off the force axis

Near a smooth, isolated band minimum the first derivative of energy vanishes. In a low-symmetry crystal its second derivatives can depend on direction. Choose principal axes x, y in the plane near the minimum $\mathbf{k}_0$,

and write $\mathbf{q} = \mathbf{k} - \mathbf{k}_0$. Taylor expansion defines two positive effective masses through the curvatures,

$$E(\mathbf{k}_0 + \mathbf{q}) \simeq E_0 + \frac{\hbar^2 q_x^2}{2m_x} + \frac{\hbar^2 q_y^2}{2m_y}.$$

Treat a narrow wave packet semiclassically in a uniform force $\mathbf{F}$, over a time short enough for scattering to be negligible. Neglect interband transitions and any anomalous velocity. Differentiating energy component by component gives group velocity. Combining this with the crystal-momentum equation gives its acceleration,

$$v_i = \frac{1}{\hbar}\frac{\partial E}{\partial q_i} = \frac{\hbar q_i}{m_i},$$
$$\hbar \dot{q}_i = F_i,$$
$$a_i = \frac{F_i}{m_i}.$$

Thus a force rotated through the principal axes traces an ellipse of accelerations. To describe a longitudinal measurement, project the full vector onto the applied force. If the force makes angle θ to the x axis and ϕ is the direction of the resulting acceleration in the same quadrant, then

$$\frac{a_\parallel}{F} = \frac{\cos^2\theta}{m_x} + \frac{\sin^2\theta}{m_y},$$
$$\tan\phi = \frac{m_x}{m_y}\tan\theta.$$

The longitudinal inverse mass is a weighted average of inverse principal masses. The difference $\phi - \theta$ vanishes along principal axes and for an isotropic band. The expansion requires the whole accelerated packet to stay inside the region of roughly constant curvature.

WORKED EXAMPLE

Follow an oblique force pulse

Change the inputs and investigate this derivation.

Let $m_x = 0.20 m_e$, $m_y = 0.80 m_e$, where m_e is the free electron mass. A force pulse of magnitude $1.00 \times 10^{-17}\,\mathrm{N}$ acts at 45° for $50.0\,\mathrm{fs}$ on a packet initially at rest. Its velocity gains are $F\Delta t/(\sqrt{2}m_i)$, or $1.94\,\mathrm{m/s}$ along x and $0.485\,\mathrm{m/s}$ along y. The velocity points only 14.0° above the light-mass axis.

SECOND APPLICATION

Recover an unseen principal curvature

Change the inputs and investigate this derivation.

A second band has longitudinal mass $0.300 m_e$ for a force at 45°, and a measurement along its x axis gives

$$m_x = 0.200 m_e.$$

In units of $1/m_e$, inversion yields $m_e/m_y = 2/0.300 - 1/0.200 = 1.667$, so $m_y = 0.600 m_e$. A pulse along y should produce one third of the acceleration measured along x for the same force, providing an independent check of the principal-axis assignment.

Try it · Mass tensor

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Principal mass mx/me **0.2**

Principal mass my/me **0.8**

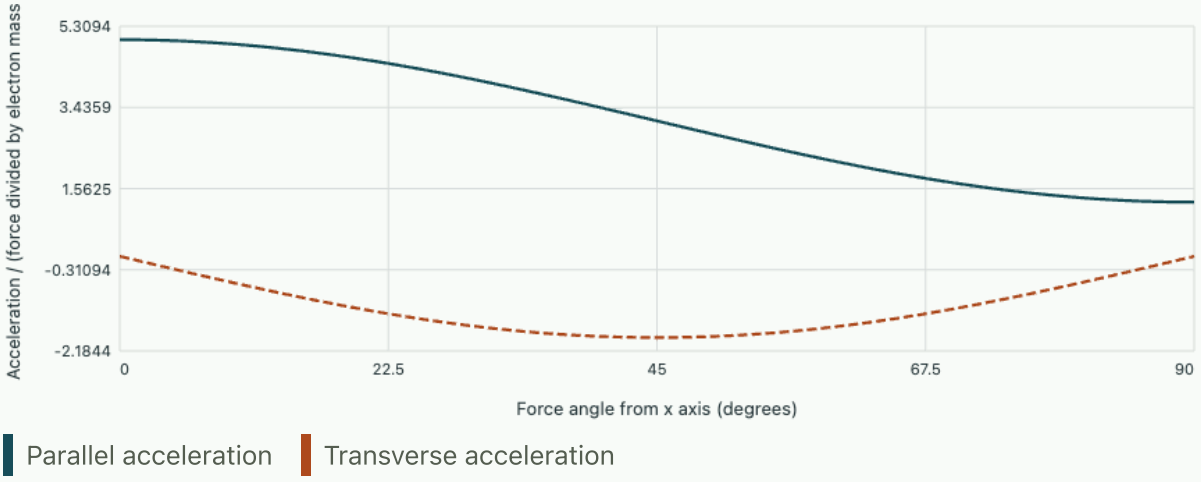

Acceleration need not align with force. The longitudinal inverse mass is cos²θ/mx + sin²θ/my.

Why label a state by crystal momentum?

A lattice translation by a commutes with the Hamiltonian when $V(x+a) = V(x)$. Translation eigenvalues have modulus one because translation preserves norm. Write them e^{ika}; then $\psi_k(x) = e^{ikx} u_k(x)$ with $u_k(x+a) = u_k(x)$. Since $e^{i(k+G)a} = e^{ika}$ for $G = 2\pi h/a$, a single Brillouin zone suffices. The index distinguishing states at the same k is the band index. Crystal momentum labels translation symmetry and can change by a reciprocal vector in scattering.

A weak potential couples a degenerate pair

Near $k = G/2$, free states $|k\rangle$ and $|k - G\rangle$ have nearly equal kinetic energies. The periodic potential has Fourier coefficients V_G that couple these waves. Absorbing the mean potential in both energies, the two-state characteristic equation is

$$(E - E_k)(E - E_{k-G}) = |V_G|^2,$$

with solutions

$$E_\pm = \tfrac{1}{2}(E_k + E_{k-G}) \pm \sqrt{\tfrac{1}{4}(E_k - E_{k-G})^2 + |V_G|^2}.$$

At equality the splitting is $2|V_G|$. For a real coefficient the states are standing waves whose densities peak at different locations relative to the potential. Far from the degeneracy, the square root is dominated by the kinetic splitting and one obtains two nearly free waves again. A Fourier component that vanishes by symmetry cannot open a first-order gap at that crossing. Check other bands before interpreting a local splitting as a gap for the entire solid.

Build a band from a localised orbital

Start instead with one localised orbital $|n\rangle$ per cell, all of energy ϵ_0, with overlap neglected in the basis and nearest-neighbour hopping matrix element $-t$. A translation eigenstate has site amplitudes $c_n = e^{ikna}/\sqrt{N}$. Acting with H on one site amplitude gives $Ec_n = \epsilon_0 c_n - t(c_{n-1} + c_{n+1})$, and hence $E(k) = \epsilon_0 - 2t\cos ka$. The bandwidth $4|t|$ measures the energy spread produced by hopping. In the limit of vanishing hopping the local energy repeats for every k. More than one orbital per cell generates a matrix eigenproblem even in this simple construction.

The group velocity is $v = \hbar^{-1} dE/dk$. Under a weak force with $\hbar\dot{k} = F$, acceleration is $FE''(k)/\hbar^2$. Define $m^* = \hbar^2/E''(k)$, a local measure of curvature. Near the bottom of the cosine band with $t > 0$,

$$E \simeq \epsilon_0 - 2t + t(ka)^2,$$

$$m^* = \hbar^2/(2ta^2).$$

Near the top the curvature changes sign. Describing the missing electrons in a nearly full band as positively charged holes with positive mass simplifies transport calculations.

Use filling to decide whether current can flow

For a complete time-reversal-symmetric band, the contributions of k and $-k$ to current cancel. More generally the integral of the derivative of a periodic band over a whole zone is zero. A weak field can redistribute carriers near a partially filled band edge, whereas a fully filled isolated band has no vacant neighbouring states. With one spin-degenerate orbital per primitive cell there are two states per cell. One electron per cell half fills this noninteracting band; two fill it. Symmetry breaking, interactions and additional bands can change the conclusion. A single measured zone-edge splitting therefore gives incomplete information about conduction.

For further work on three-dimensional bands, Fermi surfaces and electron motion, see David Tong's Solid State Physics lectures at Cambridge.

Open a gap by doubling a chain's period

A uniform one-dimensional chain with one spin-degenerate orbital and one electron per site has a half-filled band. A distortion can change both the spectrum and the periodicity used to count states. Let the sites be separated by a before the distortion. Group successive sites into A, B pairs and allow alternating hopping amplitudes $t_1, t_2 > 0$. The distorted unit cell has length $2a$. For equal on-site energies, chosen as zero, the amplitude on A couples to B inside its own cell and to B one cell away. Bloch translation supplies a phase e^{-2ika} for the latter coupling. The sublattice equations give a two-by-two matrix with zero diagonal and off-diagonal elements $-(t_1 + t_2 e^{\mp 2ika})$. Its characteristic equation yields

$$E_{\pm}(k) = \pm\sqrt{t_1^2 + t_2^2 + 2t_1 t_2 \cos(2ka)}, \qquad |k| \le \pi/(2a).$$

At the zone centre the hoppings add. At the new zone edge they subtract, opening a gap $2|t_1 - t_2|$. In the uniform limit the folded branches simply touch. Merely drawing a smaller zone did not change the crystal. It takes a physical alternation to split the states. With two electrons per doubled cell, the lower band is exactly filled when spin is counted. Its electronic energy falls on dimerization at fixed mean hopping: writing $t_{1,2} = t \pm \delta$ gives

$$E_{-} = -2\sqrt{t^2\cos^2(ka) + \delta^2 \sin^2(ka)},$$

which is at most the undistorted lower energy. The lattice also pays an elastic cost. Quantitative distortion amplitudes require that competing energy.

For $t_1 = 1.20\,\mathrm{eV}$ and $t_2 = 0.80\,\mathrm{eV}$ the zone-centre energies remain $\pm 2.00\,\mathrm{eV}$, while the zone-edge energies become $\pm 0.40\,\mathrm{eV}$. The minimum direct excitation costs $0.80\,\mathrm{eV}$, corresponding to a vacuum wavelength of about $1.55\,\mu\mathrm{m}$ if the optical transition is allowed. The cause of this insulating filling is the pairing of bonds. Away from the commensurate filling the chemical potential can lie inside a band even if the same structural splitting remains visible.

Follow a wave packet through the whole Brillouin zone

In a weak static field a band electron's crystal momentum changes even though its group velocity need not keep rising. Take one isolated tight-binding band $E(k) = E_0 - 2t\cos(ka)$, positive hopping energy t, lattice spacing a and electron charge $-e$. Let the wave packet be narrow in k compared with the zone and neglect scattering and transitions to other bands. The semiclassical force equation is $\hbar \dot{k} = -e\mathcal{E}$, while velocity follows from the band gradient:

$$v = (2ta/\hbar)\sin(ka).$$

Starting at $k = 0$ yields $ka = -\omega_B t_{\mathrm{lab}}$ with Bloch angular frequency $\omega_B = e\mathcal{E}a/\hbar$. Integrating velocity over laboratory time gives

$$x(t_{\mathrm{lab}}) - x(0) = \frac{2t}{e\mathcal{E}}[\cos(\omega_B t_{\mathrm{lab}}) - 1],$$

$$T_B = \frac{h}{e\mathcal{E}a}.$$

Although the field never changes sign, velocity reverses when the packet samples regions of opposite band slope. Over a complete zone traversal there is no net displacement because the band energy is periodic. Energy conservation offers another derivation: the electrostatic potential-energy change of a negatively charged electron in a uniform field is $+e\mathcal{E}\Delta x$, offsetting the change in band energy. The maximal excursion equals the band width $4t$ divided by the force magnitude.

Estimate an observable oscillation in a superlattice

Use a one-dimensional superlattice with $a = 8.00$ nm, hopping $t = 5.00$ meV and field $\mathcal{E} = 2.00 \times 10^5$ V/m. The potential-energy change per lattice period is $e\mathcal{E}a = 1.60$ meV. Thus $T_B = h/(1.60\,\mathrm{meV}) = 2.58$ ps and the peak-to-peak excursion is $20.0\,\mathrm{meV}/(e\mathcal{E}) = 100$ nm, about 12.5 lattice periods. The maximum speed from the band derivative is

$$2ta/\hbar = 1.22 \times 10^5\,\mathrm{m/s}.$$

At one quarter cycle velocity reaches its largest negative magnitude, while at half a cycle the electron reaches its furthest position and momentarily stops.

A momentum-relaxation time of 0.20 ps would interrupt most trajectories before a cycle completes, obscuring a Bloch oscillation despite the periodic dispersion. Conversely a scattering time several times longer than 2.58 ps would make coherent oscillatory transport plausible if the field does not tunnel carriers into other bands. In the weak-field limit the period and excursion both grow, so the mathematical limit of arbitrarily large excursions cannot be taken in a finite imperfect specimen. Recording the oscillation frequency as field is varied tests the predicted linear scaling separately from the much more sample-dependent amplitude.

The motion also clarifies the role of negative effective mass near a band top. The electron's crystal momentum continues to change in the direction of the applied force throughout the period. Its acceleration is proportional to the local band curvature, however, and can point in the opposite direction once that curvature changes sign. This reverses packet acceleration without a sudden elastic collision with a particular lattice plane. A spectrum measured along the direction of motion can predict the velocity through its local slope, giving a direct cross-check of the acceleration inferred from a time-domain trace.

The probe must resolve the oscillations before relaxation for that comparison to be meaningful.

Predicting a gap and testing it in a periodic medium

Condensed matter and band theory

Combine the translation argument, a measurable band edge and an independent numerical calculation to distinguish a gap from a folded crossing.

FROM FIRST PRINCIPLES

Infer pass bands and stop bands from a repeated cell

For a normally incident monochromatic plane wave in a lossless layered dielectric, record its tangential electric and magnetic fields in a two-component state. Interface conditions transmit these same tangential fields between adjacent layers. Propagation through one complete cell of period a is a linear map represented by a matrix $M(\omega)$. Forward and backward propagation in each lossless layer have reciprocal phase factors, so the cell transfer matrix has determinant one. A Bloch state reproduces its field after one cell up to a multiplier:

$$
\begin{aligned}
Mv &= \lambda v, \\
\lambda &= e^{ika}, \\
\lambda^2 - (\operatorname{tr} M)\lambda + 1 &= 0.
\end{aligned}
$$

The second root is $1/\lambda$. Adding the two roots and using the phase relation gives the band condition

$$
\cos(ka) = \tfrac{1}{2}\operatorname{tr} M(\omega).
$$

For a real lossless layered system the trace is real. Where its half lies between –1 and 1 there are real wave numbers. Outside this interval the two multipliers have different magnitudes and one solution decays with each successive cell. In a uniform material the trace simply reproduces homogeneous propagation: folding the dispersion into the first zone can make branches touch without opening a gap. Index contrast allows repeated reflections and can create finite stop bands.

DERIVATION AND APPLICATIONS

Find the long-wavelength optical index of a layered cell

A finite wavelength can resolve two different dielectrics in a repeating stack. Far below the first band edge, what uniform medium would reproduce the phase accumulated per cell? Consider lossless, nonmagnetic layers of thickness d_1, d_2, real refractive indices n_1, n_2, and cell length $a = d_1 + d_2$. At normal incidence the field is tangential to the interfaces. Define $k_0 = \omega/c$ and the phase thickness $\delta_j = k_0 n_j d_j$. For Bloch wavevector K, the exact trace relation is

$$
\cos(Ka) = \cos\delta_1 \cos\delta_2 - \frac{1}{2}\left(\frac{n_1}{n_2} + \frac{n_2}{n_1}\right)\sin\delta_1 \sin\delta_2.
$$

On the acoustic-like branch through zero frequency, all three phase angles are small. Keep the quadratic terms in each cosine and the product of the linear terms in the sines. Constant terms on both sides cancel. Substitution of the optical phases gives

$$(Ka)^2 \simeq k_0^2 \left[n_1^2 d_1^2 + n_2^2 d_2^2 + (n_1^2 + n_2^2) d_1 d_2 \right].$$

The mixed term records the interfaces even at long wavelengths. Factor the bracket as $a(n_1^2 d_1 + n_2^2 d_2)$ and take the forward-propagating root. Writing the first layer's thickness fraction as $f = d_1/a$ yields

$$n_{\mathrm{eff}} = \lim_{\omega \to 0} \frac{K}{k_0} = \sqrt{f n_1^2 + (1-f) n_2^2}.$$

This is the thickness-weighted average of the relative permittivities under a square root. It agrees with a tangential electric field that is nearly uniform across a small cell. As a thickness fraction tends to unity the remaining material's index is recovered. Near a band edge the omitted fourth-order terms in the trace matter, so a constant effective index cannot predict the gap.

WORKED EXAMPLE

Measure the phase of a deeply subwavelength stack

Change the inputs and investigate this derivation.

A cell contains $10.0\,\mathrm{nm}$ of index 2.00 and $30.0\,\mathrm{nm}$ of index 1.50. Its effective index is

$$\sqrt{0.25(4) + 0.75(2.25)} = 1.64.$$

At vacuum wavelength $4.00\,\mu\mathrm{m}$, ten cells accumulate Bloch phase $2\pi n_{\mathrm{eff}}(0.400/4.00) = 1.03$ radians. The larger layer phase per cell is just 0.071 radians, consistent with the expansion.

SECOND APPLICATION

Infer a layer fraction from a low-frequency slope

Change the inputs and investigate this derivation.

Another stack uses indices 3.00 and 1.50. Its long-wavelength measured phase velocity is $c/2.10$. Solving the squared-index average gives

$$f = (2.10^2 - 1.50^2)/(3.00^2 - 1.50^2) = 0.320.$$

A $50.0\,\mathrm{nm}$ cell should contain $16.0\,\mathrm{nm}$ of the high-index material. A microscope can check that inference independently of the optical phase measurement.

Try it · Homogenised stack

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

High-index layer fraction **0.25**

High refractive index **2**

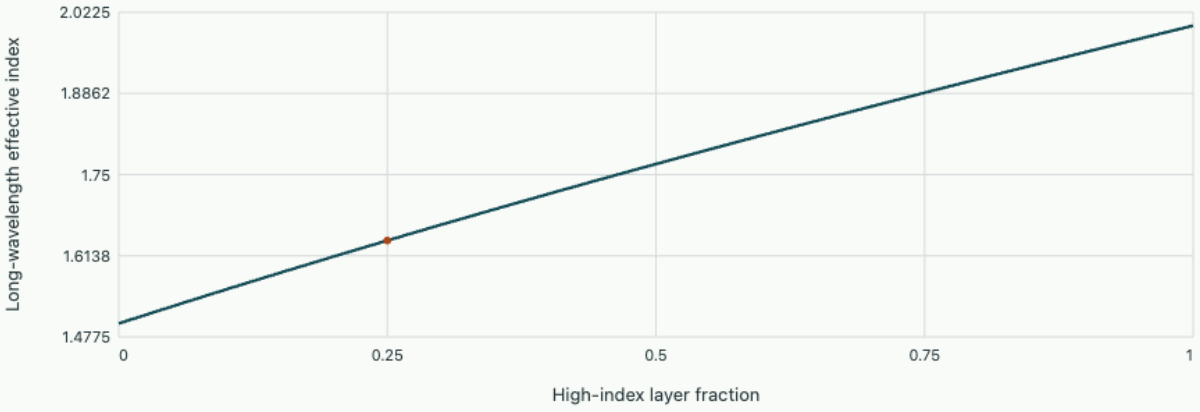

Normal-incidence dielectric stack Selected fraction

Low index is 1.5. The squared indices, rather than indices, average in this long-wavelength transverse-field limit. The cell must be much shorter than the wavelength.

Transfer a field by one period

Consider a lossless stack of two nonmagnetic dielectrics at normal incidence. Layer j has real refractive index n_j, thickness d_j and optical phase $\delta_j = 2\pi f n_j d_j / c$. Within a layer the electric field is a sum of forward and backward waves; the corresponding magnetic field has the opposite sign for the backward wave. Eliminating the wave amplitudes between the two faces gives the layer transfer matrix for electric field and magnetic field scaled by the vacuum impedance (one time-phase convention):

$$M_j = \begin{pmatrix} \cos\delta_j & i\sin\delta_j/n_j \\ i n_j \sin\delta_j & \cos\delta_j \end{pmatrix}, \qquad \det M_j = 1.$$

The tangential fields are continuous at a simple dielectric interface. Thus the matrix of one complete cell of size $a = d_1 + d_2$ is the product of the two layer matrices. Its eigenvalues multiply to one. For propagation in an infinite periodic stack the Bloch multipliers must be $e^{\pm ika}$. Their sum fixes the half-trace:

$$\cos(ka) = \cos\delta_1\cos\delta_2 - \tfrac{1}{2}(n_1/n_2 + n_2/n_1)\sin\delta_1\sin\delta_2.$$

If the right side lies outside $[-1, 1]$, no real k exists. One solution decays and the reciprocal one grows along the chosen direction. A semi-infinite passive medium selects the decaying solution inside a stop band. A finite sample can still transmit an evanescent field, especially when it contains few cells.

Derive a check that the numerical solver could fail

Set $n_1 = n_2 = n$. The right side reduces to $\cos(\delta_1 + \delta_2)$ by the addition formula. The first two folded branches meet where the combined phase is π, at $fa/c = 1/(2n)$. There is no interface reflection and no forbidden interval. With $n = \sqrt{13}$ both branches should meet at about 0.138675. Two nearly equal eigenvalues on a finite mesh must be compared with the numerical resolution before reporting a narrow gap.

As another analytic limit, take each layer a quarter wavelength thick at the same frequency, so

$$\delta_1 = \delta_2 = \pi/2.$$

The half-trace is $-\frac{1}{2}(r + r^{-1})$, $r = n_1/n_2$. Its magnitude exceeds one whenever $r \neq 1$. For $r = 2$ the decaying multiplier has magnitude $1/2$ per cell. The attenuation here arises from coherent reflections in a lossless model. Small transmission need not mean that electromagnetic energy was dissipated.

Read the finite-element eigenvalues against the calculation

In the OneSimulation band experiment, keep the layer fractions fixed and change only the second permittivity. For the pair (13,12) the two reported zone-edge frequencies are 0.1396 and 0.1433 in reduced units; the gap is about 2.6% of their mean. Compare this separation with the zero-contrast result and with the two roots of the half-trace equation at -1. For the much larger contrast (13,1), using a weak-perturbation approximation would be hard to justify. The transfer matrix still includes repeated scattering exactly for these ideal layers.

When investigating an unfamiliar contrast, halve the spatial mesh size and repeat the edge calculation. Track modes by their fields as well as by sorted eigenvalue if two branches nearly meet. For the scalar one-dimensional normal-incidence problem the observed stop band has a precise meaning. It is an interval of optical frequencies for that propagation geometry. Inferring a complete three-dimensional gap would require other directions and polarisations. The link to electronic bands is the shared translation eigenvalue and coherent coupling at a crossing; electronic filling and the optical transfer-matrix test answer different physical questions.

Turn an infinite stop band into finite-stack reflectance

The forbidden Bloch interval of an infinite periodic medium does not imply perfect reflection from a finite sample. Its termination and number of periods matter. Consider normal incidence on alternating lossless dielectric layers with refractive indices n_H, n_L, each of optical thickness $\lambda_0/4$ at the selected vacuum wavelength. Treat all permeabilities as μ_0 and use normalized optical admittance $Y = n$. Across a quarter-wave layer the characteristic matrix mapping the fields on its far face to its entrance has zero diagonal and off-diagonal entries i/Y and iY, with an overall phase convention fixed consistently.

For a terminating medium of admittance Y_{load} the resulting input admittance is $Y_{\mathrm{layer}}^2/Y_{\mathrm{load}}$. Applying this inversion first to the rear low-index layer and then to the front high-index layer multiplies load admittance by

$(n_H/n_L)^2$. After N identical pairs, with the high-index layer first as seen by incident light,

$$Y_{\mathrm{in}} = n_s (n_H/n_L)^{2N},$$

$$R = \left| \frac{n_0 - Y_{\mathrm{in}}}{n_0 + Y_{\mathrm{in}}} \right|^2.$$

Here n_s is the substrate index and n_0 the incident medium index. This design-frequency result follows from phase accumulation in each layer. For a general wavelength the matrices are no longer simple inversions and both diagonal entries must be retained. Interchanging the order of high and low layers in a finite stack also changes the input admittance and reflection phase.

Select a period count for a required mirror

Let $n_0 = n_s = 1.00$, $n_H = 2.00$ and $n_L = 1.50$ at $\lambda_0 = 600$ nm. Quarter-wave physical thicknesses are

$$600/(4 \times 2.00) = 75.0\,\mathrm{nm},$$

$$600/(4 \times 1.50) = 100\,\mathrm{nm}.$$

Four pairs give

$$Y_{\mathrm{in}} = (4/3)^8 = 9.99,$$

$$R = [(9.99 - 1)/(9.99 + 1)]^2 = 0.669.$$

Eight pairs give $Y_{\mathrm{in}} \simeq 99.8$ and $R = 0.961$. Thus a fairly clear infinite-medium stop band can coexist with substantial transmission through a short specimen.

For at least 99.0 percent reflectance at the design wavelength the required admittance ratio is $(1 + \sqrt{0.990})/(1 - \sqrt{0.990}) \simeq 398$. Solving $(4/3)^{2N} \geq 398$ gives $N \geq 10.4$, hence eleven complete pairs. If the indices become equal and match both external media, $Y_{\mathrm{in}} = n_0$ for any period count and the reflection vanishes. This finite-stack result tests a different observable from an infinite cell eigenfrequency. It connects the gap calculation to a transmission or reflectance measurement while exposing the additional data, particularly termination and thickness, needed for that prediction.

At the chosen wavelength the reflection amplitude for high-index-first ordering is negative because the effective admittance exceeds the incident admittance. Measuring reflected intensity alone loses this phase sign. Interference with a calibrated reference reflector can recover the sign and identify termination differences even when two samples have similar reflectances. This matters when the periodic mirror forms one end of a resonant cavity: the reflection phase contributes to the total round-trip phase and can move the resonance frequency. A pair of finite mirrors can therefore share a stop-band range while giving different cavity phase conditions.

Derive the quarter-wave stop-band width

The finite mirror is easy to evaluate at the design wavelength. The same stack admits an analytic expression for the width of its infinite first stop band. In both layers the optical phase thickness is $\delta = (\pi/2)(\omega/\omega_0)$ when refractive indices are treated as nondispersive. The one-period transfer matrix has unit determinant and

half-trace $\cos(Ka)$, where K is the Bloch wavenumber and a the physical period. Multiplication of the two layer matrices gives

$$\cos(Ka) = \cos^2\delta - \frac{1}{2}\left(\frac{n_H}{n_L} + \frac{n_L}{n_H}\right)\sin^2\delta.$$

At band centre the half-trace is less than minus one for unequal indices, so K has a nonzero imaginary component. At either neighbouring edge set the half-trace equal to minus one and solve for δ. The two edge phases are $\pi/2 \pm \arcsin[(n_H - n_L)/(n_H + n_L)]$ for $n_H > n_L$. Their separation gives fractional angular-frequency width

$$\frac{\Delta\omega}{\omega_0} = \frac{4}{\pi}\arcsin\left(\frac{n_H - n_L}{n_H + n_L}\right).$$

With indices 2.00 and 1.50 this is 0.183. A central frequency of 500 THz thus has edges near 454 and 546 THz. The frequency interval is symmetric about the centre in this ideal quarter-wave construction; the corresponding wavelength interval is asymmetric because wavelength is reciprocal to frequency. As the two indices merge, the edge separation vanishes even though an artificially repeated unit cell still has a Brillouin-zone boundary.

At centre the decaying Bloch amplitude falls by $n_L/n_H = 0.750$ per period. A physical period of 175 nm therefore gives a field-amplitude decay length $175/\ln(2/1.5) = 608$ nm. This infinite-medium length explains the strong dependence of the finite mirror on period count. The external-medium matching factors are still needed to turn that internal decay into the actual reflected intensity.

How alternating bonds split a one-dimensional band

Condensed matter and band theory

Two periodic chains contain precisely the same number of atoms at the same spacing. One has equal bonds and the other alternates strong and weak bonds. The difference changes which electron energies can propagate, even without changing the on-site energy of any atom.

Choose a primitive cell that repeats the bonds

Let a cell of length a contain orbitals A and B. An electron hops between A and B in the same cell with real amplitude $t_1 > 0$, and between B and A in the next cell with amplitude $t_2 > 0$. Use one orbital per site, neglect interactions and set both on-site energies to zero. We can absorb an overall hopping sign into the B basis state. Translation by one cell multiplies each component by e^{ika}, although the relative amplitudes on A and B remain unknown. Writing those amplitudes as A, B turns the real-space nearest-neighbour equations into

$$EA = (t_1 + t_2 e^{-ika})B,$$

$$EB = (t_1 + t_2 e^{ika})A.$$

The cell boundary crossed by the second hopping accounts for its phase. The two coefficients are complex conjugates as required by Hermiticity. Substitute the second relation into the first without dividing by energy. A nonzero component then requires

$$E^2 = (t_1 + t_2 e^{ika})(t_1 + t_2 e^{-ika}).$$

The two energies at each wavevector are

$$E_{\pm}(k) = \pm\sqrt{t_1^2 + t_2^2 + 2t_1 t_2 \cos ka}.$$

For positive hoppings their magnitude is largest at $k = 0$, namely $t_1 + t_2$, and smallest at $ka = \pi$, namely $|t_1 - t_2|$. The gap is $2|t_1 - t_2|$. At one spinless electron per cell the lower band is full, so this is also the independent-particle excitation gap in a large periodic chain.

Read the band formation from two limits

Take $t_1 = 1.0\,\mathrm{eV}$, $t_2 = 0.6\,\mathrm{eV}$. The lower band runs from -1.6 to $-0.4\,\mathrm{eV}$, the upper from 0.4 to $1.6\,\mathrm{eV}$. The gap is $0.8\,\mathrm{eV}$, while each bandwidth is $1.2\,\mathrm{eV}$. If t_2 tends to zero, the cell becomes an isolated pair with two discrete bonding energies $\pm t_1$. Bands flatten because cells can no longer exchange probability. At $t_1 = t_2 = t$, the spectrum becomes $\pm 2t|\cos(ka/2)|$ in the reduced zone. The two pieces meet at its edge. The two-site construction then folds the simple uniform chain into a smaller wavevector interval.

One can also read off a velocity. Away from a crossing, differentiate $E_{\pm}$ to obtain

$$v_{\pm} = \mp a t_1 t_2 \sin(ka)/[\hbar |E_{\pm}|].$$

At extrema of a gapped band it vanishes. At the equal-hopping crossing the derivative written as a ratio has a removable directional ambiguity, because the ordered upper and lower bands meet with opposite slopes.

Exchange the two bonds

Keep t_1 fixed and move t_2 through it. Compare both the gap and the bandwidth. Interchanging the numerical values of the two hoppings exactly preserves the bulk energies. The experiment in the next chapter distinguishes these patterns by opening the chain at a specified bond. A termination supplies physical information that the unordered hopping pair and the energy graph alone cannot encode.

Open this experiment online

Try the calculation

Intracell hopping t_1 (eV) **1**

Intercell hopping t_2 (eV) **0.6**

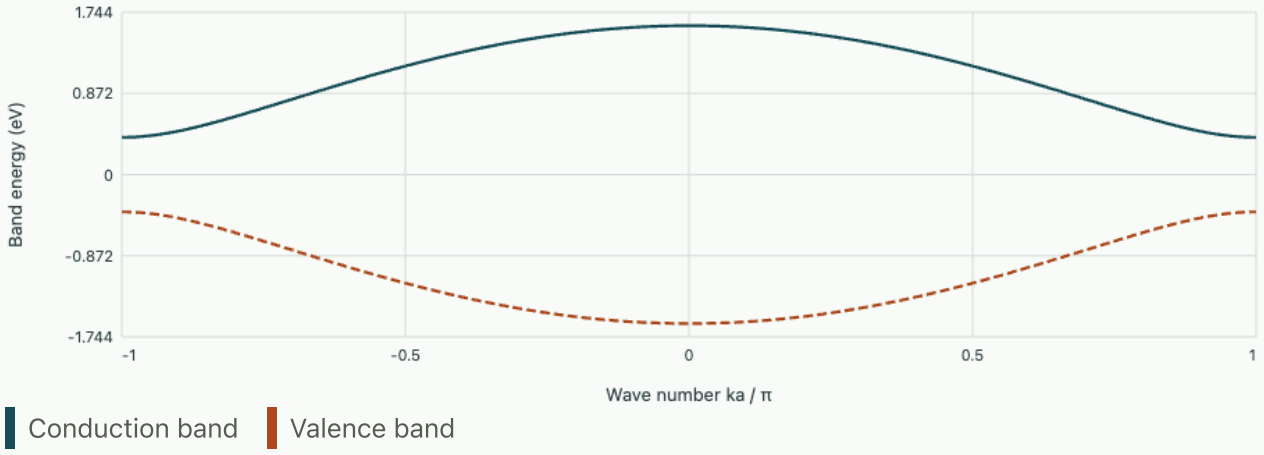

Direct gap 0.8 eV. Each band spans 1.2 eV.

Tutorial on inferring the hopping strengths

A measured dimer band has gap $0.50\,\mathrm{eV}$ and reaches $1.45\,\mathrm{eV}$ above the symmetry centre at $k = 0$. Find the two hopping amplitudes and predict the upper-band width.

▼ Work through the solution

The sum is 1.45 eV and the absolute difference is 0.25 eV. Adding and subtracting gives the unordered pair 0.85 and 0.60 eV. The upper bandwidth is

$$1.45 - 0.25 = 1.20\,\mathrm{eV}.$$

Bulk energies do not determine which hopping is interrupted by a particular surface cut.

Expand the squared band energy near the Brillouin-zone boundary for unequal hoppings. Identify the small-q dispersion and its range.

▼ Work through the solution

Set $ka = \pi + qa$. Since $\cos(\pi + qa) \simeq -1 + (qa)^2/2$,

$$E^2 \simeq (t_1 - t_2)^2 + t_1 t_2 (qa)^2.$$

For sufficiently small q the upper edge has quadratic curvature $E_+ \simeq |t_1 - t_2| + t_1 t_2 (qa)^2 / [2|t_1 - t_2|]$. The last expansion fails when its second term in the squared energy competes with the gap term.

If the same constant on-site energy U is added to every site, which inferred quantities would move in an experiment measuring energies against a fixed reference?

▼ Work through the solution

The Hamiltonian gains UI, moving both bands and their centre by U. Their gap, widths, group velocities, and the hopping estimates from differences remain unchanged. The half-filling chemical potential can still lie inside the shifted gap.

Further reading · Asbóth, Oroszlány and Pályi, SSH model and bulk bands.

Finding an end state one bond at a time

Condensed matter and band theory

The bulk spectrum of an alternating chain is unchanged when strong and weak bonds exchange places. A boundary with a named first site breaks this equivalence. The zero-energy equations provide its wavefunction without diagonalising a large matrix.

Start at an exposed A site

Use the alternating chain with intracell link t_1 and intercell link t_2, both positive. End the left edge at A_1. The first B site connects A_1 and A_2. Seek an eigenstate at $E = 0$ supported only on A sites. Each A-site equation then holds because all neighbouring B amplitudes vanish. Each B-site equation requires cancellation between its two neighbouring A amplitudes. Hence

$$0 = t_1 A_j + t_2 A_{j+1},$$

$$A_j = (-t_1/t_2)^{j-1} A_1.$$

Relative signs matter in the cancellation, even though a local probability measurement sees only squared magnitudes. On a semi-infinite chain the geometric series for the norm converges only for $r = t_1/t_2 < 1$. Normalising gives

$$|A_1|^2 = 1 - r^2.$$

The amplitude decay length satisfies $e^{-a/\ell} = r$, so $\ell = a/\ln(t_2/t_1)$. The probability decay length is half that. As the bulk gap closes, the decay length diverges and the bound state spreads into the interior.

Impose the far boundary too

A finite chain can have M A sites and $M - 1$ B sites, ending on A at both edges. Every B site still imposes exactly the recurrence above, leaving one overall free amplitude. The resulting state has exactly zero energy for any positive hoppings. Its finite norm is the sum of M successive powers of r^2. At $r = 1$ the apparent zero over zero in the closed formula is resolved by summing M unit terms. The state then has probability $1/M$ at every A site.

When $r > 1$ its norm remains finite in this finite sample and the weight accumulates at the right edge. An equal-sublattice finite chain ending on B instead has an additional boundary equation. Two exponentially localised end solutions can then hybridise, shifting their energies slightly above and below zero. Equal sublattice counts alone provide no exact zero eigenvalue. A change of boundary conditions can therefore alter the interpretation of a measured near-zero peak.

Connect the recurrence to the bulk phase

The off-diagonal matrix element $h = t_1 + t_2 e^{ika}$ traces a circle centred at t_1 with radius t_2 as k crosses the Brillouin zone. If $t_2 > t_1$ it winds once around zero. If $t_2 < t_1$ it misses zero. Passing between the two with fixed cell choice requires the gap $2|h|$ to close. A diagonal potential on just one sublattice would break the off-diagonal symmetry used to pin energies to zero. The winding statement and the boundary argument require a consistent assignment of the unit cell and termination.

For $r = 0.5$ the first three semi-infinite A probabilities are 0.75, 0.1875 and 0.046875. For $r = 0.90$, the first site has only 0.19 and the amplitude decay length increases from $1.44a$ to $9.49a$. Use the finite odd chain below to move the density from one end to the other. At reciprocal hopping ratios the two distributions must be mirror images. Check this using the endpoint readouts.

Open this experiment online

Try the calculation

Hopping ratio t_1/t_2 **0.5**

Number of A sites **12**

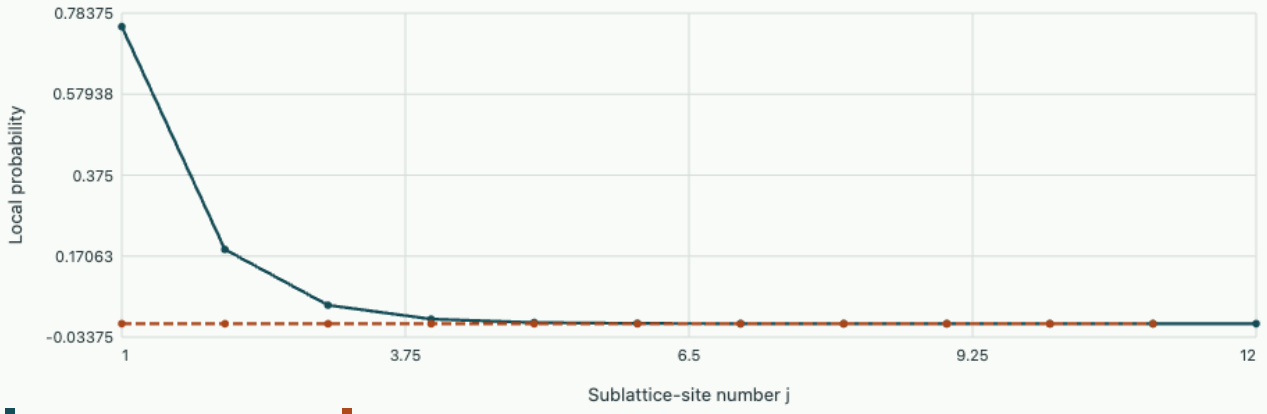

Probabilities at the first and last A sites are 0.75 and 1.788e-7. The amplitude decay length from the favoured edge is 1.443 cell spacings.

Tutorial on inspecting a local measurement

Four successive A sites have relative intensities $1, 0.36, 0.1296, 0.0467$ near a left edge. Infer the hopping ratio and amplitude decay length. Would a local intensity probe detect the alternating signs of the state?

▼ **Work through the solution**

Ratios of neighbouring probabilities equal r^2, hence $r = 0.60$. The amplitude length is

$$-a/\ln(0.60) = 1.96a.$$

A measurement proportional to the absolute square loses the sign. Interference with a phase reference or a measurement sensitive to the amplitude across a bond can access the phase relation.

A finite odd chain has five A sites and hopping ratio $r = 0.8$. What fraction of zero-mode probability lies on its first two A sites?

▼ **Work through the solution**

Sum the finite geometric sequence. The requested fraction is

$$(1+r^2)/\sum_{j=0}^{4} r^{2j} = (1-r^4)/(1-r^{10}) = 0.661.$$

Taking the semi-infinite formula would give

$$1-r^4 = 0.590.$$

The far boundary changes the normalisation substantially for this short array.

Can bond disorder that preserves positive nonzero nearest-neighbour hoppings destroy the exact zero solution of a finite chain ending on A at both ends?

▼ **Work through the solution**

At each B site retain the local recursion $A_{j+1} = -(t_{1j}/t_{2j})A_j$. It produces a finite nonzero solution after choosing A_1, with all B amplitudes zero. The profile is a product of local ratios and may cease to look like a single exponential. A generic A-site potential would invalidate the zero-energy equation.

Further reading · Asbóth, Oroszlány and Pályi, edge states and bulk-boundary correspondence.

Why a honeycomb lattice has Dirac cones

Condensed matter and band theory

In an alternating chain the two paths from A to B can cancel at a special wavevector. The honeycomb lattice adds a third path. Their cancellation now occurs at isolated points in a two-dimensional Brillouin zone and produces a linear low-energy dispersion.

Add the three Bloch phases

Each honeycomb A site has three nearest B neighbours at displacements $d(0,1)$, $d(\sqrt{3}/2,-1/2)$, and $d(-\sqrt{3}/2,-1/2)$, where d is the bond length. For equal real hopping magnitude t, sum the translation phase for each displacement. The common off-diagonal factor is

$$f(\boldsymbol{k}) = e^{ik_y d} + 2e^{-ik_y d/2}\cos(\sqrt{3}k_x d/2),$$
$$E_\pm = \pm t|f|.$$

This uses a basis phase convention that includes the orbitals' intra-cell positions. A rephasing can change the phase of f but leaves its magnitude invariant. Along $k_y = 0$, cancellation occurs where $1 + 2\cos(\sqrt{3}k_x d/2) = 0$. Choose $K = (4\pi/(3\sqrt{3}d), 0)$. The counterpart at $-K$ is inequivalent up to primitive reciprocal vectors. Near either valley many wavevectors represent distinct states even when their energies almost coincide.

Linearise at a zero, then test the approximation

Write $\boldsymbol{k} = K + \boldsymbol{q}$ with $qd \ll 1$. At K the cosine equals $-1/2$ and the sine $\sqrt{3}/2$. Differentiating the cosine term with respect to k_x consequently gives $-3d/2$. Differentiating the two exponential factors with respect to k_y gives

$$id + id/2 = 3id/2.$$

The constant terms cancel. Therefore

$$f(K+\boldsymbol{q}) \simeq \tfrac{3d}{2}(-q_x + iq_y),$$
$$E_\pm \simeq \pm\hbar v_F|\boldsymbol{q}|,$$
$$v_F = \frac{3td}{2\hbar}.$$

With $t = 2.7\,\mathrm{eV}$, $d = 0.142\,\mathrm{nm}$ and $\hbar = 0.6582\,\mathrm{eV\,fs}$, the speed is

$$0.874\,\mathrm{nm/fs} = 8.74\times10^5\,\mathrm{m/s}.$$

At $qd = 0.05$ the leading energy magnitude is $0.2025\,\mathrm{eV}$. The energy varies linearly with momentum displacement even though the microscopic Hamiltonian contains only hopping between neighbours. The two components describe sublattice amplitude, distinct from the real electron spin.

Break the equivalence of the sublattices

Assign on-site energies $+\Delta$ to A and $-\Delta$ to B. The characteristic equation of the two-by-two Hamiltonian becomes

$$E^2 = \Delta^2 + t^2|f|^2.$$

A gap $2|\Delta|$ opens at K. For $\Delta = 0.10\,\mathrm{eV}$, the edge sits at $0.10\,\mathrm{eV}$ and the predicted conduction energy at $qd = 0.05$ becomes about $0.226\,\mathrm{eV}$. Far enough from the origin the slopes recover their nearly linear values. Close to it, expansion of the square root gives

$$E_+ \simeq |\Delta| + \hbar^2 v_F^2 q^2/(2|\Delta|).$$

A band-edge mass $|\Delta|/v_F^2$ has emerged. The expansion requires $\hbar v_F q \ll |\Delta|$.

Compare the exact positive-energy cut along q_x with the local cone in the graph. Unequal deviations to the right and left of K are allowed by the lattice and cannot be captured by the lowest-order isotropic cone. The reflected cut at the other valley restores the time-reversal relation.

Open this experiment online

Try the calculation

Sublattice energy offset Δ (eV) **0.1**

Selected displacement $q_x d$ **0.05**

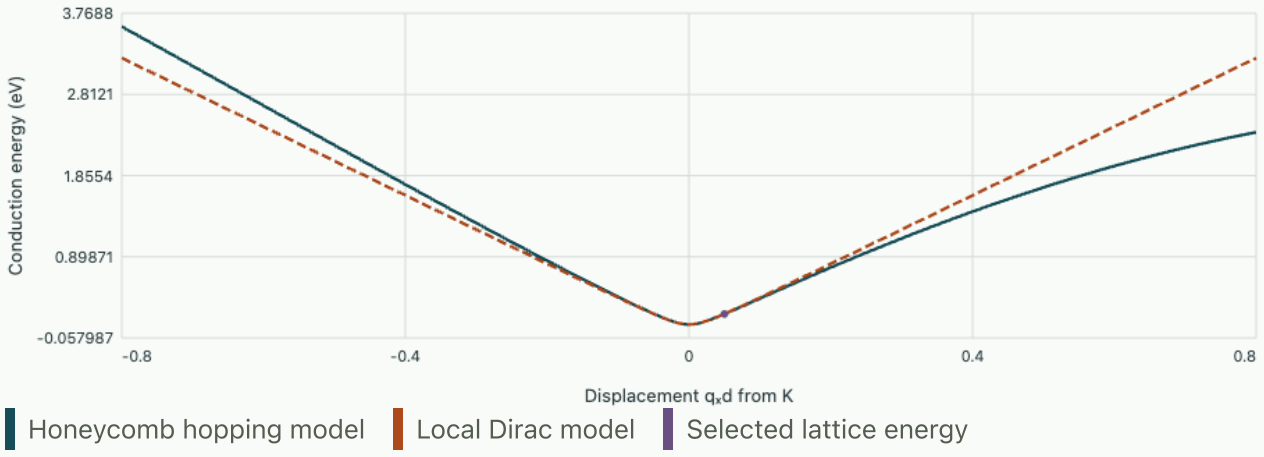

Direct gap 0.2 eV. At the selected displacement the lattice gives 0.2235 eV and the local model 0.2258 eV.

Tutorial on a measured slope

A near-valley spectrum has slope $0.55\,\mathrm{eV\,nm}$. Infer v_F and the nearest-neighbour hopping for a bond length $d = 0.142\,\mathrm{nm}$. A gap of $80\,\mathrm{meV}$ then appears with the same asymptotic slope. What is $|\Delta|$?

▼ **Work through the solution**

The slope equals $\hbar v_F$. Dividing by $0.6582\,\mathrm{eV\,fs}$ gives $0.836\,\mathrm{nm/fs}$. Since the slope also equals $3td/2$,

$$t = 2(0.55)/(3 \times 0.142) = 2.58\,\mathrm{eV}.$$

The two edges shift symmetrically by half the full gap, giving $|\Delta| = 40\,\mathrm{meV}$.

Consider a gapless honeycomb cone exactly at the valley centre. What does the tight-binding Hamiltonian predict for the energy and eigenvectors, and what information comes from approaching along a direction?

▼ **Work through the solution**

At $f = 0$ the matrix vanishes, so every two-component vector in the degenerate subspace has energy zero. For a nearby nonzero f, a chosen eigenbranch fixes a relative phase between its A and B components according to the argument of f. The limiting eigenvector therefore depends on the direction of approach even though the limiting energy is always zero.

At $q = 0$, repeat the eigenvector argument with $\Delta > 0$. Predict a consequence for sublattice-resolved spectroscopy at the two band edges.

▼ **Work through the solution**

At the valley centre the Hamiltonian is diagonal with entries $+\Delta$, $-\Delta$. The upper state lives on A and the lower on B in this basis. A probe selective to A should see more weight at the conduction edge than at the valence edge. Farther from the valley centre the hopping mixes the sublattices again.

Further reading · Delft Topology in Condensed Matter, honeycomb Hamiltonian and Dirac points.

When a periodic dielectric has no propagating Bloch wave

Fields and the response of matter

Derive a band gap using two lossless transparent materials and distinguish evanescent Bloch propagation from absorption of optical power.

Construct a unit cell whose waves interfere

Stack two transparent, lossless, nonmagnetic dielectrics periodically in one dimension. Their refractive indices are n_1, n_2, their physical thicknesses d_1, d_2, and the period $a = d_1 + d_2$. Work at normal incidence, so polarisation directions are equivalent for isotropic layers. Each separate material admits a propagating wave at the chosen optical frequency, yet the periodic stack may not. The effect comes from coherent interface reflections, not from a negative refractive index or an absorptive constituent.

Describe the local field by its tangential electric and magnetic components. Propagation through one homogeneous layer connects them by a 2×2 transfer matrix with diagonal $\cos\phi_j$, off-diagonal terms proportional to $iZ_j \sin\phi_j$ and $i\sin\phi_j/Z_j$, where $\phi_j = \omega n_j d_j/c$ and $Z_j = Z_0/n_j$. Changing the chosen propagation direction or time convention may reverse the two imaginary signs together. The real trace and unit determinant of the cell matrix below are unaffected.

Read the dispersion relation from the trace

At an interface both tangential components are continuous in the absence of imposed sheet currents. Multiply the two propagation matrices to form one cell. Its determinant is one because each layer's determinant is

$$\cos^2\phi + \sin^2\phi = 1.$$

Thus the two eigenvalues are reciprocal. For a Bloch mode they are $e^{\pm ik_B a}$ on a pass band. Taking half the trace gives

$$\cos(k_B a) = \cos\phi_1\cos\phi_2 - \frac{1}{2}\left(\frac{n_1}{n_2} + \frac{n_2}{n_1}\right)\sin\phi_1\sin\phi_2.$$

When the right side lies between –1 and 1 a real Bloch phase exists. If its magnitude exceeds one, a decaying and a growing Bloch factor replace a pair of unit-modulus factors. For a bounded excitation of a semi-infinite stack choose the decaying direction. The attenuation per cell is $\operatorname{arcosh}\left|\operatorname{trace}(M)/2\right|$. A lossless semi-infinite crystal in a stop band reflects incident power. Spatial attenuation of a Bloch amplitude in this region is compatible with zero volumetric heat production.

Obtain a quantitative stop-band width

Choose $n_1 d_1 = n_2 d_2 = \lambda_0/4$. Writing $s = \omega/\omega_0$, both layer phases are $\phi = \pi s/2$. At $s = 1$ the half trace is $-\frac{1}{2}(n_2/n_1 + n_1/n_2)$. Recognising $\cosh(\log r) = (r + r^{-1})/2$ gives centre attenuation $|\log(n_2/n_1)|$ per cell. For $n_1 = 1.5, n_2 = 2.0$ this is 0.2877, so a decaying eigencomponent has amplitude factor 0.75 per cell. This statement about an eigencomponent should not be confused with the exact transmittance of a specific finite stack, which also depends on entrance and exit impedances.

The first gap edges obey half trace −1. With $\cos^2\phi = 1 - \sin^2\phi$, solve for the edge sine or cosine. The two edge phases are symmetric about $\pi/2$ and satisfy

$$|\cos\phi| = |n_2 - n_1|/(n_1 + n_2).$$

Hence the fractional angular-frequency width is $(4/\pi)\arcsin(|n_2 - n_1|/(n_1 + n_2))$, about 0.183 for the example. The lowest gap width tends continuously to zero as index contrast vanishes. A gap in a plotted numerical spectrum at $n_1 = n_2$ would indicate an implementation or plotting problem.

Compare a frequency scan with a finite mirror

Open the live experiment

Try the calculation

High index n_2 (n_1 = 1.5) 2

Probe frequency / design frequency 1

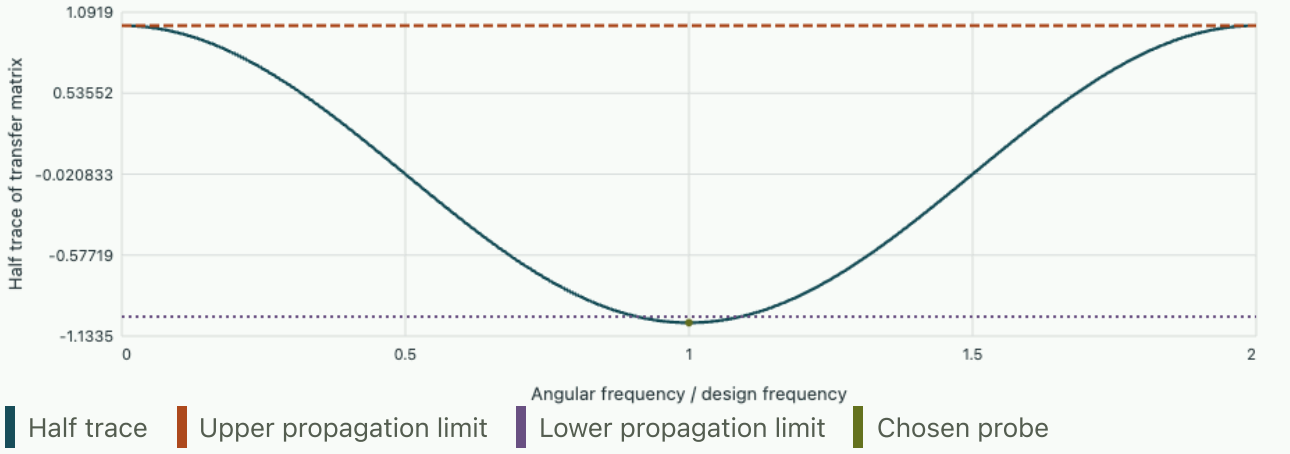

First gap width / design frequency 0.18252. Probe half trace -1.0417. Attenuation per cell 0.28768.

The graph plots half the transfer-matrix trace against normalised frequency and marks the admissible interval for a real Bloch phase. Change the index contrast first, keeping the layers quarter wave at the same reference frequency. The physical layer thickness adjusts in this family; it is a different experiment from changing a material while manufacturing all thicknesses fixed. Use the frequency control to read an attenuation exponent at the centre and just outside the gap.

At low frequency a wavelength spans many cells and a homogenised description can predict the long-wave speed. At the first gap the field changes phase by order one across a cell, exactly where the homogenisation scale separation fails. Conversely, an electromagnetic calculation showing multiple bands at very low frequency should be checked against a uniform-medium or long-wavelength effective-index limit. Compare a finite stack's reflection with the Bragg-mirror experiment. The bulk trace determines whether a propagating state can exist. A finite sample adds interference between boundaries and possible surface or defect effects.

Tutorial

If the two refractive indices become equal, simplify the trace formula for arbitrary thicknesses.

▼ **Work through the solution**

The impedance factor becomes one, so the half trace is $\cos\phi_1\cos\phi_2 - \sin\phi_1\sin\phi_2 = \cos(\phi_1+\phi_2)$. It never exceeds unit magnitude. Labelling a uniform medium by an artificial period folds the dispersion into zones but does not create a physical gap.

At the central quarter-wave frequency for index ratio 2, how many periods reduce a decaying Bloch component's amplitude by a factor of 100?

▼ **Work through the solution**

The amplitude factor per cell is

$$e^{-\log 2} = 1/2.$$

Require $2^{-N} \le 0.01$, giving $N \ge 7$. A prediction for transmitted optical power also needs termination and impedance matching.

A measured stack exhibits both a transmission minimum and measurable heating. Can the minimum be attributed entirely to a band gap?

▼ **Work through the solution**

A lossless band gap can reflect power without heating. Measure reflected and transmitted flux under consistent normalisation, then infer absorption from $1 - R - T$ after accounting for diffuse scattering and other exit channels. The bulk stop-band model and material absorption can coexist.

Compare with band calculations

MIT's mathematical nanophotonics course computes full periodic Maxwell modes. The one-dimensional trace above is a small benchmark with known pass and stop boundaries that a numerical eigenproblem should reproduce.

Action density, stress waves and the boundary that reflects them

Analytical mechanics from constraints to invariants

A chain of coupled masses becomes an elastic field when the wavelength is long compared with the spacing. Derive the field equation and its boundary traction from the same variation, follow energy and stress across a junction, and decide whether a measured pulse is limited by material waves or by numerical dispersion.

Pass from a discrete rod to an action density

Let equal masses of spacing a move longitudinally by $u_j(t)$, with nearest-neighbour spring constant k_s and mass m_s per site. Their Lagrangian is $\sum_j [m_s \dot{u}_j^2/2 - k_s (u_{j+1} - u_j)^2/2]$. For variations smooth on the spacing, $(u_{j+1} - u_j)/a \to u_x$ and $\sum_j a\, f_j \to \int f(x)\, dx$. Hold the line density $\rho A = m_s/a$ and axial rigidity $EA = k_s a$ fixed. The continuum Lagrangian density per unit reference length is

$$\mathcal{L} = \rho A\, u_t^2/2 - EA\, u_x^2/2,$$

with A area, ρ volume density and E Young modulus. Both terms have units of joules per metre. This is a slender-rod model with small strain, primarily uniaxial stress and wavelengths long compared with the cross-sectional dimensions.

Vary the field in the action over $[0, L] \times [t_1, t_2]$. Integration of $\rho A\, u_t \delta u_t$ by parts in time, with fixed endpoint fields, produces $-\rho A\, u_{tt} \delta u$. Integration of $-EA u_x \delta u_x$ by parts in space produces $+\partial_x (EA u_x) \delta u$ and a remaining term $-[EA u_x \delta u]_0^L$. Thus the interior equation is $\rho A u_{tt} = \partial_x (EA u_x)$, and the end force conjugate to displacement is axial force $N = EA u_x$. Adding external work $F_L(t) \delta u(L)$ at the right face yields

$$N(L) = F_L.$$

A prescribed displacement instead imposes $\delta u(L) = 0$, leaving its reaction as an output. Keeping this boundary term prevents prescribing displacement and independent traction simultaneously at the same end of a second-order rod.

Determine the energy flux and characteristic impedance

Assume time-independent coefficients. Multiply the field equation by u_t and rearrange to obtain

$$\partial_t [\rho A u_t^2/2 + EA u_x^2/2] + \partial_x [-EA u_x u_t] = 0.$$

The energy flux is $-N u_t$, with positive sign denoting energy moving towards increasing x. The minus sign follows the convention of positive tensile N: positive mechanical power entering through a right face is $N u_t$, which is inward with respect to the rod on its left. This identity controls energy transport across both smooth material changes and sharp joints.

For a homogeneous right-moving wave $u = F(x - ct)$, $c = \sqrt{E/\rho}$, particle velocity $v = u_t = -cF'$ and axial force $N = EAF' = -Zv$, where $Z = \rho A c = EA/c$. A left-moving wave has $N = +Zv$. The impedance is force divided by particle velocity and has units $\mathrm{kg/s}$. If one instead uses stress, impedance is ρc per unit area. Doubling the area of the same material keeps wave speed unchanged but doubles the force–

velocity impedance. A transit-time measurement and a reflection measurement therefore provide independent material or geometry information.

For a rod with $E = 70\,\mathrm{GPa}$, $\rho = 2700\,\mathrm{kg/m^3}$ and $A = 1.0 \times 10^{-4}\,\mathrm{m^2}$, $c \approx 5092\,\mathrm{m/s}$ and $Z \approx 1375\,\mathrm{kg/s}$. A right-going compressive pulse with particle speed $0.020\,\mathrm{m/s}$ carries axial force $-27.5\,\mathrm{N}$, compressive stress $0.275\,\mathrm{MPa}$, and instantaneous power $Zv^2 = 0.550\,\mathrm{W}$ across a section. A rectangular approximation lasting $100\,\mu\mathrm{s}$ carries about $55\,\mu\mathrm{J}$ and occupies $0.509\,\mathrm{m}$. The pulse speed is kilometres per second; the material particles move only millimetres per second. Treating these as the same velocity would spoil force, impulse and energy estimates.

Match a bonded interface and distinguish amplitude from power

At a perfect massless bond displacement and axial force are continuous. For an incident right-going velocity amplitude v_i in rod 1, let v_r be reflected left-going velocity and v_t transmitted right-going velocity. Velocity continuity gives

$$v_i + v_r = v_t.$$

Force continuity gives

$$-Z_1 v_i + Z_1 v_r = -Z_2 v_t.$$

Thus

$$r_v = \frac{Z_1 - Z_2}{Z_1 + Z_2},$$

$$t_v = \frac{2Z_1}{Z_1 + Z_2},$$

$$\mathcal{R} = r_v^2,$$

$$\mathcal{T} = \frac{Z_2}{Z_1} t_v^2 = \frac{4Z_1 Z_2}{(Z_1 + Z_2)^2}.$$

The power fractions add to one for real positive impedances. The reflected force coefficient has opposite sign to the reflected velocity coefficient because the wave direction reverses. At a free end $Z_2 \to 0$, reflected velocity equals incident velocity, so end velocity doubles while total force vanishes. At a rigid end $Z_2 \to \infty$, velocity reverses and force doubles. The sign a strain gauge reports can therefore differ from that of a velocity-sensitive optical measurement at the same reflection.

A step from the worked rod into one of the same material with twice the cross-sectional area gives $Z_2 = 2Z_1$, $r_v = -1/3$, $t_v = 2/3$, and power fractions $1/9$ and $8/9$. A reduced transmitted velocity amplitude does not imply that most energy is reflected. If the interface includes a concentrated mass M_b, force jumps instead obey $N_2 - N_1 = M_b \dot{v}_b$ with signs determined by forces exerted on that mass from either side. At high frequency the reaction of that mass can mimic a much stiffer termination. Frequency-dependent reflection is then an expected boundary effect, rather than evidence that the bulk modulus varies with frequency.

Connect continuum modes to their lattice predecessor

For a periodic chain the exact discrete equation is

$$m_s \ddot{u}_j = k_s(u_{j+1} - 2u_j + u_{j-1}).$$

Substitute $u_j = U e^{i(kja-\omega t)}$ to obtain

$$\omega = 2\sqrt{k_s/m_s}\,|\sin(ka/2)|.$$

In the first positive Brillouin zone, $\omega/(ck) = \sin(ka/2)/(ka/2)$ and group velocity $d\omega/dk = c\cos(ka/2)$, where $c = a\sqrt{k_s/m_s}$. The continuum rod retains only the leading low-ka term. A wave with eight sites per wavelength has phase speed about $0.9745c$ and group speed $0.9239c$. A narrowband packet's peak can arrive later than the continuum prediction even when a single crest is timed fairly accurately.

A finite-difference code approximating the continuum rod can produce the same spatial dispersion with its mesh spacing in place of a physical lattice spacing. Distinguish them experimentally or numerically by changing the computational mesh while holding the physical structure fixed. Lattice effects persist as the numerical grid is refined; artificial spatial dispersion diminishes. A separate time-step study is needed if a time integrator also disperses the wave. Comparing both peak arrival and phase at two frequencies tests more of the dynamics than fitting a single transit time.

Derive transverse beam waves from curvature energy

A slender beam bending laterally is governed by a different elastic energy from the axial rod. Let $w(x,t)$ denote small transverse deflection of its neutral axis. Under the Euler–Bernoulli assumption that cross sections remain plane and normal to that axis, the axial strain a distance z from the neutral axis is approximately $-zw_{xx}$. Integrating the local elastic energy $E\epsilon_{xx}^2/2$ across the section gives $EI_A w_{xx}^2/2$ per unit length, where $I_A = \int z^2\,dA$ is a geometrical second moment of area, with units m^4. It is distinct from a mass moment of inertia. The kinetic contribution uses line density ρA as before.

Varying $\mathcal{L} = \rho A w_t^2/2 - EI_A w_{xx}^2/2$ requires integrating twice by parts in space. With constant coefficients the interior variation gives

$$\rho A w_{tt} + EI_A w_{xxxx} = 0.$$

The boundary terms are $-[EI_A w_{xx}\delta w_x] + [(EI_A w_{xx})_x \delta w]$. Thus slope variations pair with bending moments and displacement variations with shear forces, with signs set by the chosen face orientation. A clamped end prescribes both deflection and slope; a free end sets moment and shear to zero. A simply supported end prescribes deflection and zero bending moment. Using only one boundary condition per end, as for the axial rod, leaves the fourth-order beam problem underdetermined.

A sinusoid $e^{i(kx-\omega t)}$ gives

$$\omega = \sqrt{EI_A/(\rho A)}\,k^2.$$

Phase speed grows proportionally to wavenumber and group speed is twice phase speed in this model. A narrow transverse pulse in a long beam consequently spreads even in a perfectly homogeneous non-dissipative continuum. By contrast the axial rod was nondispersive in its long-wavelength continuum approximation. At high enough frequency the simple bending model omits rotary inertia, shear deformation and eventually three-dimensional elastic behaviour. Extrapolating its unbounded speed to arbitrary k would use the slender, low-frequency approximation far beyond its derivation.

For a simply supported beam with length L, the sine modes satisfy both end conditions and give

$$\omega_n = (n\pi/L)^2 \sqrt{EI_A/(\rho A)}.$$

For a rectangular section of width w_b and thickness d in the bending direction, $I_A = w_b d^3/12$ and $A = w_b d$. Doubling thickness doubles all the natural angular frequencies at fixed L, E, ρ, whereas doubling width alone leaves the ideal frequencies unchanged. The stiffness increases and so does the inertia. A beam with $L = 0.50\,\mathrm{m}$, $d = 5.0\,\mathrm{mm}$ and the earlier aluminium properties has $\sqrt{EI_A/(\rho A)} \approx 7.35\,\mathrm{m^2/s}$, giving fundamental $46.2\,\mathrm{Hz}$. The second bending mode is near four times that frequency, a clear difference from the approximately harmonic axial mode sequence.

Match a force pulse to a support spring and dashpot

A right termination exerts force on the rod through a dashpot with mechanical impedance Z_L. For a real positive dashpot impedance the reflected velocity fraction is $r = (Z - Z_L)/(Z + Z_L)$. Taking $Z_L = Z$ gives zero reflection at the termination even though energy leaves a finite rod: the dashpot dissipates the arriving power. With $Z_L = 3Z$ one obtains a reflected velocity of $-1/2$ of the incident value, a reflected energy fraction one quarter and absorbed power fraction three quarters. A free or rigid end also reflects energy, but it absorbs no net power in the ideal limit. A measured weak reflection alone is therefore not proof of weak damping inside the rod.

For a pure end spring of stiffness k_L attached to a rigid support, the velocity impedance is $Z_L = k_L/(i\omega)$ under $e^{i\omega t}$. It is imaginary, so the reflection coefficient has unit magnitude at every frequency and a frequency-dependent phase. At low frequency the spring impedance is large and the end looks nearly fixed. At high frequency its impedance is small and the end looks nearly free within the lumped approximation. An impulsive signal contains many frequencies and can return with a changed waveform even without any transmitted or dissipated energy. A support combining a resistor-like dashpot and a spring gives both amplitude loss and phase variation.

A measurement with strain gauges at two positions can separate forward and backward waves if their path delays are resolved. In a narrow frequency band the two records are linear combinations of advancing and returning phase factors. If the gauge spacing corresponds to an integer number of half wavelengths, these combinations become linearly dependent and the separation is ill-conditioned. Change the spacing, use a broadband pulse, or combine a velocity measurement with a strain measurement at one station so that the two characteristic directions enter with opposite signs. The relation $v = (v_+ + v_-)$, $N = Z(-v_+ + v_-)$ then gives a direct local decomposition.

Derive sound propagation from mass and momentum conservation

The displacement-action argument has a fluid analogue. Take an initially uniform fluid of density ρ_0 and pressure p_0, moving little from equilibrium in one spatial dimension. Write small perturbations ρ', p', v. Linearised mass conservation is

$$\partial_t \rho' + \rho_0 \partial_x v = 0.$$

Momentum balance for a short fluid parcel is $\rho_0 \partial_t v = -\partial_x p'$. If compression over an acoustic cycle is approximately isentropic, $p' = c_s^2 \rho'$ with $c_s^2 = (\partial p/\partial \rho)_s$. Differentiate the first equation in time, eliminate the velocity derivative with momentum balance, and obtain $\partial_t^2 p' = c_s^2 \partial_x^2 p'$.

The thermodynamic derivative matters because pressure can respond differently when compression is slow enough for heat exchange. For a perfect gas with ratio of specific heats γ, isentropic perturbations give

$$c_s^2 = \gamma p_0 / \rho_0.$$

Taking an isothermal derivative would give p_0/ρ_0 instead. For $p_0 = 101\,\mathrm{kPa}$, $\rho_0 = 1.20\,\mathrm{kg/m^3}$, and $\gamma = 1.40$, the isentropic speed is about $343\,\mathrm{m/s}$, while the isothermal calculation gives $290\,\mathrm{m/s}$. A measured transit time can discriminate between these constitutive conditions. Wall heat exchange in a very narrow tube can alter which approximation is appropriate over a chosen frequency band.

For a right-moving sound wave $p' = \rho_0 c_s v$, with the pressure positive in compression. This differs in sign from the tensile-positive axial force used for the rod, so match physical force conventions before borrowing formulas across the two examples. The acoustic power per area is $p' v$. For sinusoidal pressure with rms amplitude $1.0\,\mathrm{Pa}$ in the stated gas, the mean intensity is

$$p_{\mathrm{rms}}^2 / (\rho_0 c_s) = 2.43 \times 10^{-3}\,\mathrm{W/m^2}.$$

A peak amplitude of $1.0\,\mathrm{Pa}$ would yield half this number. Writing rms or peak explicitly prevents a routine factor-of-two error when comparing a microphone measurement with an energy flux.

Determine whether a measured transit reveals a modulus or a geometry

Suppose two axially loaded slender rods have measured wave speeds agreeing within their uncertainties. That alone determines only the ratio E/ρ under the slender-rod assumptions. Measuring force–velocity impedance $Z = \rho A c$ and cross section as well gives $\rho = Z/(Ac)$ and $E = Zc/A$. With $Z = 1375\,\mathrm{kg/s}$, $c = 5092\,\mathrm{m/s}$ and $A = 1.0 \times 10^{-4}\,\mathrm{m^2}$, the inversion gives

$$\rho = 2700\,\mathrm{kg/m^3},$$

$$E = 70.0\,\mathrm{GPa}.$$

If area were overestimated by five percent, both inferred quantities would be underestimated by almost five percent while their ratio remained correct.

Now examine a thick elastic solid under a narrow compressional disturbance with wavelengths short compared with its transverse dimensions. Lateral deformation can be constrained by surrounding material. In

an isotropic three-dimensional continuum the longitudinal plane-wave modulus is $\lambda + 2\mu_s$, where λ, μ_s are Lamé coefficients, rather than the uniaxial-stress Young modulus of a slender rod. The shear-wave speed satisfies $c_T^2 = \mu_s/\rho$ and longitudinal speed $c_L^2 = (\lambda + 2\mu_s)/\rho$. Measuring both, together with density, can recover the two independent elastic moduli. Using only one longitudinal speed to infer Young modulus without specifying the deformation geometry can give a physically incorrect value.

For $\rho = 2500\,\mathrm{kg/m^3}$, $c_L = 5000\,\mathrm{m/s}$ and $c_T = 3000\,\mathrm{m/s}$, these relations give

$$\mu_s = 22.5\,\mathrm{GPa},$$

$$\lambda = 17.5\,\mathrm{GPa}.$$

The implied Young modulus $E = \mu_s(3\lambda + 2\mu_s)/(\lambda + \mu_s) = 54.8\,\mathrm{GPa}$, while $\rho c_L^2 = 62.5\,\mathrm{GPa}$. Both numbers represent material stiffness, but they enter different mechanical boundary conditions. A long thin specimen of the same material should have an axial speed closer to $\sqrt{E/\rho} = 4680\,\mathrm{m/s}$. Comparing the bulk and rod experiments tests the constitutive and geometrical assumptions behind the measured speed.

Derive orthogonality when a rod changes material along its length

For a rod with positive varying density $\rho(x)$, area $A(x)$ and modulus $E(x)$, a separated free oscillation $u(x,t) = \phi(x)\cos\omega t$ satisfies

$$-[EA\phi']' = \omega^2 \rho A \phi.$$

Suppose two modes ϕ_n, ϕ_m satisfy the same homogeneous fixed or free boundary conditions and the correct continuity of displacement and axial force across every perfectly bonded interface. Multiply the first equation by ϕ_m, the second by ϕ_n, subtract and integrate. Interior terms produced by integration by parts cancel pairwise, including the interface flux terms, and the endpoint terms vanish. Thus

$$(\omega_n^2 - \omega_m^2)\int \rho A \phi_n \phi_m \, dx = 0.$$

Mode shapes with distinct frequencies are orthogonal under the mass-weighted inner product, not generally under a simple average of displacements measured at uniformly spaced stations. An initial velocity distribution $v_0(x)$ projects onto mode n with coefficient $\int \rho A v_0 \phi_n \, dx / \int \rho A \phi_n^2 \, dx$. In a rod with a heavy section, the same imposed velocity there stores more energy and carries more modal weight. An experiment interpreting raw camera amplitudes without line-density weighting can misassign energy among modes. Normalising each measured mode shape to unit maximum does not fix the mass weighting needed for energies.

Multiplying one eigenvalue equation by its own shape gives

$$\omega_n^2 = \int EA(\phi_n')^2 dx \Big/ \int \rho A \phi_n^2 dx.$$

For an admissible trial function satisfying prescribed displacement conditions, the same quotient bounds the fundamental squared frequency from above when rigid translations are excluded. Consider a uniform fixed–free rod and trial $\phi = x/L$. The numerator is EA/L and denominator $\rho AL/3$, giving $\omega_{\mathrm{trial}} = \sqrt{3}\,c/L$. The exact fundamental is $\pi c/(2L)$, so this one-parameter shape overestimates frequency by about 10.3%. It

need not satisfy the free-end traction condition itself because that is a natural condition of the variational minimiser; it does need to satisfy the essential fixed-end displacement.

For a free–free rod the constant displacement mode has zero frequency. Using an arbitrary positive trial function without removing its mass-weighted mean can drive a quotient towards zero by adding a large constant. To estimate the first deforming mode impose orthogonality to the rigid translation, $\int \rho A \phi \, dx = 0$. This has an immediate mechanical interpretation: the deformation mode should not translate the centre of mass of an isolated rod. The null mode and its mass-weighted constraint reappear in numerical eigenproblems built from discretised stiffness and mass matrices.

Find where a discrete elastic chain ceases to look continuous

Shorten the wavelength relative to the lattice spacing. Compare the phase and group velocities with those inferred from a continuum action.

Open this experiment online

Try the calculation

Sites per wavelength **10**

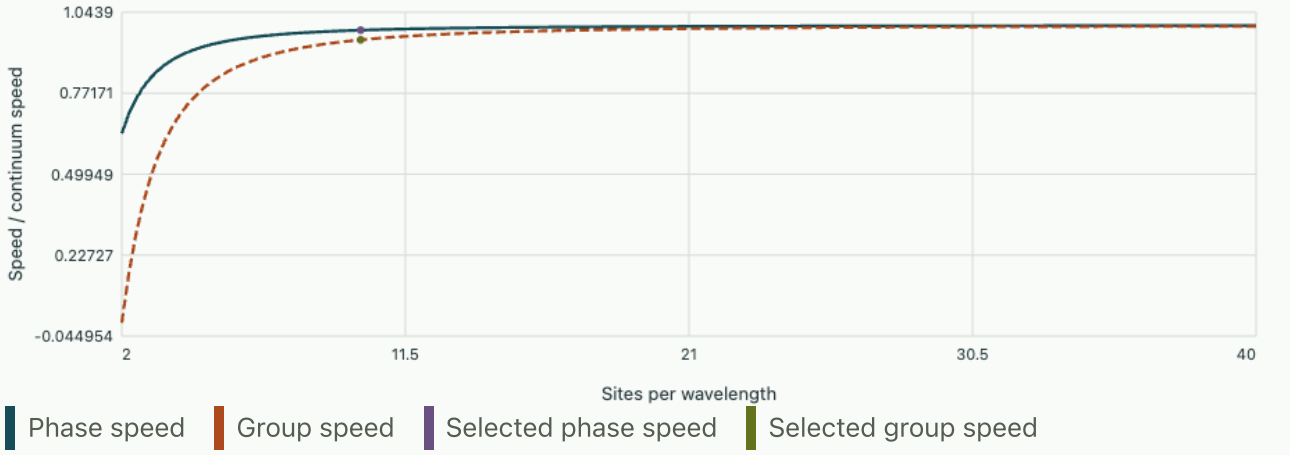

Frequency and phase speed are 98.36% of the continuum prediction. Group speed is 95.11%.

Track work and impulse in a one-dimensional pulse experiment

A finite displacement pulse provides a useful warning about integrating strain and velocity records. A right-going strain pulse of one sign can carry net momentum $\int \rho A v \, dx$. Because $v = -c u_x$, its integrated material momentum is $-\rho A c[u(+\infty) - u(-\infty)]$ at a fixed time. A localised displacement bump that returns to the same displacement on both sides has zero net ordinary material momentum, even though its energy is positive and travels to the right. Its strain and particle-velocity records contain positive and negative parts. A unipolar velocity pulse leaves behind a displacement step, produced by a source that has supplied net impulse.

For a right-moving rectangular velocity pulse of amplitude $0.020\,\mathrm{m/s}$, duration $100\,\mu\mathrm{s}$ in the worked rod, the local displacement jump measured in time is $2.0\,\mu\mathrm{m}$. Its impulse magnitude across a section is $Zv\Delta t = 2.75\times10^{-3}\,\mathrm{N\,s}$, while its transmitted energy is

$$Zv^2\Delta t = 55.0\,\mu\mathrm{J}.$$

A second equal pulse of opposite velocity sign launched later cancels the net displacement and material momentum but doubles the transmitted energy. Thus integrating a signed velocity signal and integrating its square distinguish different properties of the same experimental record.

A displacement sensor has a small linear drift during a burst. Which inferred observables are particularly sensitive?

▼ **Work through the solution**

Differentiating adds a constant offset to particle velocity. Integrating that apparent velocity over a long window produces a growing spurious displacement jump and momentum estimate. Squaring it adds both a cross term with the physical velocity and a positive baseline energy. Measure a pre-pulse baseline, choose an integration interval covering the physical pulse, and test that the result is stable when the window is moderately extended. An independent strain gauge gives another characteristic combination; comparison with $N = -Zv$ can expose drift or an overlapping reflection.

Can a finite-element rod model with a plausible first eigenfrequency still fail an impulse test?

▼ **Work through the solution**

The first frequency mostly probes a long-wavelength weighted stiffness-to-mass ratio. A brief pulse contains higher modes and is sensitive to temporal resolution, dispersion, boundary tractions and the way discrete source work enters. Verify modal frequencies at several wavelengths, check the pulse's positive energy and signed impulse separately, and compare the measured incoming and reflected characteristic fluxes. A single standing-wave eigenvalue cannot validate all of those physically independent balances.

Recover a jump condition at an interface with a point mass

For a concentrated mass M_b joining two rods include $M_b u_t(x_b)^2/2$ as a separate contribution to the Lagrangian. Its time variation contributes $-M_b u_{tt}\delta u$ at the joint. The left and right rod boundary terms contribute respectively $-N_1\delta u$ and $+N_2\delta u$. Requiring stationarity for an independently varying shared displacement gives exactly $N_2 - N_1 = M_b u_{tt}$. This derivation identifies both the sign and the constitutive assumption of a perfectly bonded, nonslipping joint.

For identical rods and a sinusoidal drive at angular frequency ω, the reflected velocity coefficient becomes $r = -i\omega M_b/(2Z + i\omega M_b)$ under the stated phasor convention. At $\omega M_b = 2Z$ it reflects half the power and transmits half. The mass itself absorbs no cycle-averaged energy, though it stores kinetic energy during a cycle. For $Z = 100\,\mathrm{kg/s}$ and $M_b = 0.10\,\mathrm{kg}$, the equal-power frequency is about $318\,\mathrm{Hz}$. A real joint that loses appreciable power at that frequency needs a damping or radiation term in addition to its measured mass.

Tutorial on stress waves

A gauge $0.80\,\mathrm{m}$ from a free end records the echo of a sharp disturbance in the worked rod. What echo delay and strain sign should it expect?

▼ **Work through the solution**

The extra path from gauge to end and back is $1.60\,\mathrm{m}$, giving $1.60/5092 = 0.314\,\mathrm{ms}$. A free end has zero net axial force, so the reflected force and strain have opposite signs to the incident ones. The velocity reflection keeps its sign. Finite pulse width and any change of material between gauge and end can obscure that simple result.

A bonded step has measured reflected velocity coefficient -0.20 and no detectable dissipation. Infer its impedance ratio and transmitted energy fraction.

▼ **Work through the solution**

Rearranging $r_v = (1 - r)/(1 + r)$ gives

$$r = Z_2/Z_1 = (1 - r_v)/(1 + r_v) = 1.50.$$

The energy reflection is 0.04 and transmission 0.96. The velocity transmission is 0.80, whose square alone is only 0.64. Multiplying by the transmitted-to-incident impedance ratio recovers the correct energy fraction.

A rod model with a prescribed displacement at its right end reports nonzero work there, even though the variational boundary term vanished in deriving its equation. Is there a contradiction?

▼ **Work through the solution**

Vanishing of $N\delta u$ follows because the allowed virtual displacement at a prescribed boundary is zero. Actual boundary velocity can be nonzero when the prescribed displacement varies in time. The actuator supplies power Nu_t at the right boundary. Compare the time-integrated input with the change of stored energy plus any losses or outgoing energy rather than identifying an admissible variation with actual displacement.

Compare a finite chain with its long-wavelength continuum.

Fließbach includes string vibrations, beam bending and sound in his continuum mechanics treatment.

Matter and fields in quantum equilibrium

Derive quantum occupations from particle counting. Use them to predict responses of electrons, acoustic modes and photons. Look for the boundary between a one-particle and a collective description.

1. Partition functions and thermal fluctuations
2. Indistinguishable particles and quantum gases
3. Thermal radiation in a cavity
4. Vibrations and the Debye heat capacity
5. Why only a narrow shell of a Fermi gas carries heat
6. Filling electronic states and measuring the Fermi surface
7. Magnetic moments and susceptibility
8. Pairing, excitation gaps and what a spectrum measures
9. Why a polar crystal has two optical phonon frequencies
10. Thermal radiance, étendue and detector calibration
11. Counting, absolute entropy and the entropy of mixing
12. Ensembles, thermodynamic potentials and measurable curvature
13. Rotational and vibrational populations in a molecular gas
14. Bose condensation, dimensionality and a finite trap
15. Degenerate fermions, heat and magnetic response

Partition functions and thermal fluctuations

Thermodynamics and statistical physics

Calculate the occupation, heat capacity and energy fluctuations of a two-level system from a single partition function.

FROM FIRST PRINCIPLES

Obtain Boltzmann weights by counting the bath's available states

Put a small system in weak thermal contact with a much larger bath. Treat the combined pair as isolated with energy E_{tot}. If the small system occupies a particular microstate i of energy E_i, the bath must have energy $E_{\mathrm{tot}} - E_i$. The probability of that microstate is proportional to the number of bath states compatible with the remaining energy,

$$P_i \propto \Omega_{\mathrm{bath}}(E_{\mathrm{tot}} - E_i) = \exp[S_{\mathrm{bath}}(E_{\mathrm{tot}} - E_i)/k_B].$$

A large bath changes temperature negligibly over the small system's range of exchanged energies. Expand its entropy to first order, using $\partial S_{\mathrm{bath}}/\partial E = 1/T$; an arbitrary energy reference contributes only a common factor. Normalizing the resulting probabilities gives

$$P_i = \frac{e^{-\beta E_i}}{Z},$$
$$Z = \sum_j e^{-\beta E_j},$$
$$\beta = \frac{1}{k_B T}.$$

The sum called the partition function counts every microstate, so equal-energy states contribute separately. Differentiating it yields the mean energy $U = -\partial_\beta \ln Z$ when the energy levels are temperature independent. A constant shift of all levels changes Z and U but leaves the normalized probabilities unchanged.

Background · Inverse temperature

$$\beta = 1/(k_B T)$$

It is conjugate to energy in a canonical weight and has units of reciprocal energy. Temperatures in Boltzmann factors are absolute.

DERIVATION AND APPLICATIONS

Measure a static response from spontaneous fluctuations

Consider a coordinate x that can exchange energy with a heat bath at fixed temperature T. A weak constant force f changes its energy from H_0 to $H_f = H_0 - fx$. This sign gives lower energy to positive displacement under a positive force. Assume that the accessible states themselves do not depend on the force and that the instrument allows enough time for the system to sample equilibrium. Write $d\Gamma$ for a sum or integral over states and $\beta = 1/(k_B T)$. Differentiation of the Boltzmann weight gives

$$Z(f) = \int e^{-\beta(H_0 - fx)} d\Gamma,$$

$$\frac{\partial \ln Z}{\partial f} = \beta \langle x \rangle_f.$$

A second derivative acts on both the weighted numerator for the mean and its normalization. The quotient rule subtracts the square of the mean. Hence the static compliance, measured with temperature held fixed, obeys

$$\left(\frac{\partial \langle x \rangle_f}{\partial f} \right)_T = \beta \left(\langle x^2 \rangle_f - \langle x \rangle_f^2 \right).$$

For a harmonic confining potential $H_0 = kx^2/2$, with spring constant $k > 0$, completing the square makes the interpretation especially direct. The position distribution remains Gaussian and its centre shifts with the load,

$$\tfrac{1}{2} k x^2 - f x = \tfrac{1}{2} k (x - f/k)^2 - \frac{f^2}{2k},$$

$$\langle x \rangle_f = \frac{f}{k},$$

$$\operatorname{var} x = \frac{k_B T}{k}.$$

The same stiffness controls response and equilibrium noise. For an anharmonic potential the general variance identity still gives the local response at the current force, although the compliance can change at larger loads. A finite exposure time can average away rapid motion, so measured positions must resolve the fluctuating coordinate.

WORKED EXAMPLE

Calibrate a bead trap at thermal equilibrium

Change the inputs and investigate this derivation.

A bead suspended in a harmonic optical trap at $295\ \mathrm{K}$ has a measured one-coordinate standard deviation of $40.0\ \mathrm{nm}$. Its stiffness is

$$k = k_B T / (40.0 \times 10^{-9})^2 = 2.55 \times 10^{-6}\,\mathrm{N/m}.$$

A steady force of $0.100\,\mathrm{pN}$ should shift the mean by $39.3\,\mathrm{nm}$. Resolving such a shift requires averaging over many statistically independent positions.

SECOND APPLICATION

Predict noise from a measured mechanical compliance

Change the inputs and investigate this derivation.

A second device near a fixed bias force responds to an extra $0.500\,\mathrm{pN}$ with a mean displacement of $10.0\,\mathrm{nm}$. Its local compliance is $2.00 \times 10^{4}\,\mathrm{m/N}$. At $310\,\mathrm{K}$ the expected displacement standard deviation is

$$\sqrt{k_B T (2.00 \times 10^{4})} = 9.25\,\mathrm{nm}.$$

This inference uses only local linear response and thermal equilibrium, so it remains applicable even if the device is nonlinear over a larger displacement range.

Try it · Trap fluctuations

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Temperature (K) **295**

Trap stiffness (μN/m) **2.5**

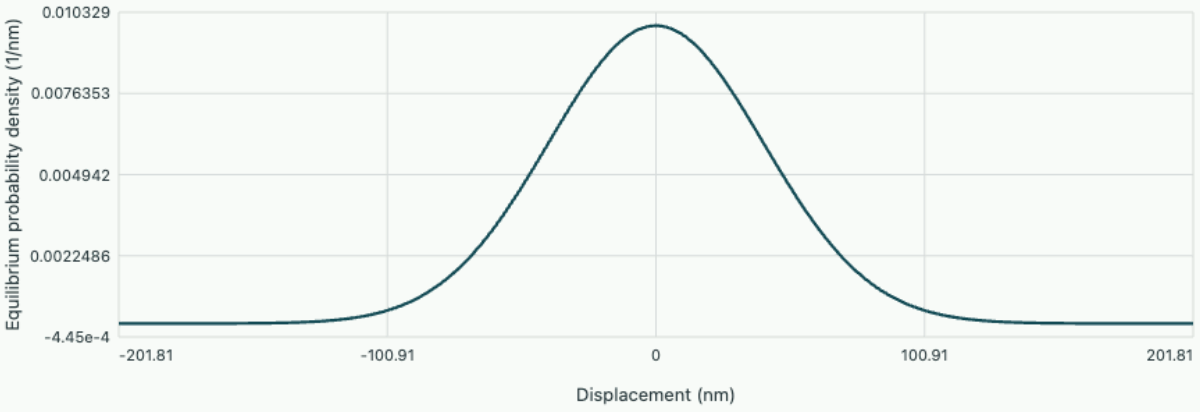

Harmonic equilibrium distribution

Rms displacement 40.363 nm; compliance 1/k = 400000 m/N. Variance = kBT × compliance.

Put a small system in a heat bath

A spectroscopic probe may resolve two energy levels even when a bulk thermometer gives only one temperature. Consider many independent units with ground energy zero and excitation energy $\Delta > 0$. Each unit has two nondegenerate states and is weakly coupled to a much larger bath at temperature T. Weak coupling lets us assign energies to the unit while still permitting energy exchange over the measurement time.

Under the canonical distribution the probability of a microstate j is $p_j = e^{-\beta E_j}/Z$, with $\beta = 1/(k_B T)$. The factor $Z = \sum_j e^{-\beta E_j}$ normalises probabilities. This form comes from expanding the bath entropy about its energy before the exchange: the number of bath states compatible with system energy E_j is proportional to $\exp[-E_j/(k_B T)]$. A bath with appreciable temperature change after one exchange would call for a finite-reservoir correction.

Differentiate the normalisation

The mean energy is $\sum_j p_j E_j$. Differentiating $\ln Z$ with respect to β inserts a factor $-E_j$ in the weighted sum, giving $\langle E \rangle = -\partial_\beta \ln Z$. Differentiating a second time produces $\langle E^2 \rangle - \langle E \rangle^2$, which is the energy variance. Changing variables from β to T produces the fluctuation relation

$$C = \frac{\partial \langle E \rangle}{\partial T} = \frac{\langle E^2 \rangle - \langle E \rangle^2}{k_B T^2}.$$

Hold fixed the external parameters that set the energy levels while taking this derivative. A change of magnetic field or volume can shift the levels and introduce an additional energy response. For independent identically prepared units, variances add. The standard deviation of total energy grows as the square root of the number of units, while its mean grows in proportion to that number. Relative fluctuations consequently shrink as a macroscopic sample is assembled.

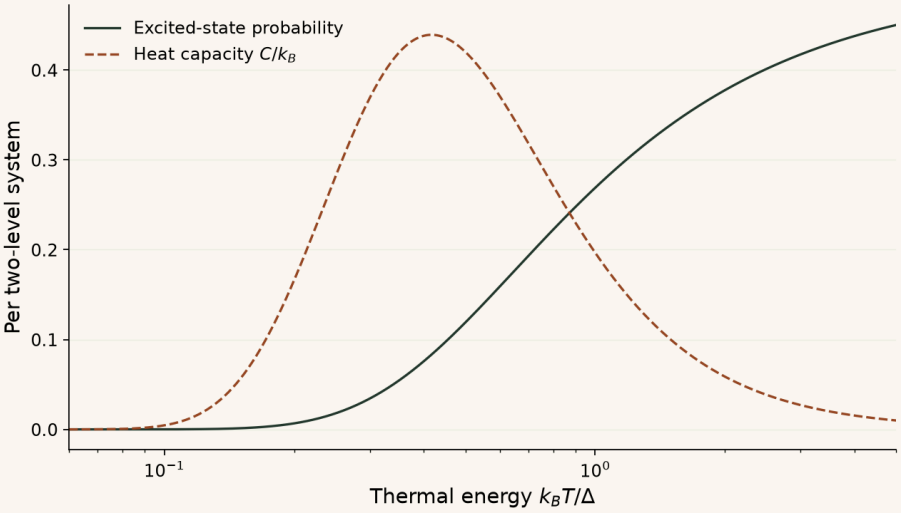

The two-level heat-capacity maximum marks the temperature range where changing the temperature most rapidly redistributes occupation between the levels.

Try it · Two level

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Energy gap / kB (K) **200**

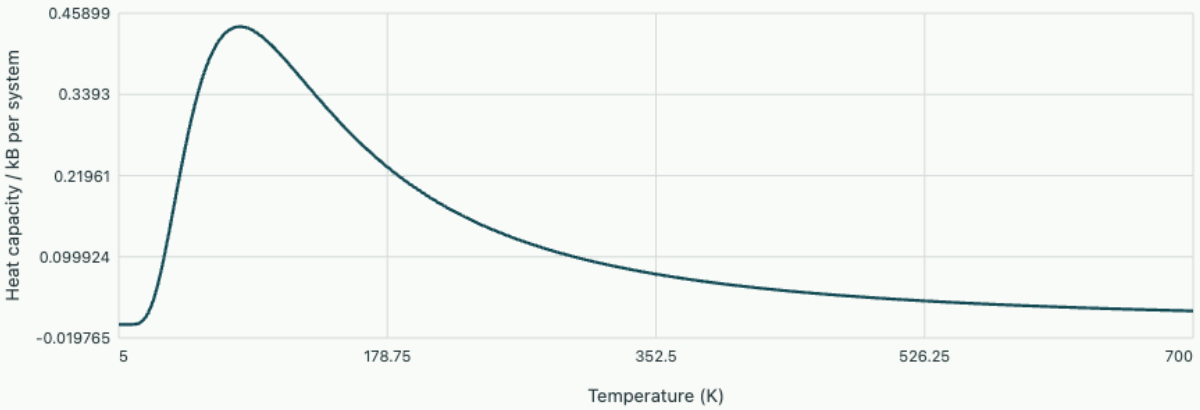

Two-level heat capacity

The nondegenerate two-level maximum lies at T = gap/2.399357 = 83.356 K.

OneSimulation · A finite ring has two routes for a spin correlation

Take N classical spins $s = \pm 1$ on a ring with nearest-neighbour energy $-Js_i s_{i+1}$ and zero external field. Assign the Boltzmann weight of one bond to a two-by-two matrix $T_{ss'} = \exp[Jss'/(k_B T)]$. Matrix multiplication sums over each intermediate spin. Taking the trace joins the last spin back to the first, so $Z = \mathrm{tr}(T^N)$. To measure two spins insert their signs into the same trace at the chosen sites.

$$C(r) = \langle s_0 s_r \rangle = \frac{\mathrm{tr}(S T^r S T^{N-r})}{\mathrm{tr}(T^N)},$$

$$S = \mathrm{diag}(1, -1).$$

The OneSimulation eigensystem computes all separations from the same two modes. At fixed separation, moving the far side of the ring farther away weakens the contribution that wraps around. Distances r and $N - r$ give the same pair correlation even when the infinite-chain approximation predicts very different answers.

```
"""Finite periodic Ising chain: spin correlations from a transfer matrix."""
import matplotlib.pyplot as plt
import numpy as np
```

```
from onesim.numerics.linear.matrices.dense_linear_algebra import symmetric_eigensystem

coupling_over_kT = 0.90
spin = np.array([1.0, -1.0])
transfer = np.exp(coupling_over_kT * np.outer(spin, spin))
eigenvalues, vectors = symmetric_eigensystem(transfer)
ratios = eigenvalues / eigenvalues[-1]  # common scale cancels from the correlation
spin_in_modes = vectors.T @ np.diag(spin) @ vectors
u = np.tanh(coupling_over_kT)
fig, ax = plt.subplots(figsize=(6.6, 3.5), layout="constrained")
for length in (10, 24):
    distances = np.arange(length + 1)
    values = []
    for r in distances:
        numerator = np.sum(spin_in_modes**2 * ratios[:, None]**r * ratios[None, :]**(length-r))
        values.append(numerator / np.sum(ratios**length))
    exact = (u**distances + u**(length-distances)) / (1 + u**length)
    assert np.allclose(values, exact)
    print(f"{length} sites: antipodal correlation {values[length//2]:.4f}")
    ax.plot(distances, values, "o-", markersize=3, label=f"Ring of {length} sites")
ax.plot(distances, u**distances, "--", label="Infinite chain")
ax.set(xlabel="Separation (lattice spacings)", ylabel="Spin correlation")
ax.grid(alpha=0.2)
ax.legend()
fig.savefig("thermodynamics-statistical-physics.png", dpi=180)
```

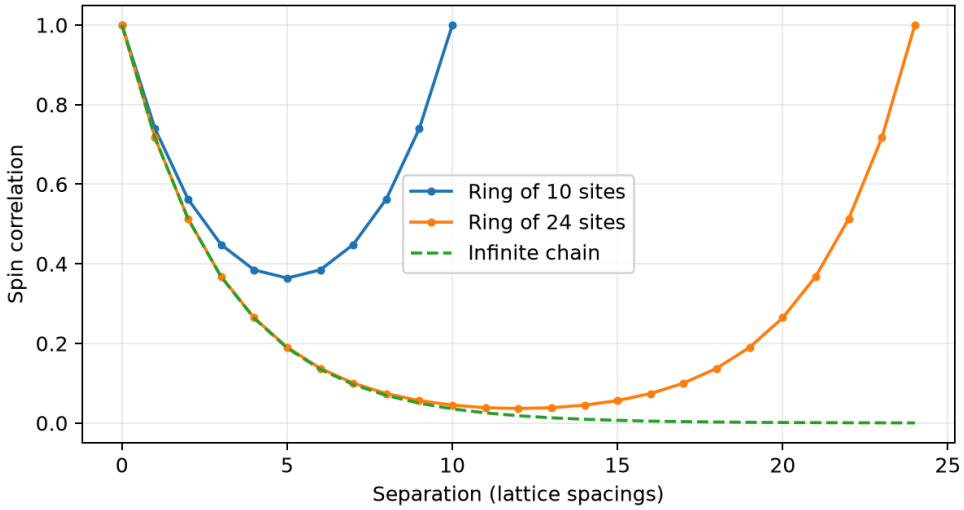

For J/(kBT) = 0.90 the antipodal correlation is 0.364 for ten spins and 0.0365 for 24. At a full turn the observable is the original spin squared.

Try it · Ph ising ring

Explore the same physical relation with adjustable inputs. This independent parameter study has its own displayed reference settings; compare its assumptions and units before matching the static example.

Open this live calculation

Try the calculation

Temperature $k_\beta T/J$ **1**

Sites around the ring **32**

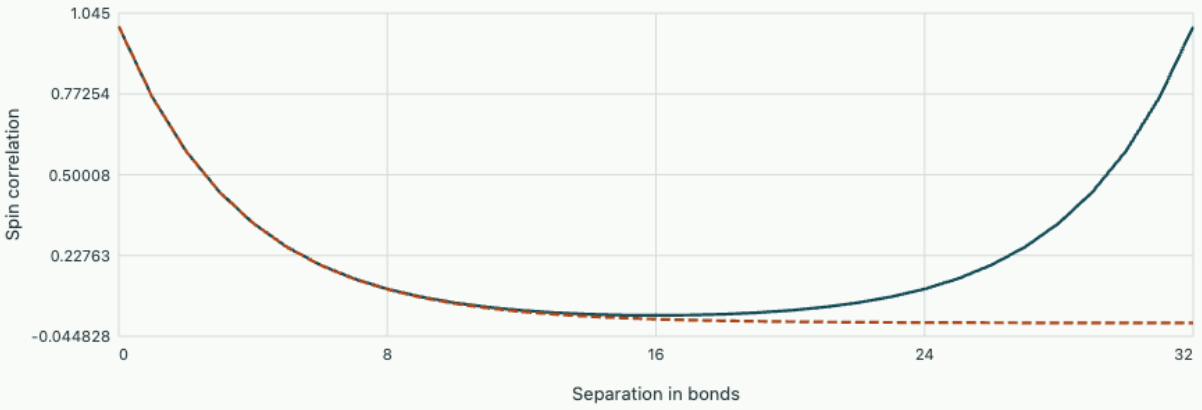

Closed finite ring Infinite chain

Infinite-chain correlation length 3.672 lattice spacings. The opposite site on the 32-site ring has correlation 0.02562.

Download the Python example

Locate the heat-capacity peak

For one two-level unit, $Z = 1 + e^{-x}$ with $x = \Delta/(k_B T)$. Its excited-state probability is

$$p = 1/(1 + e^{x}).$$

Hence $\langle E \rangle = \Delta p$, the variance is $\Delta^2 p(1-p)$ and the heat capacity is $C/k_B = x^2 e^x/(1+e^x)^2$. At low temperature the excited population is exponentially small and little energy can be absorbed by a small temperature increase. At high temperature the two populations approach one half, so changing temperature redistributes them only weakly. A maximum must occur between these limits.

Setting the derivative of $\ln C$ with respect to x to zero gives

$$2/x = \tanh(x/2).$$

Its positive solution is about 2.40, placing the peak at $T \simeq 0.417\Delta/k_B$. This broadened heat-capacity feature is the Schottky anomaly of a two-level population. No collective phase transition is required to produce it. Degeneracy of either level would change the population ratio and modify the peak position.

Infer an energy scale and observe a fluctuation

Take $\Delta = 12.0$ meV, so $\Delta/k_B \simeq 139$ K. A peak near 58 K is expected. At $T = 139$ K one has $x \simeq 1$, $p = 0.269$, mean energy 3.23 meV and $C \simeq 0.197 k_B$ per unit. For a single unit the standard deviation of the energy is

$$\Delta\sqrt{p(1-p)} = 5.32\,\mathrm{meV}.$$

The standard deviation exceeding the mean is possible because most single observations return the ground energy and a minority return the full excitation energy.

For a million independent units at the same temperature, the mean total energy scales up by a million and the standard deviation by a thousand. A calorimeter measuring the bulk value will therefore see a much smoother response than a single-unit spectroscopic probe. If measured fluctuations scale faster than this independent-unit prediction, correlations, common temperature fluctuations or instrumental noise deserve investigation before assigning a microscopic heat capacity.

For a noninteracting collection the total partition function factorizes into one factor per unit. Taking its logarithm converts that product into a sum, so mean energy and heat capacity become extensive. The variance of total energy also grows in proportion to the number of units when correlations vanish. The standard deviation grows only as the square root of that number, explaining how a large specimen can display a reproducible macroscopic energy while its microscopic constituents exchange energy continuously with the bath.

Predict particle-number fluctuations in an open subsystem

The canonical partition function sums over energies for a fixed particle number. A small region that exchanges both particles and energy with a large bath requires another weight. If the region gains energy E and N particles at fixed volume, the bath entropy changes to first order by $(-E + \mu N)/T$, from $dS = (dU - \mu\, dN)/T$. With $\beta = 1/(k_B T)$ and fixed reservoir chemical potential μ, sum over the states of each particle-number sector,

$$\Xi(T, V, \mu) = \sum_{N=0}^{\infty} e^{\beta\mu N} Z_N(T, V),$$

$$\langle N \rangle = \partial_{\beta\mu} \ln \Xi,$$

$$\operatorname{var} N = \partial_{\beta\mu}^2 \ln \Xi.$$

These derivatives are taken at fixed temperature and volume. Each derivative brings down a factor of N from the exponential, and differentiating the normalization supplies the subtraction of $\langle N \rangle^2$. The grand potential is $-k_B T \ln \Xi$. Number fluctuations therefore measure a thermodynamic response,

$$\operatorname{var} N = k_B T (\partial \langle N \rangle / \partial \mu)_{T,V}.$$

Suppose 200 independent adsorption sites are exposed to a large particle reservoir. Each site can be empty with zero energy or contain one particle of energy ϵ. The single-site grand sum is $1 + e^{\beta(\mu-\epsilon)}$ and the sample sum is its 200th power. Choose $\mu - \epsilon = -k_B T \ln 3$: occupancy probability is $1/4$, mean population is 50,

variance is $200(1/4)(3/4) = 37.5$ and standard deviation is 6.12 particles. Doubling the number of sites doubles mean and variance, while relative fluctuations fall by $\sqrt{2}$. If the entire sample were sealed with exactly 50 particles, its total-number variance would instead be zero. A small observation window inside it could still fluctuate.

Count rotational states of a diatomic molecule

A partition sum can also show how a continuous-looking thermal degree of freedom switches on. Model a heteronuclear diatomic molecule as a rigid rotor of fixed bond length and moment of inertia I about an axis perpendicular to the bond. Ignore rotation about the bond, vibrational excitation, electronic excitation and intermolecular interactions. Its quantum levels are $E_J = \hbar^2 J(J+1)/(2I) = k_B \Theta_r J(J+1)$, with degeneracy $2J+1$ from the independent projections of angular momentum. Thus

$$Z_{\mathrm{rot}} = \sum_{J=0}^{\infty} (2J+1) e^{-\Theta_r J(J+1)/T},$$

$$\Theta_r = \frac{\hbar^2}{2 I k_B}.$$

At temperatures large compared with Θ_r, many J values contribute and the sum is well approximated by an integral. Substitute $u = J(J+1)$, which supplies $du = (2J+1)dJ$. The integral is then T/Θ_r. Differentiating $\ln Z$ gives $U_{\mathrm{rot}} \simeq k_B T$ and $C_{\mathrm{rot}} \simeq k_B$ per molecule. The two classical quadratic rotational degrees of freedom each supply one half of this energy.

At very low temperature the first excited level lies $2k_B\Theta_r$ above the ground state and has degeneracy three. Keeping just $J = 0, 1$ gives $Z \simeq 1 + 3e^{-2\Theta_r/T}$. Its excitation probability is the second term divided by Z. The mean rotational energy and heat capacity become exponentially small. If nuclei are identical, exchange symmetry can change allowed rotational populations through nuclear-spin weights; the deliberately heteronuclear model avoids that extra counting condition.

Infer when the first rotational line becomes populated

Use an illustrative rotor with $\Theta_r = 25.0$ K. At $T = 10.0$ K the weight of $J = 1$ is $3e^{-5} = 0.0202$ and that of $J = 2$ is

$$5e^{-15} = 1.53 \times 10^{-6}.$$

Neglecting the latter gives excited fraction 0.0198. The gap is $50.0 k_B$ K, giving mean energy per molecule approximately $0.990 k_B$ K. The two-level variance formula gives

$$C_{\mathrm{rot}}/k_B = (50/10)^2 p(1-p) = 0.486.$$

The heat capacity is a derivative and can already be appreciable even when the excited-state fraction is only two percent.

At $T = 250$ K many rotational states contribute and the leading high-temperature sum gives $Z \simeq 10$ and molar rotational heat capacity near R. The line positions also constrain structure. Inverting $\Theta_r = \hbar^2/(2Ik_B)$

gives $I = 1.61 \times 10^{-47}$ kg m². With independently known reduced nuclear mass $\mu = 1.50 \times 10^{-27}$ kg, the effective bond length $\sqrt{I/\mu} = 0.104$ nm. Calorimetry and rotational spectroscopy can thus test the same moment of inertia through distinct observables.

Absolute line intensities provide an additional check on the rotor model. A line originating from the first excited rotational manifold samples its occupation, but a measured signal also contains transition matrix elements and instrument response. To compare with the 1.98-percent prediction, either calibrate those factors or compare the same line across temperatures. The gap can then be estimated from the slope of a logarithmic population ratio against inverse temperature, while the ratio's intercept contains the degeneracy. This separates a genuine error in moment of inertia from an omitted factor counting the allowed projection states.

At higher temperatures vibrational excitation can add to rotational calorimetry. An energy-resolved line measurement helps identify which internal degree of freedom has actually become populated.

Indistinguishable particles and quantum gases

Thermodynamics and statistical physics

Use single-state occupancy to derive Fermi–Dirac and Bose–Einstein distributions and estimate when classical gas statistics cease to apply.

FROM FIRST PRINCIPLES

Derive the mean population of one quantum state

Consider a single-particle state of energy ϵ exchanging particles and energy with a large reservoir at temperature T and chemical potential μ. In a noninteracting gas, occupation n of this state changes $E - \mu N$ by $n(\epsilon - \mu)$. Its statistical weight is q^n, where $q = e^{-(\epsilon-\mu)/(k_B T)}$. For one specified state including its spin label, fermions allow occupations 0 and 1. Bosons allow every nonnegative integer,

$$Z_F = 1 + q,$$

$$Z_B = 1 + q + q^2 + \cdots = \frac{1}{1-q}.$$

The Bose sum requires $q < 1$. Differentiating a weight with respect to q brings down its occupancy, so the normalized average is $q\,\partial_q \ln Z$. Applying this to the two sums gives

$$\bar{n}_F = \frac{1}{e^{(\epsilon-\mu)/(k_B T)} + 1},$$

$$\bar{n}_B = \frac{1}{e^{(\epsilon-\mu)/(k_B T)} - 1}.$$

For $q \ll 1$, both means approach q; visits by additional particles to the same state are then rare. The difference between the distributions follows entirely from which occupation numbers the quantum statistics allow.

DERIVATION AND APPLICATIONS

Derive the first exchange correction to ideal-gas pressure

Before quantum degeneracy becomes strong, exchange statistics already affect the relation between pressure and density. Consider noninteracting particles of mass m in a large three-dimensional box. All g internal states have the same energy and chemical potential. Introduce fugacity $z = e^{\mu/(k_B T)}$, thermal wavelength $\lambda_T = h/\sqrt{2\pi m k_B T}$, and a sign $s = +1$ for bosons or -1 for fermions. At small fugacity the occupation of a one-particle energy ϵ is a geometric expansion,

$$\bar{n}(\epsilon) = z e^{-\beta\epsilon} + s z^2 e^{-2\beta\epsilon} + O(z^3),$$

$$\beta = (k_B T)^{-1}.$$

Integrating the first term over momentum gives g/λ_T^3 times z. In the second term, doubling the inverse temperature narrows each Cartesian Gaussian by $\sqrt{2}$. The density integral therefore gains a factor $2^{-3/2}$. At

fixed temperature the Gibbs–Duhem relation $dp = n\,d\mu$ gives

$$d[p/(k_B T)] = n\,dz/z.$$

Integration from the empty-gas limit gives the pressure series,

$$\frac{n\lambda_T^3}{g} = z + \frac{sz^2}{2^{3/2}} + O(z^3),$$
$$\frac{p\lambda_T^3}{g k_B T} = z + \frac{sz^2}{2^{5/2}} + O(z^3).$$

Experiments often fix density. Put $y = n\lambda_T^3/g$ and invert the density series once to get

$$z = y - sy^2/2^{3/2} + O(y^3).$$

Substituting into the pressure eliminates the unmeasured chemical potential,

$$\frac{p}{nk_B T} = 1 - \frac{s}{2^{5/2}}\frac{n\lambda_T^3}{g} + O(y^2).$$

At fixed density exchange raises the pressure of fermions and lowers that of bosons. The magnitude grows on cooling because the thermal wavelength increases. Real interactions can also alter the pressure, and the truncated expansion requires small y.

WORKED EXAMPLE

Estimate a pressure excess for a dilute Fermi gas

Change the inputs and investigate this derivation.

A spin-balanced ideal Fermi gas has $g = 2$, $n = 2.00 \times 10^{18}\,\mathrm{m}^{-3}$ and $\lambda_T = 400\,\mathrm{nm}$. It has $y = 0.0640$ and $p/(nk_B T) \simeq 1.0113$. A measured classical pressure scale $nk_B T = 20.0\,\mathrm{pPa}$ would therefore acquire an exchange correction of about $0.226\,\mathrm{pPa}$.

SECOND APPLICATION

Recover density from a bosonic pressure deficit

Change the inputs and investigate this derivation.

A single-state Bose gas at known $\lambda_T = 500\,\mathrm{nm}$ gives

$$p/(nk_B T) = 0.9950.$$

To leading order the deficit fixes $y = 0.0050 \times 2^{5/2} = 0.0283$. The inferred density is

$$y/\lambda_T^3 = 2.26 \times 10^{17}\,\mathrm{m}^{-3}.$$

At unchanged temperature halving this density halves the fractional pressure deficit, whereas the absolute pressure correction falls by a factor of four.

Try it · Quantum virial

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Selected degeneracy nλ³/g **0.065**

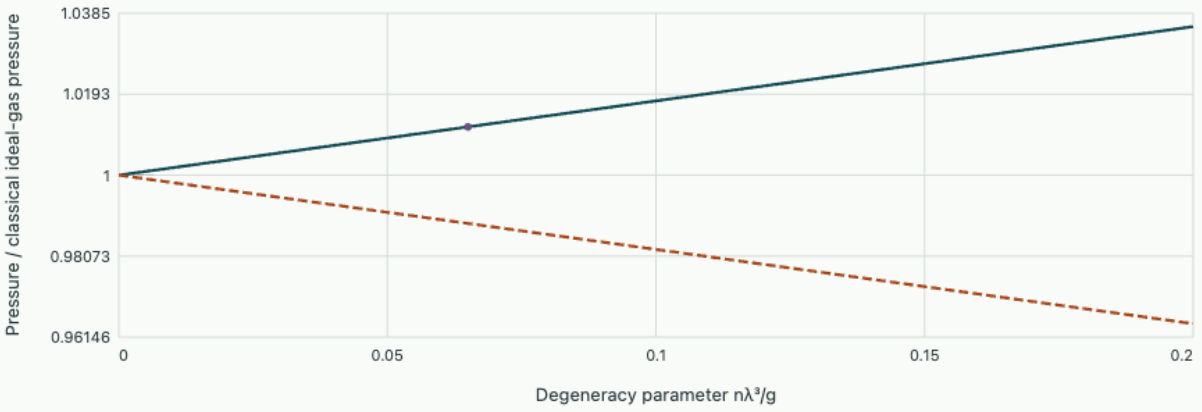

Fermi leading correction | Bose leading correction | Selected Fermi gas

These are first quantum virial corrections, valid for nλ³/g≪1. Exchange raises the dilute fermion pressure and lowers the dilute boson pressure.

Decide which exchanges leave the state unchanged

An ideal classical gas can often be treated with Maxwell–Boltzmann statistics. Cooling the gas or increasing its density eventually makes wave packets overlap appreciably. For indistinguishable quantum particles, merely assigning names to exchanged particles must not create extra states. In an ideal gas one can instead specify the occupation of each single-particle state. Identical fermions occupy a given quantum state at most once, including its spin label. Identical bosons can accumulate in the same state.

Use a reservoir that exchanges both energy and particles with the system. Its temperature T and chemical potential μ are fixed. A microstate with energy E and particle number N has grand-canonical weight proportional to $\exp[-\beta(E-\mu N)]$. This follows by applying the thermodynamic entropy differential to a reservoir that loses both E and N. The chemical potential sets the mean particle number through these statistical weights.

Sum the allowed occupancies of one state

For a single-particle state of energy ϵ, define $z_\epsilon = \exp[-\beta(\epsilon-\mu)]$. An occupancy of n contributes weight z_ϵ^n. For fermions the single-state grand partition sum is $1+z_\epsilon$, so the mean occupancy is $z_\epsilon/(1+z_\epsilon)$. For bosons the sum runs over all nonnegative integers and equals $1/(1-z_\epsilon)$, provided $z_\epsilon<1$. Differentiating the logarithm of each sum with respect to $\ln z_\epsilon$ gives

$$\bar{n}_F = \frac{1}{e^{\beta(\epsilon-\mu)}+1},$$

$$\bar{n}_B = \frac{1}{e^{\beta(\epsilon-\mu)}-1}.$$

These expressions describe one state, not the total population within an energy interval. To get the latter, multiply by the density of states and integrate. If the exponential is large then both occupancies reduce to $e^{-\beta(\epsilon-\mu)}$. The upper sign suppresses double occupancy while the lower sign enhances occupation near the lowest available energy.

Count the occupied sphere of a cold electron gas

In a large cubic volume V with periodic boundaries, allowed wave vectors fill reciprocal space with one state per volume $(2\pi)^3/V$ for each spin. At zero temperature the noninteracting electron gas fills all states inside a sphere of radius k_F and leaves higher ones empty. Including two spin states per wave vector gives $N = 2V(4\pi k_F^3/3)/(2\pi)^3$, so $k_F = (3\pi^2 N/V)^{1/3}$. For electrons of mass m_e, the Fermi energy is $\hbar^2 k_F^2/(2m_e)$.

With density $n = 8.0\times10^{28}$ m⁻³, $k_F = 1.33\times10^{10}$ m⁻¹ and $E_F \simeq 6.77$ eV. The associated temperature $T_F = E_F/k_B$ is approximately 7.9×10^4 K. At room temperature k_BT is small compared with this energy, so only a narrow region around the Fermi surface changes occupation noticeably. A classical equipartition estimate of the electron contribution to heat capacity would ignore this restriction.

Test the dilute limit with a cold atomic gas

A useful scale for translational wave packets is the thermal wavelength $\lambda_T = h/\sqrt{2\pi m k_B T}$. For a single internal state the dimensionless product $n\lambda_T^3$ compares the volume per particle with the thermal wave volume. Small values support the dilute classical approximation. Near unity the occupancy of low-energy states can become appreciable and the statistics of indistinguishability matter.

For particles of mass 6.6×10^{-27} kg at 1.0 K the thermal wavelength is approximately 0.88 nm. At density 10^{25} m⁻³, $n\lambda_T^3 \simeq 0.0068$. Reducing temperature to 0.040 K at unchanged density multiplies that product by $25^{3/2} = 125$, giving about 0.85. This is a degeneracy estimate, not a prediction of the phase of a real interacting gas. Once quantum statistics matter, particle interactions, dimensionality and confinement can also affect which phase the experiment actually finds.

Internal degeneracy enters these comparisons through the density per available internal state. Suppose the same total particle density is spread equally across two spin states with identical spatial dispersion. Each can fill the same momentum levels independently. In a cold ideal Fermi gas the Fermi wavevector at fixed total density is then smaller than in a fully spin-polarized gas. A measured pressure change after altering spin

populations can therefore contain a statistical contribution even if the mass, volume and total particle count remain fixed.

Determine the capacity of excited Bose states

Ideal bosons can accumulate macroscopically in one spatial state. Consider a uniform three-dimensional noninteracting gas of identical spinless particles with fixed particle number in a large box. Set the lowest energy to zero. The excited-state occupation at energy ϵ is $1/[\exp((\epsilon-\mu)/(k_B T))-1]$ and requires $\mu \le 0$. In the thermodynamic limit, counting momentum states in a sphere gives an excited-state density per volume proportional to $\sqrt{\epsilon}$. Consequently, as the density rises, taking μ toward zero from below can increase the excited population only up to a finite value.

Explicitly integrate $d^3p/(2\pi\hbar)^3$ over positive energies $\epsilon = p^2/(2m)$. At $\mu = 0$, setting $x = \epsilon/(k_B T)$ reduces the integral to $\int_0^\infty x^{1/2}/(e^x-1)dx = \Gamma(3/2)\zeta(3/2)$. The result is

$$n_{\mathrm{ex,max}} = \frac{\zeta(3/2)}{\lambda_T^3},$$

$$\lambda_T = \frac{h}{\sqrt{2\pi m k_B T}},$$

$$\zeta(3/2) = 2.6124\ldots.$$

Below the temperature where this capacity equals the total density, additional particles occupy the lowest state. Solving gives

$$T_c = (2\pi\hbar^2/m k_B)[n/\zeta(3/2)]^{2/3},$$

$$N_0/N = 1 - (T/T_c)^{3/2}.$$

The state count explains why a high phase-space density matters. Reaching this regime does not require individual particles to lose their identity at a sharply defined minimum separation.

Estimate the onset temperature in a box

Take $n = 1.00 \times 10^{20}$ m⁻³ and particle mass 1.44×10^{-25} kg. At the ideal transition the thermal wavelength is

$$(2.6124/n)^{1/3} = 297\,\mathrm{nm}.$$

Inverting its definition gives $T_c = h^2/(2\pi m k_B \lambda_T^2) = 0.398$ μK. At $T = 0.250$ μK, the excited fraction is $(0.250/0.398)^{3/2} = 0.498$ and the condensate fraction is approximately 0.502. Doubling the density raises T_c by $2^{2/3} = 1.59$ at fixed particle mass.

The predicted fraction tends to zero continuously as T approaches T_c from below, and tends to one as temperature approaches zero for the ideal gas. A harmonic trapping potential has a different density of states and hence a different temperature exponent from a uniform box. Finite size smooths the mathematical onset; interactions and inaccurate thermometry can shift its inferred location. Comparing the number in the narrow

lowest-momentum component with the thermal-cloud population requires that the analysis resolve these two contributions rather than simply count all particles near the centre of a spatial image.

The dimensionality of the gas is as essential as Bose exchange itself. In a uniform two-dimensional ideal box the single-particle density of states tends to a constant at low energy. At chemical potential approaching zero, Bose occupation varies as $k_B T/\epsilon$ there. Its energy integral consequently diverges logarithmically rather than approaching a finite excited-state capacity. The three-dimensional derivation of a nonzero transition temperature cannot then be carried across unchanged. A finite box, a confining trap or interparticle interactions can all alter the observed low-energy behaviour and must be identified when comparing different experimental geometries.

Repeat the counting in a harmonic trap

A harmonic trap changes the number of available excited states even for particles with the same mass and statistics. For three oscillator frequencies $\omega_x, \omega_y, \omega_z$, measure energy above the lowest level and suppose $k_B T$ greatly exceeds each $\hbar\omega_i$. The quantum labels n_i are nonnegative integers with energy $\sum_i \hbar\omega_i n_i$. The count below energy E is approximately the volume of the positive-coordinate tetrahedron bounded by this plane, namely $E^3/(6\hbar^3\omega_x\omega_y\omega_z)$. Differentiating gives an energy density of states proportional to E^2.

At the onset of condensation integrate this density against $1/(e^{E/k_B T}-1)$. The identity $\int_0^\infty x^2/(e^x-1)\,dx = 2\zeta(3)$ then gives

$$N_{\mathrm{ex,max}} = \zeta(3)\left(\frac{k_B T}{\hbar\bar{\omega}}\right)^3,$$

$$\bar{\omega} = (\omega_x\omega_y\omega_z)^{1/3}.$$

For $N = 1.00\times10^5$ particles in an isotropic 100 Hz trap, $\hbar\bar{\omega}/k_B = 4.80$ nK and $[N/\zeta(3)]^{1/3} = 43.7$, so the predicted transition temperature is 0.210 μK. At 0.150 μK the condensate fraction is

$$1 - (0.150/0.210)^3 = 0.636.$$

An analysis using the uniform-box exponent at the same temperature ratio would instead give 0.396. The difference comes entirely from the confinement-dependent state count.

Reducing all three trap frequencies at fixed particle number lowers the onset temperature in direct proportion. This differs from enlarging a uniform box at fixed mean density, which leaves its bulk critical temperature unchanged. Specifying how particle number, volume and external confinement are controlled is therefore essential when interpreting a measured condensate fraction.

Thermal radiation in a cavity

Thermodynamics and statistical physics

Count electromagnetic states, derive the temperature scaling of radiation energy and connect a cavity measurement to emitted power.

FROM FIRST PRINCIPLES

Derive the cavity spectrum by counting electromagnetic modes

Consider a large empty cavity at uniform temperature T. With periodic boundaries each independent wavevector occupies reciprocal volume $(2\pi)^3/V$, and each nonzero wavevector admits two transverse polarizations. Counting a spherical shell of radius k and thickness dk, then using $\omega = ck$, gives the number of modes per unit cavity volume and angular-frequency interval,

$$\frac{g(\omega)}{V}\,d\omega = \frac{2(4\pi k^2 dk)}{(2\pi)^3} = \frac{\omega^2}{\pi^2 c^3}\,d\omega.$$

The cavity walls can create and absorb photons, so equilibrium photon number is unrestricted and its chemical potential is zero. A mode has mean occupation $1/[e^{\hbar\omega/(k_B T)} - 1]$. Multiplying by its photon energy $\hbar\omega$ and the density of modes gives the thermal spectrum,

$$u_\omega(T) = \frac{\hbar\omega^3}{\pi^2 c^3 [e^{\hbar\omega/(k_B T)} - 1]}.$$

This quantity is energy density per unit angular frequency. Integrating over frequency and changing variable to $x = \hbar\omega/(k_B T)$ gives a total proportional to T^4. The continuum mode count assumes a cavity large compared with the wavelengths contributing most of that energy.

DERIVATION AND APPLICATIONS

Locate the peak that a spectrometer actually measures

A cavity at temperature T emits thermal radiation through a small aperture. An instrument can report energy in equal intervals of frequency or equal intervals of wavelength. The energy in any given physical band agrees, but the positions of the spectral-density maxima depend on the bin coordinate. Assume an ideal blackbody and a calibrated detector reporting energy, with response corrections applied. The spectral energy density per ordinary frequency ν has the Planck form

$$u_\nu = \frac{8\pi h\nu^3}{c^3}\,\frac{1}{e^{h\nu/(k_B T)} - 1}.$$

The energy in corresponding bins must be equal. Substitute $\nu = c/\lambda$ and include the magnitude of its derivative to obtain the density per unit wavelength,

$$u_\lambda = u_\nu(c/\lambda)\left|\frac{d\nu}{d\lambda}\right| = \frac{8\pi hc}{\lambda^5}\frac{1}{e^{hc/(\lambda k_B T)}-1}.$$

For frequency bins set $x = h\nu/(k_B T)$; for wavelength bins set $x = hc/(\lambda k_B T)$. Temperature-dependent prefactors can be ignored in locating the maximum at fixed T. Both maximizations reduce to differentiating $x^m/(e^x - 1)$, with different powers m. Its logarithmic derivative vanishes when

$$\frac{m}{x} - \frac{1}{1-e^{-x}} = 0,$$
$$x = m(1 - e^{-x}).$$

The positive interior roots are $x_\nu = 2.82144$ for $m = 3$ and $x_\lambda = 4.96511$ for $m = 5$. The stationary points are maxima because both tails fall below the interior peak. Bin widths and what the detector counts are part of the observable. A photon counter without energy weighting or a spectrometer recording per logarithmic bandwidth would require its own density transformation.

WORKED EXAMPLE

Compare two instruments viewing the same furnace

Change the inputs and investigate this derivation.

For $T = 1200\,\mathrm{K}$, the wavelength instrument peaks at $hc/(4.96511 k_B T) = 2.41\,\mu\mathrm{m}$. The frequency instrument peaks at $2.82144 k_B T/h = 70.5\,\mathrm{THz}$, corresponding to $4.25\,\mu\mathrm{m}$. Both spectra describe the same beam. A wavelength interval of a given width spans more frequency bandwidth at shorter wavelengths, which shifts its maximum toward higher photon energies.

SECOND APPLICATION

Infer temperature and forecast another peak

Change the inputs and investigate this derivation.

A calibrated equal-frequency instrument finds a maximum at $35.0\,\mathrm{THz}$ from a second cavity. This gives

$$T = h\nu/(2.82144 k_B) = 595\,\mathrm{K}.$$

An equal-wavelength instrument should then peak at $2.898 \times 10^{-3}\,\mathrm{m\,K}/595\,\mathrm{K} = 4.87\,\mu\mathrm{m}$. The wavelength $c/\nu = 8.57\,\mu\mathrm{m}$ still identifies the same photons as the frequency maximum and can be used to match a narrow band between instruments.

Try it · Blackbody peaks

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Temperature (K) **1200**

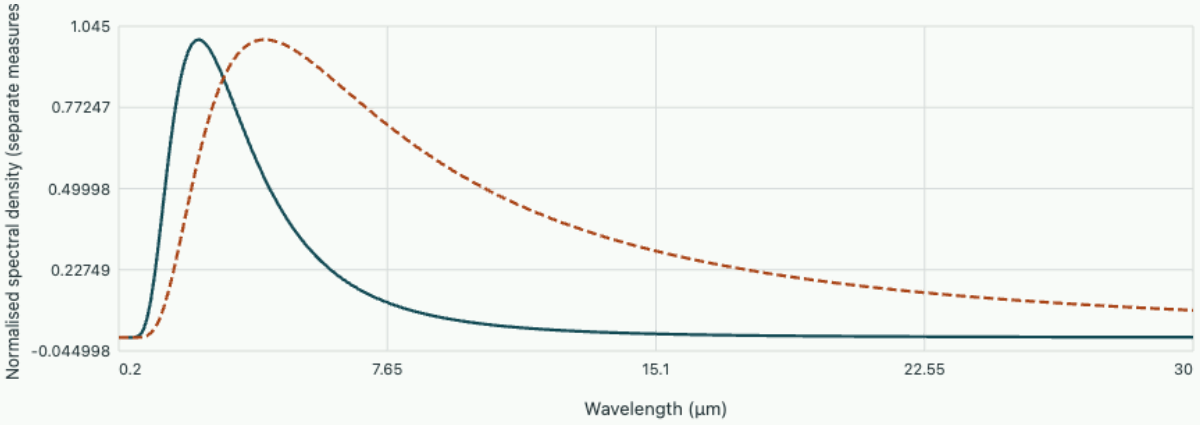

Equal-wavelength channels Equal-frequency channels, labelled by c/ν

Wavelength-density peak 2.4148 µm; frequency-density peak corresponds to 4.2495 µm. The two densities use different channel widths and are not related by relabelling alone.

Photons in thermal equilibrium

An empty cavity with walls at temperature T contains electromagnetic radiation. The walls absorb and emit photons, so equilibrium does not constrain their number. The chemical potential of equilibrium photons is zero. Each mode of angular frequency ω carries a thermal mean energy $\hbar\omega/[\exp(\hbar\omega/k_BT)-1]$ above its temperature-independent vacuum energy. The spectrum follows by multiplying this occupation by the number of modes in a frequency interval.

Take a large box of volume V and periodic boundary conditions for counting. One state per polarization occupies a reciprocal-space cell $(2\pi)^3/V$. A thin spherical shell contributes $4\pi k^2 dk$ times two independent transverse polarizations. Using $\omega = ck$ gives the state density per volume $\omega^2/(\pi^2 c^3)$. Hence

$$u_\omega(T) = \frac{\hbar\omega^3}{\pi^2 c^3}\frac{1}{e^{\hbar\omega/k_BT}-1}.$$

At low frequency the mean mode energy approaches k_BT as in classical equipartition. Extending that classical formula to arbitrarily high frequencies would produce an infinite integral. Quantum occupation

suppresses these modes exponentially when a single photon costs much more than the thermal energy.

Energy, pressure and flux

Integrate over positive ω, putting $x = \hbar\omega/(k_B T)$. The integral $\int_0^\infty x^3/(e^x - 1)\,dx = \pi^4/15$ gives energy density $u = aT^4$ with $a = \pi^2 k_B^4/(15\hbar^3 c^3)$. Pressure follows from radiation momentum: a photon carries momentum E/c, while crossing a plane weights directions by their projected speed. Evaluating the normal momentum flux over isotropic directions introduces $\langle \cos^2\theta \rangle = 1/3$, hence $p = u/3$.

A small hole in a black cavity admits only outgoing directions. The half-sphere average of the normal speed gives net emitted flux $cu/4$, which is the Stefan–Boltzmann law σT^4 with $\sigma = ac/4$. The different factors of three and four describe different angular averages and observables. A wall facing another large black surface also receives thermal radiation, so its net radiative heat loss depends on the fourth powers of both temperatures.

Observe cooling power

A 1.00 cm² aperture faces a large black surrounding enclosure at 300 K. The cavity temperature is 600 K. The outgoing power is $\sigma A T^4 = 0.735$ W while incoming power is 0.0459 W. The difference, 0.689 W, must be replaced to maintain temperature if all other thermal exchanges are negligible. The interior radiation energy density is only 9.81×10^{-5} J/m³; repeated replenishment at the walls sustains the outgoing flux.

A wavelength spectrometer and a frequency spectrometer plot different spectral densities. To convert use $u_\lambda = u_\nu |d\nu/d\lambda|$, since both must assign the same energy to corresponding intervals. Merely substituting $\nu = c/\lambda$ into a density per hertz shifts the plotted maximum incorrectly. The integrated energy and the predicted net aperture power are independent of which spectral coordinate is chosen.

The photon equation of state also predicts a reversible adiabatic temperature change. Write $U = aVT^4$ and $p = aT^4/3$, then set $dU = -p\,dV$. Collecting terms gives $dT/T = -dV/(3V)$, or $TV^{1/3} = \text{constant}$. Halving the linear dimension of a cubical equilibrium radiation cavity would double its temperature. This scaling agrees with the shift of every electromagnetic mode frequency under uniform contraction.

Determine the entropy carried by equilibrium radiation

A beam emitted through a small black-cavity opening carries entropy as well as energy. For large equilibrium photon systems, extensivity gives the Euler relation $U = TS - pV + \mu N$. The photon chemical potential is zero and $p = u/3$, where $u = U/V = aT^4$. Solving for entropy density gives

$$s = (u + p)/T = 4aT^3/3.$$

The same result follows independently by integrating $dS = C_V dT/T$ at fixed volume, since $C_V = 4aVT^3$.

Inside the cavity the radiation is isotropic. Photons leaving through a small aperture sample each outgoing direction in proportion to its projected velocity, giving $c/4$ times the interior density for either energy or entropy. Thus the emitted energy flux is σT^4 and entropy flux is $(4/3)\sigma T^3$. The latter exceeds energy flux divided by T by a factor of $4/3$. Opening the cavity to a colder environment is an irreversible emission

process, not the reversible transfer of an equal amount of heat through a material boundary at the source temperature.

Consider a device intercepting ideal black radiation from a hot reservoir at T_h while exchanging a counterpropagating equilibrium field with an environment at T_0. The net incoming energy flux per area is $\sigma(T_h^4 - T_0^4)$ and the corresponding net entropy flux is $(4/3)\sigma(T_h^3 - T_0^3)$. A reversible converter must return this entropy to the ambient reservoir, requiring at least T_0 times the net entropy flux as rejected heat. The remaining energy is the maximum possible work for these specified radiation streams,

$$\frac{\dot{W}_{\max}}{A} = \sigma T_h^4 \left[1 - \frac{4}{3}\frac{T_0}{T_h} + \frac{1}{3}\left(\frac{T_0}{T_h}\right)^4\right].$$

Evaluate an optical energy and entropy budget

Take a 1.00 cm² aperture, $T_h = 600$ K and $T_0 = 300$ K. Outgoing energy is 0.735 W and incoming ambient radiation 0.0459 W. Net energy input is therefore 0.689 W. The net entropy supply is $(4/3)\sigma A(600^3 - 300^3) = 1.43 \times 10^{-3}$ W/K. Returning this to 300 K requires at least 0.429 W, leaving approximately 0.260 W available as work. The result is about 35.4 percent of the hot aperture's outgoing radiation power.

As T_h tends to T_0, net energy and net entropy transport vanish and the maximum work vanishes as well. A grey, directional or spectrally filtered source may have different relationships between intensity and entropy from these two ideal black streams. Using only a bolometer reading and a colour temperature to infer available work is therefore justified only if the assumed angular and spectral distributions are also credible.

A check on the limiting expression is to compare it with an ideal material heat engine fed at T_h. Such an engine can in principle extract a fraction $1 - T_0/T_h$ of a reversible heat input. In the optical problem the incoming stream already has a specified entropy-to-energy ratio. Delivering it to a converter cannot undo the entropy generated by producing that stream in the first place. This accounts for the smaller work fraction at the same two nominal temperatures and makes clear which boundary is being used when assigning a maximum work to emitted radiation.

Vibrations and the Debye heat capacity

Thermodynamics and statistical physics

Count vibrational modes in a solid and explain why a crystal's low-temperature heat capacity follows a cube law.

FROM FIRST PRINCIPLES

Derive the low-temperature heat capacity from acoustic modes

In a large three-dimensional crystal, long-wavelength acoustic vibrations have approximately linear dispersion $\omega = vk$. Count allowed wavevectors inside a sphere of radius k. Its volume grows as k^3, so the number of modes below ω grows as ω^3. The frequency density is its derivative. In a single-speed model normalized to the $3N$ displacements of N atoms,

$$g(\omega) = \frac{9N\omega^2}{\omega_D^3},$$

$$\int_0^{\omega_D} g(\omega)\, d\omega = 3N.$$

Each mode stores mean thermal energy $\hbar\omega/[e^{\hbar\omega/(k_BT)} - 1]$. Multiply by the mode count and substitute $x = \hbar\omega/(k_BT)$ to expose the temperature dependence,

$$U_{\mathrm{th}} = \frac{9N(k_BT)^4}{\hbar^3\omega_D^3} \int_0^{\hbar\omega_D/(k_BT)} \frac{x^3}{e^x - 1}\, dx.$$

When $T \ll \hbar\omega_D/k_B$, the omitted tail above the upper limit is exponentially small. The integral approaches a constant, giving $U_{\mathrm{th}} \propto T^4$ and $C_V = (\partial U/\partial T)_V \propto T^3$. The cube law relies on both three-dimensional mode counting and linear low-frequency dispersion.

DERIVATION AND APPLICATIONS

Recover the Debye velocity from distinct sound branches

A three-dimensional solid usually carries one longitudinal and two transverse acoustic polarizations. Their speeds can differ substantially. Consider a large elastically isotropic crystal of volume V, with linear dispersion $\omega = v_s k$ for each polarization s in the thermally populated range. Impose periodic boundaries; one wavevector state occupies reciprocal volume $(2\pi)^3/V$. For each polarization count all points in a sphere of radius ω/v_s, which includes opposite propagation directions. Differentiating gives the total density of modes,

$$N_s(<\omega) = \frac{V\omega^3}{6\pi^2 v_s^3},$$

$$g(\omega) = \frac{V\omega^2}{2\pi^2}\left(v_L^{-3} + 2v_T^{-3}\right).$$

A slower sound branch has more low-frequency states because the same frequency reaches a larger wavevector sphere. Define one effective velocity v_D to preserve this acoustic density of states. Requiring an isotropic monatomic Debye model to contain $3N$ vibrations, with number density $n = N/V$, fixes its frequency cutoff and temperature scale,

$$\frac{3}{v_D^3} = \frac{1}{v_L^3} + \frac{2}{v_T^3},$$

$$\omega_D = v_D (6\pi^2 n)^{1/3},$$

$$\Theta_D = \frac{\hbar \omega_D}{k_B}.$$

At low temperature the cutoff can be moved to infinity in the thermal energy integral. Changing variables to $x = \hbar\omega/(k_B T)$ uses $\int_0^\infty x^3/(e^x - 1)\, dx = \pi^4/15$. A temperature derivative yields the directly measurable coefficient

$$\frac{C_V}{V T^3} = \frac{2\pi^2 k_B^4}{15 \hbar^3} \left(v_L^{-3} + 2 v_T^{-3} \right).$$

The inverse cubes explain why replacing the three velocities by their arithmetic mean can noticeably underestimate the heat capacity. Optical or electronic contributions must be identified separately in specimens where they are thermally active.

WORKED EXAMPLE

Predict the acoustic temperature scale

Change the inputs and investigate this derivation.

For $v_L = 6000\,\mathrm{m/s}$, $v_T = 3000\,\mathrm{m/s}$, the longitudinal branch contributes only $1/17$ of the low-temperature coefficient. The effective velocity is

$$3000 (24/17)^{1/3} = 3365\,\mathrm{m/s}.$$

With $n = 5.00 \times 10^{28}\,\mathrm{m^{-3}}$, the cutoff wavevector is $1.44 \times 10^{10}\,\mathrm{m^{-1}}$, giving $\Theta_D \simeq 369\,\mathrm{K}$. The long-wave estimates apply well below that scale.

SECOND APPLICATION

Extract the transverse speed from calorimetry

Change the inputs and investigate this derivation.

Suppose a separate specimen gives $v_D = 2500\,\mathrm{m/s}$ from its low-temperature heat capacity and $v_L = 5000\,\mathrm{m/s}$ from a longitudinal pulse. Solving for the missing speed gives

$$v_T = [2/(3/v_D^3 - 1/v_L^3)]^{1/3} = 2.22 \times 10^3\,\mathrm{m/s}.$$

A shear-wave transit across $10.0\ \mathrm{mm}$ should consequently take $4.51\ \mu\mathrm{s}$. This timing test uses a distinct mechanical measurement of the same states counted by calorimetry.

Try it · Debye speed

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Longitudinal speed (m/s) **6000**

Transverse speed (m/s) **3000**

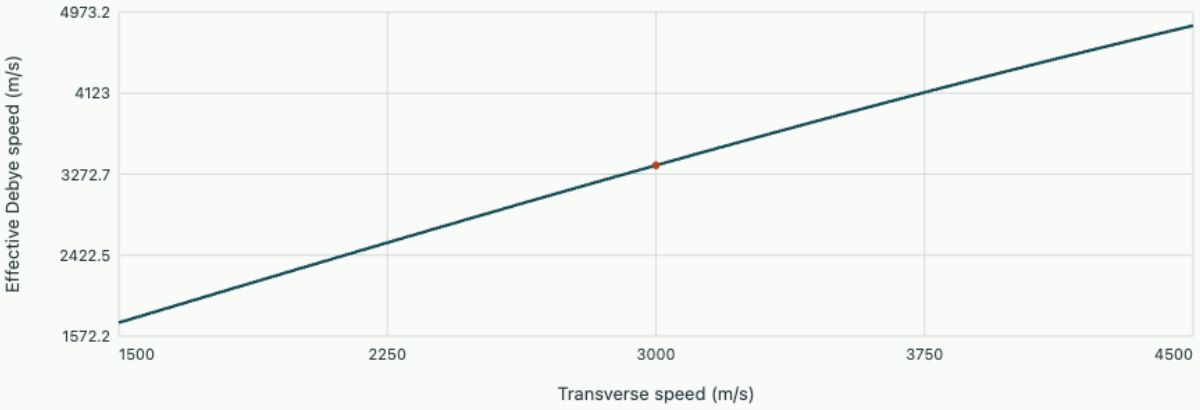

Three acoustic polarisations | Selected material

1/vD³=(1/vL³+2/vT³)/3. Longitudinal fraction of the low-temperature heat coefficient 0.058824.

What a solid can store

Consider an electrically insulating three-dimensional crystal of N atoms, volume V, sufficiently cold that its structure remains fixed and normal-mode frequencies vary weakly with temperature. In a harmonic approximation displacements about equilibrium separate into $3N$ independent oscillators. Their excitations are called phonons. Phonon number is not conserved when the solid exchanges heat, so each mode has a Bose mean occupation with zero chemical potential. The modes exist because the crystal can deform, including at very long wavelengths.

A model assigning the same frequency to every oscillator predicts an exponentially small low-temperature heat capacity, since exciting any mode would require a fixed nonzero energy. Acoustic vibrations instead approach zero frequency as their wavenumber approaches zero. Even in a cold macroscopic crystal, long-wavelength modes can remain thermally accessible. The key quantity is how many have energies of order $k_B T$ or less.

Test the single-frequency model before introducing a spectrum

The premise of the Einstein solid is deliberately restrictive. All $3N$ vibrational coordinates of a crystal are independent harmonic oscillators with one frequency ω_E. For one coordinate the levels are $E_n = \hbar\omega_E(n + 1/2)$. Write $x = \hbar\omega_E/(k_B T) = \Theta_E/T$. Summing the geometric series gives

$$Z_1 = e^{-x/2}/(1 - e^{-x}).$$

The partition function of the crystal is Z_1^{3N}, because each normal coordinate can be occupied independently and phonon number is free to change.

Apply $-\partial \ln Z/\partial\beta$, with $\beta = 1/(k_B T)$, to obtain

$$U = 3N\hbar\omega_E[1/2 + 1/(e^x - 1)].$$

The zero-point energy contributes no heat capacity at fixed volume. Differentiating the occupation term and using $dx/dT = -x/T$ gives

$$C_{V,E} = 3Nk_B \frac{x^2 e^{-x}}{(1 - e^{-x})^2}.$$

At high temperature $1 - e^{-x} \simeq x$, so the heat capacity per mole of atoms approaches $3R$. At low temperature the denominator tends to one, and the surviving factor $x^2 e^{-x}$ measures the cost of the lowest available excitation. For $\Theta_E = 240$ K the same mole would have $C_V = 18.06$ J/(mol K) at 120 K, 7.585 at 60 K and 0.536 at 30 K. Halving temperature does not give a fixed ratio. The curve crosses from a classical plateau to thermally rare excitations.

A real three-dimensional macroscopic crystal supports collective vibrations whose frequencies approach zero with wavelength increasing. Replacing the single frequency by a distribution puts modes below any fixed nonzero thermal scale. The next calculation finds the density of those modes, and therefore a low-temperature power law. High-temperature agreement alone cannot distinguish the two models since both count the same total number of coordinates.

Count to a cutoff

Use a continuum approximation with three acoustic polarizations and a common sound speed v. The modes in a shell of wavevectors grow as $k^2 dk$, as for photons. Real crystals generally have distinct longitudinal and transverse speeds. An effective Debye speed averages their inverse cubes. Converting via $\omega = vk$ gives density $g(\omega) = 3V\omega^2/(2\pi^2 v^3)$. The continuum would supply infinitely many modes, whereas a crystal with N atoms has exactly $3N$. Choose an upper cutoff ω_D to match that count,

$$\int_0^{\omega_D} g(\omega)\, d\omega = 3N,$$

$$g(\omega) = 9N\omega^2/\omega_D^3.$$

This cutoff is an approximate representation of the lattice at high wavenumber. The low-temperature prediction depends on the long-wave acoustic spectrum and is insensitive to precise details near the cutoff.

In a multi-atom cell additional optical branches change the intermediate-temperature behaviour, which cannot generally be reconstructed with a single effective speed.

Two limits and a measurement

The thermal energy is the integral of $g(\omega)\hbar\omega/[\exp(\hbar\omega/k_B T)-1]$ over the band. Define $\Theta_D = \hbar\omega_D/k_B$. When $T \ll \Theta_D$, extend the upper limit of the dimensionless integral to infinity. The resulting energy scales as T^4 and differentiation yields

$$C_V = \frac{12\pi^4}{5} N k_B \left(\frac{T}{\Theta_D}\right)^3.$$

At high temperature each mode instead approaches $k_B T$ of thermal energy and the heat capacity tends to $3Nk_B$. The acoustic model has the correct count by construction. The constant zero-point energy never enters either heat capacity because it disappears upon differentiation.

For one mole of atoms with $\Theta_D = 300$ K, the low-temperature prediction at 10.0 K is 0.0720 J mol⁻¹ K⁻¹, far below the classical $3R = 24.9$ J mol⁻¹ K⁻¹. At 20.0 K the cubic asymptote predicts 0.576 J mol⁻¹ K⁻¹. Measuring several temperatures and plotting C/T^3 tests the constant coefficient directly. In a conducting solid an electronic term proportional to T can dominate at low temperatures. Separating that term through a plot of C/T against T^2 allows the intercept and slope to reveal the two physically distinct energy stores.

Finite size eventually supplies a lower frequency cutoff related to the crystal dimensions, so the continuum counting argument presupposes a specimen large enough to have many wavelengths populated throughout the temperature interval.

Connect vibrations to thermal expansion

A purely harmonic crystal with force constants independent of volume can store vibrational energy yet has no vibrational mechanism for thermal expansion. Real vibrational frequencies usually vary as the lattice spacing changes. The quasiharmonic approximation treats each mode as harmonic at each volume while allowing its angular frequency $\omega_j(V)$ to change smoothly. Define the mode Grüneisen parameter $\gamma_j = -\partial \ln \omega_j / \partial \ln V$. A positive value means expansion softens that mode.

For one oscillator the Helmholtz free energy is

$$F_j = \hbar\omega_j/2 + k_B T \ln[1 - e^{-\hbar\omega_j/(k_B T)}].$$

Differentiation at fixed T gives

$$\partial F_j / \partial \omega_j = U_j / \omega_j,$$

where U_j includes the zero-point energy and the mean thermal excitation. Since vibrational pressure is $-\partial F/\partial V$, the chain rule gives

$$p_{\mathrm{vib}} = V^{-1} \sum_j \gamma_j U_j.$$

A temperature derivative at fixed volume suppresses the constant zero-point contribution. The mechanical relation $(\partial p/\partial T)_V = \alpha_V B_T$ finally yields

$$\alpha_V = \frac{\sum_j \gamma_j C_{V,j}}{B_T V}.$$

Here B_T is the isothermal bulk modulus and α_V the volume expansion coefficient. The mode sum weights each Grüneisen parameter by its thermally active heat capacity. A low-frequency mode that softens under compression can contribute negatively and can dominate in a cold crystal while high-frequency modes remain frozen. A single positive average parameter is therefore a material model rather than a thermodynamic necessity.

Infer a mode average from dilatometry

For a cubic crystal take measured molar heat capacity $C_{V,m} = 12.0$ J/(mol K), molar volume $V_m = 1.20 \times 10^{-5}$ m³/mol and bulk modulus $B_T = 100$ GPa. A measured linear expansion coefficient 6.00×10^{-6} K⁻¹ corresponds to $\alpha_V \simeq 3\alpha_L = 1.80 \times 10^{-5}$ K⁻¹. The effective heat-capacity-weighted parameter is then

$$\bar{\gamma} = \alpha_V B_T V_m / C_{V,m} = 1.80.$$

The factor of three comes from differentiating the cube of a length, not from counting acoustic polarizations again.

If the same positive average and bulk modulus remain approximately constant down to a regime where phonons give $C_V \propto T^3$, the volume expansion coefficient should also fall as T^3. Doubling temperature there increases the coefficient by about eight, although the accumulated length change requires integrating it over temperature. A comparison of α_V/C_V at several temperatures can test whether the active modes really share a common volume sensitivity. Molar heat capacity and molar volume must be taken on the same chemical counting basis before inferring any dimensionless Grüneisen parameter.

Hydrostatic pressure spectroscopy gives an independent way to measure a mode parameter. By the chain rule, the definition yields $\gamma_j = B_T(\partial \ln \omega_j/\partial p)_T$ within the quasiharmonic description. For the numerical bulk modulus, a mode with parameter 1.80 would increase its frequency by approximately 1.8 percent under a 1.00 GPa hydrostatic pressure increment. Agreement between this slope and the value inferred from expansion requires that the probed mode be representative of the thermally populated spectrum. A strong optical mode visible to a spectrometer need not dominate low-temperature heat capacity or expansion.

Quoted pressure derivatives must therefore specify the mode, the temperature and whether hydrostatic or directional stress was applied in the measurement.

Which lattice heat capacities survive at low temperature?

Compare a linear acoustic spectrum with a single oscillator frequency. Scale each temperature by its model’s characteristic temperature. The Debye prediction uses numerical quadrature of $3t^3 \int_0^{1/t} x^4 e^x/(e^x-1)^2\,dx$; both expressions approach unity at high temperature.

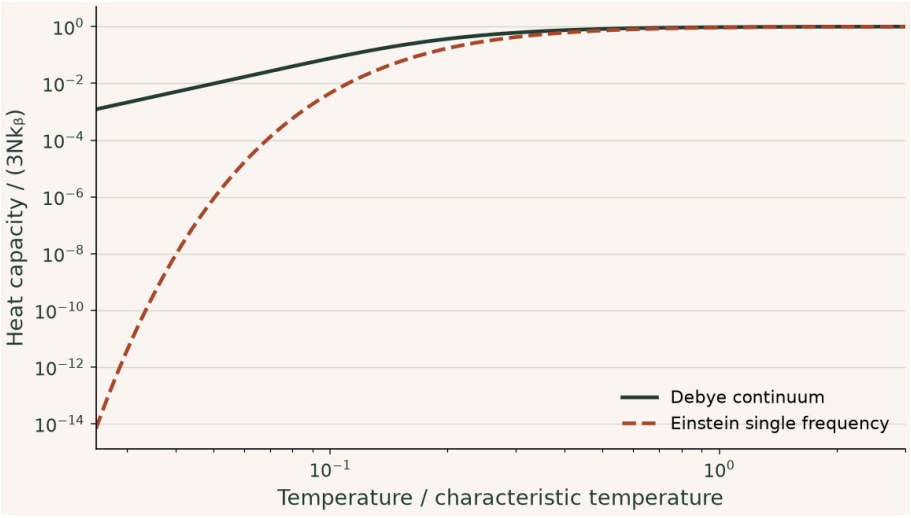

The single-frequency solid loses thermal excitations exponentially. The acoustic continuum still has low-frequency modes and its low-temperature curve follows a power law. Equal characteristic temperatures here set a comparison scale; fitting real data can give different values for the two models.

Try it · Einstein debye

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Characteristic temperature (K) **300**

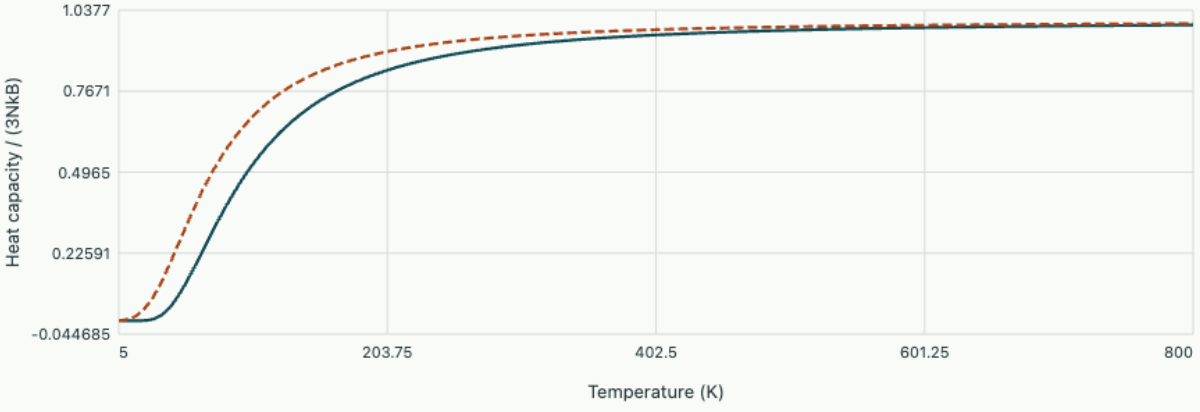

Einstein single-frequency model Debye acoustic continuum

Both approach the Dulong–Petit value one at high temperature. Debye retains low-frequency acoustic modes and gives a T^3 law at low temperature.

▶ **Read values from the plot**

Why only a narrow shell of a Fermi gas carries heat

Thermodynamics and statistical physics

Heating a filled Fermi sea does not give every particle a classical thermal energy. Deep below the chemical potential, nearby states are already occupied. Quantifying the accessible states predicts both the low-temperature heat capacity and the shift of the chemical potential at fixed particle number.

Convert a sharp Fermi surface into a narrow weight

For independent fermions with occupancy $f(E) = 1/[1 + \exp((E-\mu)/(k_B T))]$, consider integrals $I = \int_0^\infty H(E) f(E)\, dE$. Choose $F(E) = \int_0^E H(s) ds$. Integration by parts gives $I = \int_0^\infty F(E)[-\partial_E f(E)] dE$, since the boundary terms vanish. The derivative is a peak of width $k_B T$ around μ. With $x = (E-\mu)/(k_B T)$ the weight becomes

$$w(x) = 1/[4\cosh^2(x/2)].$$

For μ many thermal units above the bottom and H smooth over the thermal shell, extend the lower limit to $-\infty$ when expanding locally around μ. The weight has unit integral and is even, so the linear term in $F(\mu + k_B T x)$ integrates to zero. Its second moment is $\pi^2/3$, obtained by expanding $w(x) = \sum_{\ell \geq 1} (-1)^{\ell+1} \ell e^{-\ell x}$ on $x > 0$ and integrating x^2 termwise. Since $F'' = H'$,

$$I = \int_0^\mu H(E) dE + \frac{\pi^2}{6} (k_B T)^2 H'(\mu) + O(T^4).$$

The first correction samples the slope of the smooth weighting function. If a band edge or singularity lies within the shell, replacing the lower limit and Taylor expanding are poor approximations.

Conserve the number of particles

For a three-dimensional parabolic dispersion the density of states is $D(E) = A\sqrt{E}$. At $T = 0$, $N = (2A/3) E_F^{3/2}$. At small T, use the expansion with $H = D$ and write the chemical potential as $E_F + \delta\mu$, where $\delta\mu = O(T^2)$. Subtract the zero-temperature count. To quadratic order $0 = D(E_F)\delta\mu + (\pi^2/6)(k_B T)^2 D'(E_F)$. Because $D'/D = 1/(2E_F)$,

$$\mu/E_F = 1 - \frac{\pi^2}{12} (T/T_F)^2 + O[(T/T_F)^4],$$

$$T_F = E_F/k_B.$$

The increasing density of states offers more available levels just above the Fermi energy than just below it. The chemical potential shifts down to preserve the count. At $T/T_F = 0.10$ the predicted shift is $-0.00822 E_F$. For an ideal gas with $E_F = 5.0\,\mathrm{eV}$, this particular reduced temperature corresponds to about 5800 K, so reduced temperature and ordinary laboratory temperature should be distinguished.

Differentiate the energy at fixed number

Use $H = ED$ in the same integral, including the just-derived shift of μ. The boundary shift contributes $E_F D(E_F)\delta\mu$. In the Sommerfeld term the derivative is $D(E_F) + E_F D'(E_F)$. The terms involving the slope of D cancel after fixing number. Thus $U - U_0 = (\pi^2/6)D(E_F)(k_B T)^2$, and $C_V = (\pi^2/3)D(E_F)k_B^2 T$. In three dimensions $D(E_F) = 3N/(2E_F)$, giving $C_V/(Nk_B) = (\pi^2/2)T/T_F$. This is an electronic contribution for independent particles. Other degrees of freedom such as phonons can dominate a measured material heat capacity.

The experiment reconstructs μ by numerical integration and bisection at fixed number. It overlays the asymptotic prediction, allowing the fractional error to be checked before using it at larger T/T_F. The quadrature uses the momentum-like substitution $E/E_F = y^2$ to remove the square-root endpoint singularity.

Open this experiment online

Try the calculation

Reduced temperature T/TF **0.1**

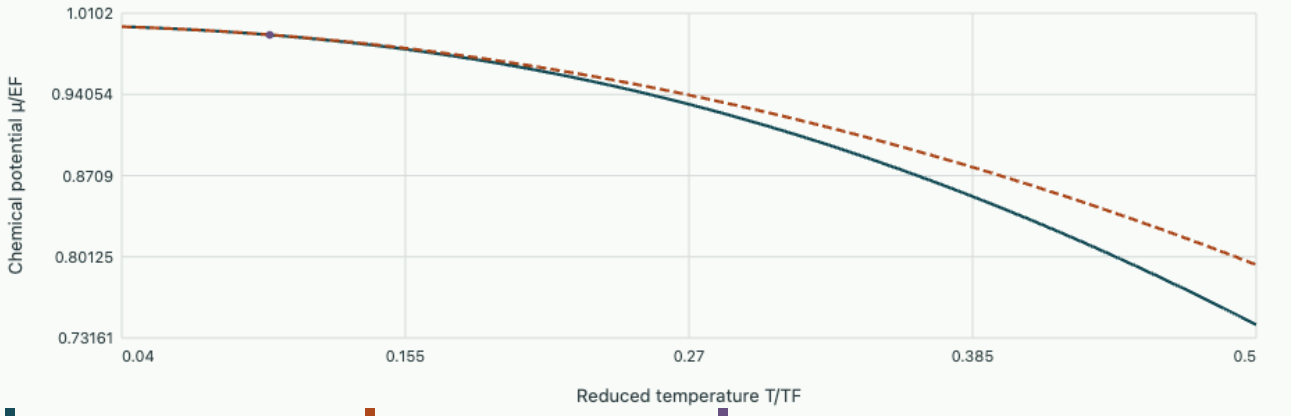

Fixed-number μ/EF = 0.9916, quadratic prediction 0.9918. Leading heat capacity $C_v/(Nk_\beta)$ = 0.4935.

Tutorial on interpreting heat-capacity data

For a gas with $E_F = 4.0\,\mathrm{eV}$ at $T = 300\,\mathrm{K}$, estimate $C_V/(Nk_B)$ and the chemical-potential change in eV. Use $k_B T = 0.02585\,\mathrm{eV}$.

▼ **Work through the solution**

The reduced temperature is 0.00646. The heat capacity ratio is approximately 0.0319, and $\delta\mu = -(\pi^2/12)(0.02585)^2/4.0 = -1.37 \times 10^{-4}\,\mathrm{eV}$.

A two-dimensional parabolic band has constant density of states above its edge. What does the quadratic expansion predict for the fixed-number chemical potential and energy?

▼ **Work through the solution**

$D' = 0$ makes the quadratic shift of μ vanish as long as the edge is distant. The energy integral still uses $(ED)' = D$, so it has a positive T^2 correction. Effects of the distant band edge can give exponentially small departures in μ.

Further reading · Cambridge statistical physics, degenerate Fermi gas and Sommerfeld expansion.

Filling electronic states and measuring the Fermi surface

Condensed matter and band theory

Use exclusion and state counting to predict carrier energies, thermal response and the effect of dimensionality.

FROM FIRST PRINCIPLES

Count the states in a filled Fermi sphere

Take independent spin-one-half electrons in a large cube of volume $V = L^3$ with periodic boundary conditions. The condition $e^{ik_x L} = 1$ permits $k_x = 2\pi n_x / L$, with the same spacing in the other directions. Thus each allowed wave vector occupies reciprocal volume $(2\pi)^3/V$. For the free-particle energy $E = \hbar^2 k^2/(2m)$, all states below energy E lie in a sphere. Counting its volume and two spin states gives

$$N(E) = 2\frac{V}{(2\pi)^3}\frac{4\pi}{3}\left(\frac{2mE}{\hbar^2}\right)^{3/2}.$$

At zero temperature the exclusion principle fills these single-particle states up to E_F. Writing $n = N/V$, solve the counting formula at that upper limit:

$$k_F = (3\pi^2 n)^{1/3},$$

$$E_F = \frac{\hbar^2}{2m}(3\pi^2 n)^{2/3}.$$

The number of available states in a thin energy shell is the derivative $D(E)dE$. In three dimensions $D(E) \propto \sqrt{E}$ because the shell area grows with k^2 while dk/dE falls as $1/k$. This counting assumes an isotropic parabolic band. Additional valleys or a different effective dimension change the result.

Background · Fermi level

$$f(E) = 1/[1 + \exp((E-\mu)/k_B T)]$$

At equilibrium a one-electron level at the chemical potential has mean occupancy one half per available state. Spin multiplicity enters the number of states separately.

Background · Density of states

$$g(E) = dN_{\text{states}}/dE$$

State density and state occupation are independent factors. Their product, integrated over a band, gives the mean particle number.

DERIVATION AND APPLICATIONS

Obtain the pressure of a filled Fermi sea by compressing its states

Even at zero temperature a box of mobile electrons exerts a pressure on its boundaries. Consider N noninteracting spin-one-half fermions in a three-dimensional volume V, with fixed isotropic mass m and quadratic energy. Take the bottom of the band as a volume-independent energy reference. Every state below E_F is occupied once per allowed spin state. Since the density of states including spin has the form $D(E) = C\sqrt{E}$, dividing two elementary integrals eliminates the normalization constant,

$$\frac{U}{N} = \frac{\int_0^{E_F} E\, C\sqrt{E}\, dE}{\int_0^{E_F} C\sqrt{E}\, dE} = \frac{3}{5} E_F.$$

State counting gives

$$E_F = (\hbar^2/2m)(3\pi^2 N/V)^{2/3}.$$

During a uniform compression at fixed N the upper occupied energy rises. At zero entropy the energy derivative with respect to box volume is the negative pressure, so

$$P = -\left(\frac{\partial U}{\partial V}\right)_N = \frac{2U}{3V} = \frac{2}{5} n E_F,$$
$$n = \frac{N}{V}.$$

Density appears both explicitly and through E_F. A second volume derivative gives the positive resistance to a small fractional compression,

$$B = -V\left(\frac{\partial P}{\partial V}\right)_N = \frac{5}{3} P = \frac{2}{3} n E_F.$$

Both P and the bulk modulus B have units of energy per volume. The factor $5/3$ distinguishes the change of pressure from pressure itself. An actual solid also has electrostatic and lattice contributions. Extracting the electronic contribution from a measured solid modulus requires information about these other stresses and the volume dependence of the band parameters.

WORKED EXAMPLE

Estimate the electronic stiffness of a simple metal

Change the inputs and investigate this derivation.

Take $n = 5.00 \times 10^{28}\,\mathrm{m^{-3}}$ and $E_F = 5.00\,\mathrm{eV}$. With $1\,\mathrm{eV} = 1.602 \times 10^{-19}\,\mathrm{J}$, the product nE_F is $4.01 \times 10^{10}\,\mathrm{Pa}$. The pressure contribution is $16.0\,\mathrm{GPa}$ and the modulus is $26.7\,\mathrm{GPa}$. A one percent volume decrease would raise this pressure by approximately $0.267\,\mathrm{GPa}$ in the linear response regime.

SECOND APPLICATION

Predict a finite compression without linearizing

Change the inputs and investigate this derivation.

Suppose a separate trapped gas starts at pressure $12.0\,\mathrm{kPa}$ and is compressed slowly to 0.800 of its original volume at fixed particle number and negligible entropy. Its pressure becomes

$$12.0(1/0.800)^{5/3} = 17.4\,\mathrm{kPa}.$$

Its energy per particle rises by only $(1/0.800)^{2/3} = 1.16$ times. The larger pressure ratio also contains the increase in particles per unit volume.

Try it · Fermi pressure

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Initial pressure (kPa) 12

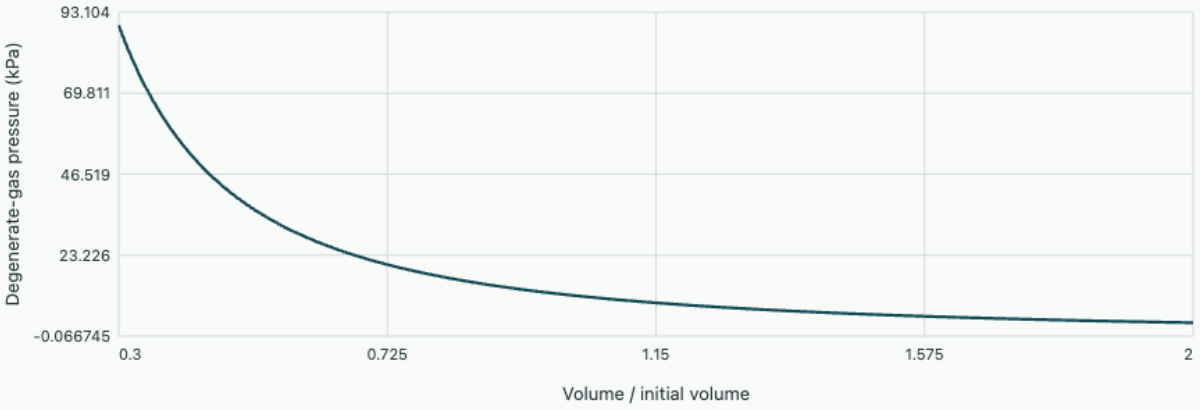

Zero-temperature 3D ideal Fermi gas

At fixed particle number P∝V^(−5/3), and the bulk modulus is 5P/3. Nonrelativistic, noninteracting fermions; thermal corrections omitted.

Count occupied states at zero temperature

Consider a cubic box of volume V with periodic boundaries and free electrons of effective model mass m. Allowed wavevectors occupy a volume $(2\pi)^3/V$ per state for each spin. At zero temperature each state

below the chemical potential is filled. For $\epsilon = \hbar^2 k^2 / 2m$ the filled region is a sphere and each k has two spin states. Hence

$$N = 2 \frac{V}{(2\pi)^3} \frac{4\pi k_F^3}{3},$$

$$k_F = (3\pi^2 n)^{1/3},$$

$$E_F = \frac{\hbar^2 k_F^2}{2m}.$$

With $n = 8.5 \times 10^{28}$ m⁻³ and the free electron mass, k_F is about 1.36×10^{10} m⁻¹ and E_F about 7.0 eV. The corresponding temperature E_F / k_B is roughly 81,000 K. Ordinary room temperature therefore changes occupation only near the outer surface of this sea in a simple metal. This is a model estimate; anisotropic bands and multiple pockets change its interpretation.

The density of states is a derivative

Replacing k_F by $k(E) = \sqrt{2mE}/\hbar$ gives the integrated state count below energy E. Its derivative is

$$D(E) = \frac{V}{2\pi^2} \left(\frac{2m}{\hbar^2} \right)^{3/2} \sqrt{E},$$

$$D(E_F) = \frac{3N}{2E_F}.$$

Check dimensions by integrating $D(E)\,dE$ to recover a count. In two dimensions the occupied area grows as k^2 and the density per energy is constant for a parabolic band. In one dimension the count grows as k and the density has an integrable $E^{-1/2}$ singularity at its bottom. A large density of states means many states fit in a narrow energy window. It does not imply a large group velocity: a flat dispersion may create the large density.

OneSimulation · What would a probe see at the last atom?

An open one-dimensional chain with one orbital per atom has standing-wave eigenvectors. Counting their energies gives a spectrum for the complete specimen. An ideal local probe couples to the amplitude on the atom below it. For a probe over site j, broaden each spectral line with a normalised Gaussian g_η of experimental width η.

$$\rho_j(E) = \sum_n |\psi_n(j)|^2 g_\eta(E - E_n),$$

$$\int \rho_j(E)\,dE = 1.$$

The integral counts one orbital per spin at either site. Near a band edge the standing waves vary slowly and have a small amplitude next to the missing neighbour. Their large bulk density of states can coexist with a weak end signal. The OneSimulation eigensolve uses the same Hamiltonian for both curves. Only the projected site changes.

```
"""Local spectroscopy at an open chain end and in its interior."""
import matplotlib.pyplot as plt
import numpy as np
from onesim.numerics.linear.matrices.dense_linear_algebra import symmetric_eigensystem

sites = 81
hopping = 1.0  # eV
broadening = 0.09  # eV, standard deviation of the energy resolution
hamiltonian = -hopping * (np.eye(sites, k=1) + np.eye(sites, k=-1))
levels, vectors = symmetric_eigensystem(hamiltonian)
energy = np.linspace(-2.6 * hopping, 2.6 * hopping, 1001)
kernel = np.exp(-0.5 * ((energy[:, None] - levels) / broadening)**2)
kernel /= np.sqrt(2 * np.pi) * broadening
fig, ax = plt.subplots(figsize=(6.6, 3.5), layout="constrained")
for site, label in [(0, "End site"), (sites // 2, "Centre site")]:
    weights = np.abs(vectors[site])**2
    density = kernel @ weights
    assert np.isclose(np.trapezoid(density, energy), 1, atol=1e-6)
    print(f"{label}: local density at zero energy = {density[len(energy)//2]:.3f} /eV")
    ax.plot(energy, density, label=label)
ax.set(xlabel="Energy (eV)", ylabel="Local density per site and spin (1/eV)")
ax.grid(alpha=0.2)
ax.legend()
fig.savefig("condensed-matter.png", dpi=180)
```

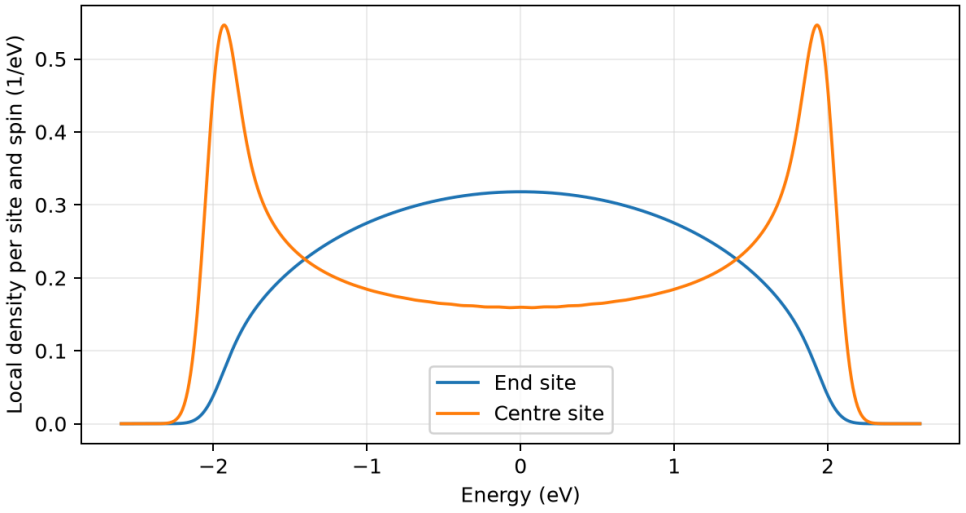

The integrated areas agree. Near zero energy the end site carries about 0.318 states/eV per spin and the centre carries 0.160. A change in probe location redistributes spectral weight.

Try it · Local dos

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Probe site in 40-site open chain **1**

Lorentzian broadening (eV) **0.08**

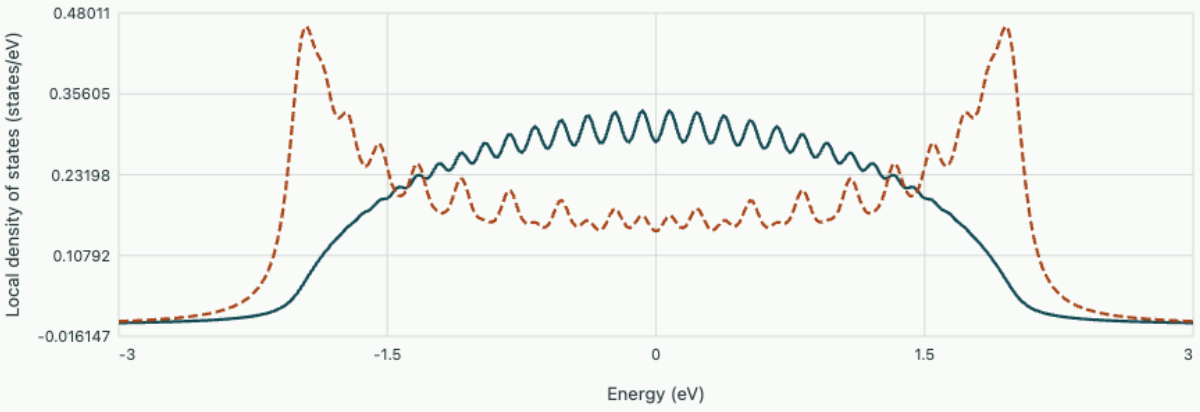

Selected site Central site

Nearest-neighbour hopping is 1 eV. Each full-energy spectral integral is one state per spin; finite plotted limits omit Lorentzian tails.

Download the Python example

Why only a shell contributes to heat capacity

At finite temperature the occupation is $f(E) = [1 + e^{(E-\mu)/k_BT}]^{-1}$. Heating can promote an electron only into a vacant state. Deep in the filled sea all nearby destinations are occupied. The active shell has width of order k_BT, involving about $D(E_F)k_BT$ states; the typical energy change is another factor k_BT. Thus the thermal energy increases quadratically and the heat capacity rises linearly with T.

A coefficient follows by expanding the smooth density of states across that shell. At fixed particle number, the relevant excess energy to leading order is $\int (E - E_F)D(E_F)[f(E) - \Theta(E_F - E)]\,dE$. The filled and empty sides contribute equally to this order. Substituting $x = |E - E_F|/k_BT$ and using $\int_0^\infty x/(e^x + 1)\,dx = \pi^2/12$ gives

$$C_{\mathrm{el}} = (\pi^2/3)D(E_F)k_B^2T.$$

This expansion requires the density to be smooth on the thermal scale; a band edge at the chemical potential needs separate treatment.

Separate two low-temperature contributions

In a three-dimensional simple metal combine $C_{\mathrm{el}} = \gamma T$ with the phonon result $C_{\mathrm{ph}} = \beta T^3$. Plot C/T against T^2; the intercept estimates γ, and the slope estimates β. If two points are $(4, 3)$ and $(16, 6)$ in K² and mJ mol⁻¹ K⁻², respectively, $\beta = 0.25$ mJ mol⁻¹ K⁻⁴ and $\gamma = 2.0$ mJ mol⁻¹ K⁻². Curvature across the fitted interval

would be evidence against this two-term description. To go beyond a free sphere we next retain the periodic potential of the ions.

Count charges and electronic compressibility in two dimensions

A two-dimensional electron layer illustrates how changing available spatial dimensions alters the density of states and an electrical measurement. Assume one parabolic subband $E = E_c + \hbar^2(k_x^2 + k_y^2)/(2m^*)$, spin degeneracy two and no valley degeneracy. A square sample of area A with periodic boundary conditions supplies one state per spin per reciprocal area $(2\pi)^2/A$. Counting the disk up to k_F gives areal density $n_s = k_F^2/(2\pi)$, hence

$$E_F - E_c = \frac{\pi \hbar^2 n_s}{m^*},$$

$$D_A(E) = \frac{dn_s}{dE_F} = \frac{m^*}{\pi \hbar^2}$$

above the subband edge at zero temperature. Unlike the square-root density in three dimensions, each occupied parabolic two-dimensional subband contributes a constant step. When an additional subband begins to fill it supplies another step, possibly with different degeneracy or mass. Temperature and disorder broaden the step in a real measurement.

An electrostatic gate must both create a field across the insulating spacer and raise the layer’s chemical potential. For areal capacitance C_g' of the spacer and increment of sheet charge magnitude $dQ' = e\,dn_s$, the required voltage is

$$dV = dQ'/C_g' + dE_F/e.$$

The second term defines the quantum capacitance $C_q' = e^2 D_A$. Their inverse capacitances add in series. For a small density of states the electronic contribution measurably lowers the total capacitance even when the geometric dimensions are known precisely.

Measure which capacitance limits a device

Take $m^* = 0.067 m_e$ and $n_s = 2.00 \times 10^{15}$ m⁻². Then $k_F = \sqrt{2\pi n_s} = 1.12 \times 10^{8}$ m⁻¹ and $E_F - E_c = 7.15$ meV. At 4 K the thermal scale $k_B T = 0.345$ meV is smaller than this separation, supporting the degenerate estimate. The electronic capacitance per area is 0.0448 F/m². For a spacer with $C_g' = 0.0100$ F/m², the measured series value is $(1/0.0100 + 1/0.0448)^{-1} = 0.00818$ F/m².

Over a gate area of 100 µm², the predicted capacitance is 0.818 pF rather than the purely geometric 1.00 pF. Treating that measured value as a change of dielectric thickness would misidentify an electronic effect. In the limit of a very high density of states, the quantum capacitance becomes large and the geometric value is recovered. Near depletion the assumptions of a sharp occupied subband and uniform layer require scrutiny. Comparing capacitance with an independent sheet-density measurement can test whether a low reading reflects depleted electronic states or simply a different insulating spacer.

Using the series result, a small gate swing can be split into two measurable energy changes. A 10.0 mV swing at the stated operating point changes sheet charge magnitude by 8.18×10^{-5} C/m², or sheet density by

5.10×10^{14} m⁻². The geometric voltage increment is 8.18 mV and the chemical-potential contribution is 1.82 mV, corresponding to a Fermi-level shift of 1.82 meV. Their sum recovers the imposed swing. This decomposition assumes the layer remains in the same subband and equilibrates through its electrical contact on the capacitance measurement timescale.

A slow measurement with significant leakage instead mixes capacitance with conductance and can corrupt both estimated voltage contributions.

Count states added by raising the Fermi level

For a parabolic band, states within a radius k fill a length, area or volume in reciprocal space. Since k scales as $\sqrt{E}$, the integrated count varies as $E^{d/2}$. Normalise at one energy to compare exponents independently of the sample dimensions.

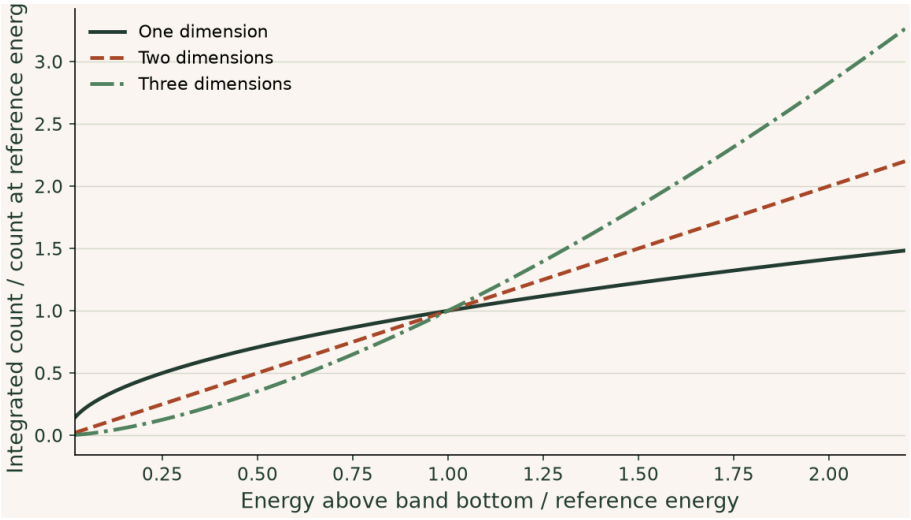

Doubling the Fermi energy raises the count by √2, 2 or 2√2. The local slope gives the density of states. Its behaviour at the band edge is quite different even when the three integrated counts agree at the reference energy.

Try it · State count

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Selected energy / reference

1

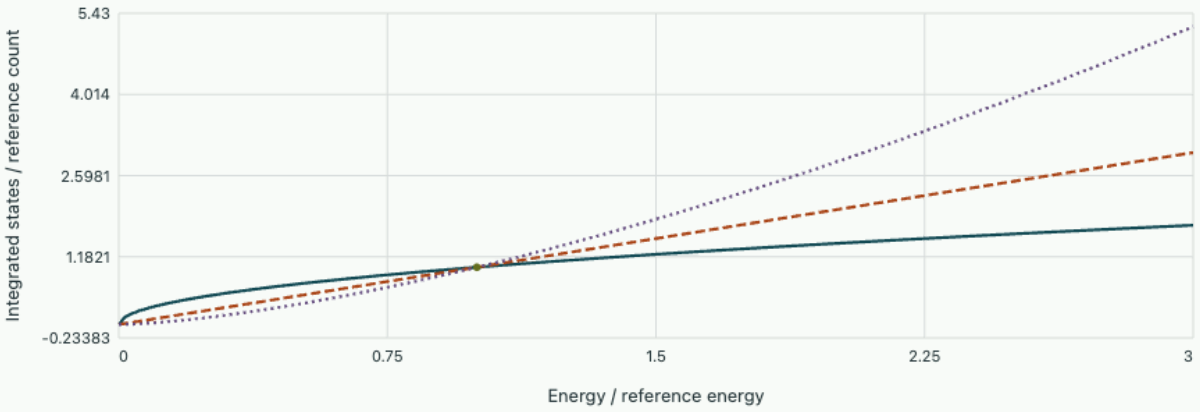

At selected energy: 1D 1, 2D 1, 3D 1. Differentiating these curves gives the corresponding density of states.

▶ **Read values from the plot**

Magnetic moments and susceptibility

Condensed matter and band theory

Derive thermal magnetisation for a two-state moment and distinguish its signature from the response of a filled electronic sea.

FROM FIRST PRINCIPLES

Obtain the Curie law by counting two magnetic orientations

Take many identical, independent localized moments per volume n, each with component $+\mu$ or $-\mu$ along an applied induction B. The magnetic energies are $-\mu B$ and $+\mu B$. At temperature T their probabilities are proportional to their Boltzmann factors. With $x = \mu B/(k_B T)$, normalize the two weights and subtract their contributions to the mean:

$$Z = e^{x} + e^{-x},$$

$$\langle m_z \rangle = \mu \frac{e^{x} - e^{-x}}{e^{x} + e^{-x}} = \mu \tanh x.$$

Multiply by number density to get magnetization. In weak fields or at high temperatures $x \ll 1$, so the leading term in $\tanh x$ gives

$$M \simeq \frac{n\mu^2 B}{k_B T},$$

$$\chi = \frac{\partial M}{\partial H} \simeq \frac{\mu_0 n \mu^2}{k_B T}.$$

The susceptibility expression uses a weakly magnetic sample with $B \simeq \mu_0 H$ and negligible demagnetizing correction. At low temperature and fixed nonzero field both the expansion and the $1/T$ extrapolation fail. The exact two-state average saturates at $n\mu$ because no moment can contribute more than μ. Exchange between neighbours can also invalidate the independent weights.

DERIVATION AND APPLICATIONS

Distinguish the saturation moment from the Curie moment

An ion with angular momentum quantum number J can have more than two projections on a magnetic field. Consider independent identical ions of number density n in thermal equilibrium, with isotropic Landé factor g and negligible splitting at zero field. Label the levels $m = -J, \ldots, J$ so that their moments along the applied induction B are $g\mu_B m$. Adjacent Zeeman levels are then separated by $g\mu_B B$. Introduce $x = g\mu_B B/(k_B T)$. Weighing each level with its Boltzmann factor gives

$$Z = \sum_{m=-J}^{J} e^{mx},$$

$$M = n g \mu_B \frac{\partial \ln Z}{\partial x}.$$

In weak fields each zero-field state is almost equally likely. The zeroth-order sum of m vanishes by symmetry. The next term in the numerator is $x \sum m^2$, while the denominator is $2J+1$ to this order. Using the sum of squares of the allowed projections gives

$$\langle m \rangle \simeq x \frac{\sum_{m=-J}^{J} m^2}{2J+1} = \frac{J(J+1)}{3} x.$$

For a dilute specimen take $B \simeq \mu_0 H$. The dimensionless volume susceptibility is the derivative of M with respect to magnetic field strength H. In strong fields the highest projection alone survives. These limits identify two different combinations of the microscopic parameters,

$$\chi = \frac{\mu_0 n (g\mu_B)^2 J(J+1)}{3 k_B T},$$

$$M_{\mathrm{sat}} = n g \mu_B J.$$

The Curie-law effective moment is $g\mu_B \sqrt{J(J+1)}$, whereas the aligned projection is $g\mu_B J$. Their ratio approaches one for large J. A susceptibility fit requires $Jx \ll 1$; a saturation measurement requires the opposite tendency and sufficiently weak interactions between moments.

WORKED EXAMPLE

Predict the response of dilute spin-one ions

Change the inputs and investigate this derivation.

Take $J = 1$, $g = 2.00$,

$$n = 1.00 \times 10^{26}\,\mathrm{m^{-3}},$$

$$T = 10.0\,\mathrm{K}.$$

The aligned magnetization is $1.85 \times 10^{3}\,\mathrm{A/m}$. The Curie formula gives $\chi = 2.09 \times 10^{-4}$. At $B = 0.100\,\mathrm{T}$, $Jx = 0.0134$, and the predicted magnetization is only $\chi B/\mu_0 = 16.6\,\mathrm{A/m}$, consistent with weak polarization.

SECOND APPLICATION

Identify a moment from two calibrated measurements

Change the inputs and investigate this derivation.

Suppose independent concentration measurements in another material yield a saturation moment of $3.00\mu_B$ per ion and an effective Curie moment of $3.87\mu_B$. The squared ratio is $1 + 1/J = (3.87/3.00)^2 = 1.66$,

giving $J \simeq 1.50$. The saturation value then gives $g \simeq 2.00$. Susceptibility alone fixes only a product of the factor g and the size of the angular momentum.

Try it · Spin magnetisation

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Spin quantum number J **1**

Temperature (K) **10**

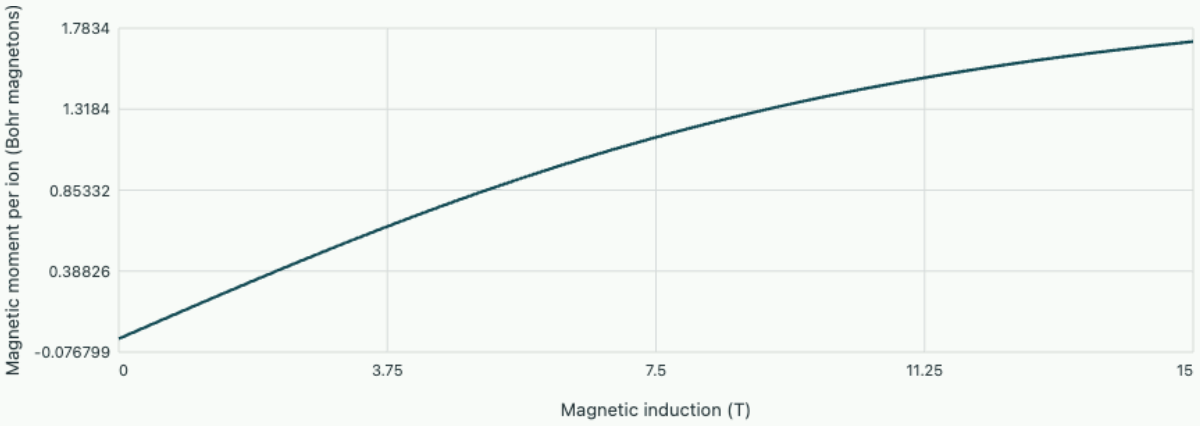

Exact spin partition sum, g=2

Saturation moment 2J=2 μB; Curie effective moment 2√(J(J+1))=2.8284 μB. Independent local moments.

A moment has an energy as well as a direction

Consider N independent identical magnetic moments in volume V, each with allowed projection $\pm\mu$ along a uniform applied induction B. The energies are $\mp\mu B$; the state aligned with the field has lower energy. For $\beta = 1/(k_B T)$ the one-moment partition function is $Z_1 = 2\cosh(\beta\mu B)$. From $\langle m_z \rangle = k_B T\, \partial_B \ln Z_1$ obtain

$$M = n\mu \tanh(\mu B / k_B T),$$

$$n = N/V.$$

The full curve approaches $n\mu$ at high field or low temperature. In the weak-field regime $\mu B \ll k_B T$, expansion of the hyperbolic tangent gives

$$M \simeq n\mu^2 B/(k_B T).$$

Under weak magnetisation and negligible demagnetising corrections use $B \simeq \mu_0 H$; the dimensionless SI susceptibility $\partial M/\partial H$ is then

$$\chi = \mu_0 n \mu^2/(k_B T).$$

Be explicit about differentiating with respect to B or H, which carry different units.

Work a moment count from two measurements

Suppose a dilute two-state sample exhibits saturation magnetisation 9.27×10^3 A/m and the moments are known to have $\mu = \mu_B = 9.27 \times 10^{-24}$ J/T. The implied density is 1.00×10^{27} m⁻³. At 10 K and 1 T the argument $\mu B/k_B T$ is 0.0672, safely in the linear range for a rough estimate: $M \simeq 622$ A/m. At the same temperature a field of about 14.9 T gives argument one and $M/M_{\mathrm{sat}} = 0.762$. Fitting a straight line up to that field would underestimate the zero-field slope.

If the moment magnitude is unknown, saturation determines $n\mu$ and the small-field slope supplies $n\mu^2$. Taken together they can separate density from moment within this model. A mixture of moment species, orbital contributions or interactions can defeat this two-equation inference. Plotting curves at several temperatures against B/T tests the independent-moment scaling without needing to know the overall density beforehand.

Compare with itinerant spin response

For a weakly perturbed degenerate metal, the two spin species shift by $\mp \mu B$. Almost all states below E_F remain doubly occupied, so only a thin shell at the Fermi surface changes its net spin. At low temperature and nearly constant total density of states $D_V(E_F)$, the spin imbalance is approximately $D_V(E_F)\mu B$, giving $M \simeq \mu^2 D_V(E_F) B$. The leading Pauli response is nearly temperature independent, in contrast to the $1/T$ scaling of independent localised moments. This argument addresses spin susceptibility; orbital diamagnetism also contributes to a real measured susceptibility.

Diagnose a small impurity contribution

If bulk data show a roughly constant offset plus a low-temperature upturn, test a simple additive model $\chi = \chi_0 + C/T$ over a weak-field interval. Plot χ against $1/T$ rather than taking an inverse before subtracting the offset. An upturn can arise from a dilute impurity population even when the host electronic sea dominates heat capacity. If scaling fails systematically, examine interactions between moments; the next chapter shows how cooperative ordering changes the small-field susceptibility.

Cool local moments by adiabatic demagnetization

Independent moments have an entropy that depends on the ratio of magnetic field to temperature. For a two-state moment of magnitude μ, energies are $\mp \mu B$. Define $x = \mu B/(k_B T)$. The one-moment partition function is $2 \cosh x$, average energy is $-\mu B \tanh x$, and the magnetic entropy follows from $S = k_B \ln Z + U/T$,

$$\frac{S_{\mathrm{spin}}}{N k_B} = \ln(2\cosh x) - x\tanh x.$$

At zero field for fixed positive temperature the moment has two equiprobable orientations and contributes $k_B \ln 2$. At large x nearly all moments favour the lower state and the spin entropy tends to zero. If field is increased at fixed temperature, this reduction of spin entropy expels heat to a bath. Now isolate the material and reduce the field quasistatically. In the ideal spin-dominated limit entropy remains fixed, so x remains fixed and T/B is constant.

Thermal isolation and sufficiently slow variation are both required. The first prevents external heat from entering, and the second allows internal degrees of freedom to follow equilibrium. Other stores of entropy, particularly lattice vibrations or mobile electrons, participate in an actual specimen. The pure spin scaling can then be modified even when the process as a whole is perfectly adiabatic.

Estimate temperature and cooling capacity

Take $N = 1.00 \times 10^{21}$ independent moments each of magnitude μ_B, initially at 2.00 K. Magnetize isothermally from negligible field to 3.00 T, using $\mu_B/k_B = 0.6717$ K/T. The final magnetic parameter is $x = 1.01$, giving spin entropy per particle about $0.362 k_B$ compared with the initial $0.693 k_B$. The heat passed to the bath in that reversible isothermal step is $T N k_B (0.693 - 0.362) = 9.14$ mJ. Work is supplied by the field source during preparation.

Isolate and reduce field from 3.00 T to 0.300 T. Holding x fixed predicts

$$T_f = 2.00 (0.300/3.00) = 0.200\,\mathrm{K}.$$

Further field reduction in the ideal noninteracting equations would approach zero temperature. Real moments interact, however, and internal effective fields or ordering lift their near-zero-field degeneracy. Thermal leaks and finite relaxation rates also become increasingly important. A practical test records temperature against field on several ramp rates; dependence on the rate can distinguish irreversible heating from a reproducible correction to the equilibrium entropy model. The final cold sample can absorb only finite heat before its moments become disordered, even if the predicted endpoint temperature is very low.

A second check comes from magnetization along the ideal isolated sweep. Since $M = N\mu \tanh x$ at fixed volume and x is conserved, the spin magnetization should stay almost constant while both the external field and the temperature fall. In contrast, an isothermal field decrease reduces x and magnetization. Measuring temperature alone cannot establish that the sample really followed an isentropic spin path. A simultaneous magnetic reading tests the fixed-population interpretation. Deviations concentrated at the weakest fields can signal interactions between moments rather than a general failure of the high-field entropy estimate.

Pairing, excitation gaps and what a spectrum measures

Condensed matter and band theory

Use a two-component excitation model to predict spectroscopic signatures, then distinguish three physically different gaps encountered across the course.

FROM FIRST PRINCIPLES

See how pairing opens a spectroscopic gap

Begin with a single normal-state band and let ξ be its energy measured from the chemical potential. A minimal mean-field model for a uniform paired state mixes an electron amplitude with a hole amplitude through a complex pairing energy Δ. The two diagonal energies have opposite signs. To find its stationary excitations write the eigenvalue condition:

$$H = \begin{pmatrix} \xi & \Delta \\ \Delta^* & -\xi \end{pmatrix}, \qquad \det(H - EI) = (\xi - E)(-\xi - E) - |\Delta|^2 = 0.$$

It follows that the two branches are $E_\pm = \pm\sqrt{\xi^2 + |\Delta|^2}$. At the former Fermi surface $\xi = 0$, even the cheapest positive excitation costs $|\Delta|$. To predict a tunnelling spectrum for constant isotropic gap and nearly constant normal density of states, count the normal states mapping into an energy interval dE. The Jacobian from $\xi = \pm\sqrt{E^2 - |\Delta|^2}$ gives the normalized density, with the electron and hole weights combined:

$$\frac{D_s(E)}{D_n(0)} = \frac{|E|}{\sqrt{E^2 - |\Delta|^2}} \quad (|E| > |\Delta|).$$

There are no states within the ideal gap at this level of description. Finite temperature, lifetime broadening and instrumental resolution change the measured conductance near its edges. The model assumes a nonzero Δ; it predicts consequences of pairing without deriving the attractive interaction that supports it.

DERIVATION AND APPLICATIONS

Predict the weights of a particle peak and its partner

A pairing gap determines excitation energies, but a momentum-resolved spectrum also measures how strongly an electron couples to each peak. Consider a clean, uniform superconductor with a real positive gap Δ, at temperature low enough to neglect thermal broadening. Measure the normal-state energy ξ at one wavevector from the chemical potential. In the electron and hole basis the positive-energy eigenvector has components u, v. The first row of its eigenvalue equation requires

$$(\xi - E)u + \Delta v = 0,$$

$$\frac{v}{u} = \frac{E - \xi}{\Delta},$$

$$E = \sqrt{\xi^2 + \Delta^2}.$$

Normalize the vector to unit probability. Since $(E-\xi)(E+\xi)=\Delta^2$, the squared amplitude ratio can be written without subtracting two nearly equal energies. The resulting coherence factors are

$$|u|^2=\frac{1}{1+(E-\xi)/(E+\xi)}=\frac{1}{2}\left(1+\frac{\xi}{E}\right),$$
$$|v|^2=\frac{1}{2}\left(1-\frac{\xi}{E}\right).$$

The negative-energy eigenvector is orthogonal to the positive one and consequently has electron weight $|v|^2$. In the ideal single-particle spectral function at this momentum the two sharp peaks have unit total weight,

$$A(\epsilon)=|u|^2\delta(\epsilon-E)+|v|^2\delta(\epsilon+E),$$
$$\int A(\epsilon)\,d\epsilon=1.$$

Thus the weights exchange as the wavevector passes through the normal Fermi surface. At $\xi=0$ the two weights are equal. Well above the surface the positive-energy peak carries almost all the weight. To compare both sides quantitatively an experiment must account for occupation and matrix elements; ordinary photoemission alone samples the occupied side most directly.

WORKED EXAMPLE

Resolve a weak partner peak

Change the inputs and investigate this derivation.

At a selected momentum let $\xi=3.00\,\mathrm{meV}$ and $\Delta=4.00\,\mathrm{meV}$. The spectral energies are $+5.00$ and $-5.00\,\mathrm{meV}$, with respective electron weights 0.800 and 0.200. At a momentum with the opposite ξ and the same gap the two peak energies stay fixed while their weights reverse. Such a weight exchange tests particle–hole mixing more specifically than the peak separation alone.

SECOND APPLICATION

Reconstruct a gap using corrected peak areas

Change the inputs and investigate this derivation.

In a different sample the peaks are separated by $6.00\,\mathrm{meV}$, and corrected electron spectral areas have positive-to-negative ratio $3:2$. Their normalized weights are 0.600 and 0.400, so

$$\xi/E=0.200.$$

With $E=3.00\,\mathrm{meV}$ this gives $\xi=0.600\,\mathrm{meV}$ and $\Delta=\sqrt{3.00^2-0.600^2}=2.94\,\mathrm{meV}$. Half the separation overestimates the gap here because the probed momentum lies away from the normal Fermi surface.

Try it · Bcs weights

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Gap (meV) 4

Normal-state energy ξ (meV) 3

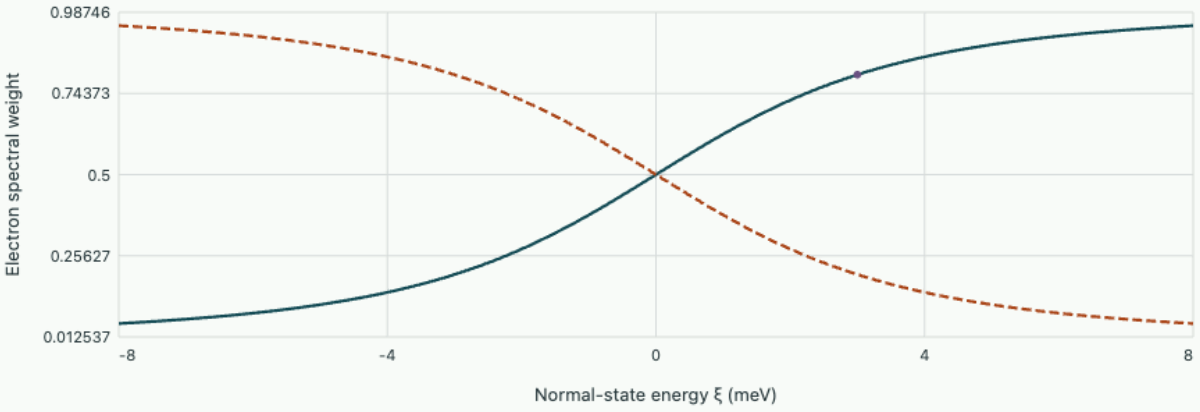

Positive-energy electron weight u² Negative-energy electron weight v²

Selected positive peak

Peak energies ±5 meV; weights sum to one. The measured electron spectrum also needs occupation and instrumental factors.

Build excitations on a paired state

For a conventional fully gapped superconductor, a mean-field description mixes particle and hole amplitudes near the Fermi surface. With $\xi_k = \epsilon_k - \mu$ and gap parameter Δ, a simplified matrix for a time-reversed pair of states is

$$\begin{pmatrix} \xi_k & \Delta \\ \Delta^* & -\xi_k \end{pmatrix} \begin{pmatrix} u_k \\ v_k \end{pmatrix} = E \begin{pmatrix} u_k \\ v_k \end{pmatrix}, \qquad E_{\pm} = \pm\sqrt{\xi_k^2 + |\Delta|^2}.$$

A positive-energy quasiparticle has minimum energy $|\Delta|$ where the normal band crossed the chemical potential. Breaking a pair into two excitations can cost at least $2|\Delta|$. These are different thresholds, so the probe and transition being discussed must be specified. The matrix displays level repulsion, reminiscent of the nearly free electron calculation, but its basis states and the physical origin of the coupling are different.

Transform a density of states

Suppose the normal-state density is approximately constant over the gap scale and the gap is isotropic. For $E > |\Delta|$,

$$|\xi| = \sqrt{E^2 - |\Delta|^2},$$

$$d|\xi|/dE = E/\sqrt{E^2 - |\Delta|^2}.$$

Counting the paired branches relative to the normal-state convention yields $D_s(E)/D_n = |E|/\sqrt{E^2 - |\Delta|^2}$ outside the gap and zero inside in the ideal zero-broadening model. The edge singularity is integrable. States removed from the gap reappear near its edges. Thermal broadening, disorder, nodes or an anisotropic gap can change the measured shape.

A low-temperature tunnel junction with a normal counterelectrode of nearly constant density can use differential conductance as a probe of this spectrum, subject to matrix-element and broadening effects. Peaks at biases ± 1.5 mV suggest $|\Delta| \simeq 1.5$ meV in this geometry. Peak separation on the voltage axis is 3.0 mV. Using the weak-coupling BCS ratio $2|\Delta(0)|/(k_B T_c) \simeq 3.53$ gives $T_c \simeq 9.86$ K. Measure T_c independently; agreement with this estimate supports the simple model, while deviation calls for closer investigation.

Select a second observation that distinguishes the mechanism

A nearly free electron band gap separates eigenvalues near a zone boundary. For conduction it also matters where the chemical potential sits. A superconducting gap opens at the Fermi surface on entering an ordered state, and has magnetic screening and a characteristic temperature dependence. A photonic stop band instead excludes real propagating wavevectors in a frequency range for the specified periodic optical medium. All three can produce missing spectral weight or reduced transmission. Their energies, frequency units, occupation rules and control parameters differ.

Return to the numerical optical stack with a prediction: making both layers identical closes its first stop band even though the artificial zone boundary remains. For a proposed superconducting gap, repeat the tunnelling and magnetisation measurements as temperature crosses the transition. For an electronic band edge, relate spectroscopy to filling, dispersion and transport. This combination of an analytic limit, a controlled change and an independent observable is what makes an interpretation of a spectrum persuasive.

For continued study, see the MIT lectures on Theory of Solids, including quasiparticles and superconducting electrodynamics.

Read the phase difference across a superconducting weak link

A single-particle tunnelling spectrum can reveal an excitation gap. A weak junction between two superconductors can also carry a coherent pair current that responds directly to their relative condensate phase. Use the gauge-invariant phase difference φ, including the vector-potential line integral across the junction. For a simple tunnel junction in the low-transparency regime, its leading coupling energy is

$$U_J = -E_J \cos\varphi.$$

Its 2π periodicity reflects invariance under an integer phase winding and the energy minimum at aligned phase for a conventional junction.

A potential difference V between the condensates causes their relative phase to advance at $\dot{\varphi} = 2eV/\hbar$, choosing a consistent orientation of current and voltage. Combining electrical power IV with the rate of change of junction energy gives $IV = E_J \sin\varphi\, \dot{\varphi}$, hence

$$I = I_c \sin\varphi,$$

$$I_c = \frac{2eE_J}{\hbar}.$$

Under a steady subcritical current the phase can remain constant while $V = 0$. At constant nonzero voltage the phase winds and the ideal coherent current oscillates with frequency $f_J = 2e|V|/h$. These two statements distinguish a DC pair current with no terminal voltage from an AC pair response under bias. A real junction includes quasiparticle conduction and capacitance, so its observed total current can contain additional components.

Convert a microwave frequency into a voltage

A junction biased at 5.00 μV has Josephson oscillation frequency $f_J = (2e/h)V = 2.42$ GHz. Illuminate the junction with a stable microwave reference of frequency 10.0 GHz. If its mean phase advance locks to one full winding per reference cycle, then $\langle\dot{\varphi}\rangle = 2\pi f_{\mathrm{ref}}$, requiring mean voltage $hf_{\mathrm{ref}}/(2e) = 20.7$ μV. Locking to two windings per cycle gives 41.4 μV. Such voltage steps probe phase dynamics and the charge associated with coherent transfer.

Suppose the same device has measured critical current $I_c = 30.0$ nA at the operating temperature. Its coupling energy is $\hbar I_c/(2e) = 9.87 \times 10^{-24}$ J, corresponding to $E_J/k_B = 0.715$ K. At a bath temperature of 50 mK thermal energy is appreciably smaller, although phase fluctuations can still depend on capacitance and the surrounding circuit. If either electrode is driven normal, the equilibrium phase-link energy disappears and the ideal supercurrent relation ceases to apply. Observing quantized phase locking and an excitation gap therefore accesses two different signatures of the superconducting state; their characteristic frequencies need not be numerically identical.

The two Josephson relations yield a further calibration test that does not require measuring energy directly. At fixed microwave frequency, the voltage separation of successive locking steps is $hf_{\mathrm{ref}}/(2e)$ independently of junction area in the ideal phase model. Increasing junction area can change total critical current and the visibility or width of current intervals over which locking persists, yet should leave the quantized voltage spacing fixed. Looking at spacing and current width separately distinguishes uncertainty in the reference frequency from a change in the junction coupling or its electromagnetic environment.

Estimate a spectral gap when peaks have finite width

Add a phenomenological lifetime broadening to the isotropic quasiparticle density of states by evaluating $z/\sqrt{z^2-1}$ at $z = (E + i\Gamma)/\Delta$ and taking the physical nonnegative real density. Hold the gap itself fixed in all three runs.

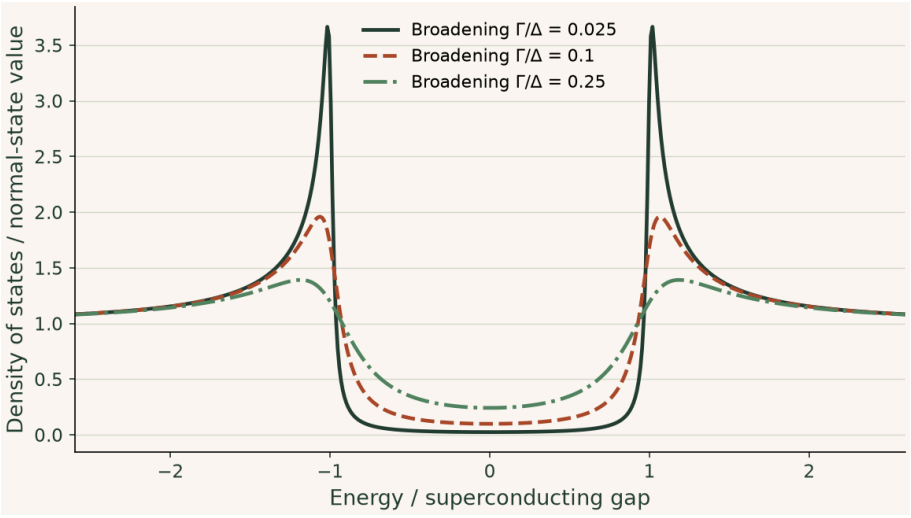

A broader peak and finite central density are possible even with the same gap parameter. At $E = 0$ the ratio is $\gamma/\sqrt{(1+\gamma^2)}$, where $\gamma = \Gamma/\Delta$. Inferring the gap from a threshold in conductance also requires the thermal and instrumental resolution.

Try it · Dynes

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Lifetime broadening Γ/Δ **0.1**

Gap Δ (meV) **1**

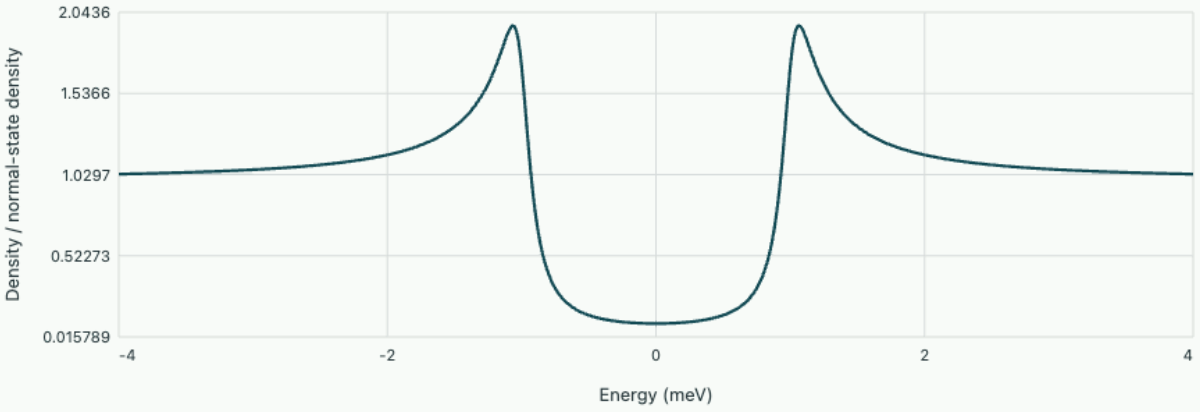

Broadened superconducting spectrum

Zero-energy density ratio 0.099504. This density-of-states model precedes thermal and instrument convolution.

▶ **Read values from the plot**

Why a polar crystal has two optical phonon frequencies

Condensed matter and band theory

A uniform relative displacement of positive and negative ions creates an electric polarisation. A longitudinal vibration induces a macroscopic restoring field while a transverse displacement has a different electrical boundary condition. An infrared spectrum can reveal both frequencies.

Separate the ionic and electronic contributions

For one infrared-active optical mode let u be relative displacement per primitive cell, M the effective reduced mass, Z^* the effective charge in coulombs, and N the number of cells per volume. The oscillator equation driven by the macroscopic field is

$$M(\ddot{u} + \gamma\dot{u} + \omega_{\mathrm{TO}}^2 u) = Z^* E.$$

The zero-macroscopic-field stiffness defines the transverse optical frequency in this oscillator model. Bound electronic response at these frequencies supplies a real relative background permittivity ϵ_∞. The ionic polarisation is $P_{\mathrm{ion}} = NZ^*u$, so $D = \epsilon_0\epsilon_\infty E + P_{\mathrm{ion}}$.

With time convention $e^{-i\omega t}$, solving for u gives

$$\epsilon_r(\omega) = \epsilon_\infty + \frac{S}{\omega_{\mathrm{TO}}^2 - \omega^2 - i\gamma\omega},$$

$$S = \frac{N(Z^*)^2}{\epsilon_0 M}.$$

The signed effective charge appears twice, once as drive and once in the dipole density. Its square therefore sets oscillator strength. At zero frequency the ions can follow the field and $\epsilon_s = \epsilon_\infty + S/\omega_{\mathrm{TO}}^2$. Far above the phonon resonance but below electronic absorption the ionic term is small, making ϵ_∞ the observed plateau.

Apply Gauss's law to a longitudinal vibration

Consider a long-wavelength, nonzero wavevector $\boldsymbol{k}$ with no external free charge. Gauss's law gives

$$\boldsymbol{k} \cdot \boldsymbol{D} = 0.$$

For a longitudinal field $D_L = 0$, so the ions create $E_L = -NZ^*u/(\epsilon_0\epsilon_\infty)$. Substituting this opposing field into the undamped oscillator equation adds a restoring term. Thus

$$\omega_{\mathrm{LO}}^2 = \omega_{\mathrm{TO}}^2 + S/\epsilon_\infty,$$

$$\frac{\omega_{\mathrm{LO}}^2}{\omega_{\mathrm{TO}}^2} = \frac{\epsilon_s}{\epsilon_\infty}.$$

Equivalently the longitudinal collective mode is the zero of the lossless dielectric function. The transverse optical frequency is its pole. For a single isotropic mode this is the Lyddane–Sachs–Teller relation. With several infrared modes, anisotropy, damping or mobile free carriers, assigning one measured peak and one dielectric ratio to a single oscillator needs further justification.

Connect a reflectance band with the two modes

In the lossless limit rewrite $\epsilon_r = \epsilon_\infty(\omega_{\mathrm{LO}}^2 - \omega^2)/(\omega_{\mathrm{TO}}^2 - \omega^2)$. Between the two frequencies this is negative. A transverse electromagnetic wave satisfies $k^2 = \epsilon_r\omega^2/c^2$, so its wavenumber becomes imaginary there in a nonmagnetic lossless crystal. At a planar vacuum boundary a semi-infinite crystal therefore has unit ideal reflectance over this interval. Finite damping makes the refractive index complex and allows absorption, smoothing the edges. The two mode frequencies have been deduced from mechanical response and electric boundary conditions; an electromagnetic wave interacting with the lattice is a coupled phonon–polariton excitation.

For $f_{\mathrm{TO}} = 6.0\,\mathrm{THz}$, $\epsilon_\infty = 4$ and $\epsilon_s = 9$, $f_{\mathrm{LO}} = 9.0\,\mathrm{THz}$. At 7.5 THz in the lossless limit $\epsilon_r = 4(81 - 56.25)/(36 - 56.25) = -4.89$. Below the mode, normal-incidence static-index reflectance is

$$[(\sqrt{9} - 1)/(\sqrt{9} + 1)]^2 = 0.25.$$

Above the phonon plateau the corresponding background reflectance approaches $1/9$. The interactive plot evaluates the passive complex refractive index and the exact single-interface reflectance. Change the static permittivity and damping; compare the high-reflection interval with the lossless zero and pole.

Open this experiment online

Try the calculation

Static relative permittivity	9
Damping ratio γ/ωTO	0.05
Probe frequency (THz)	7.5

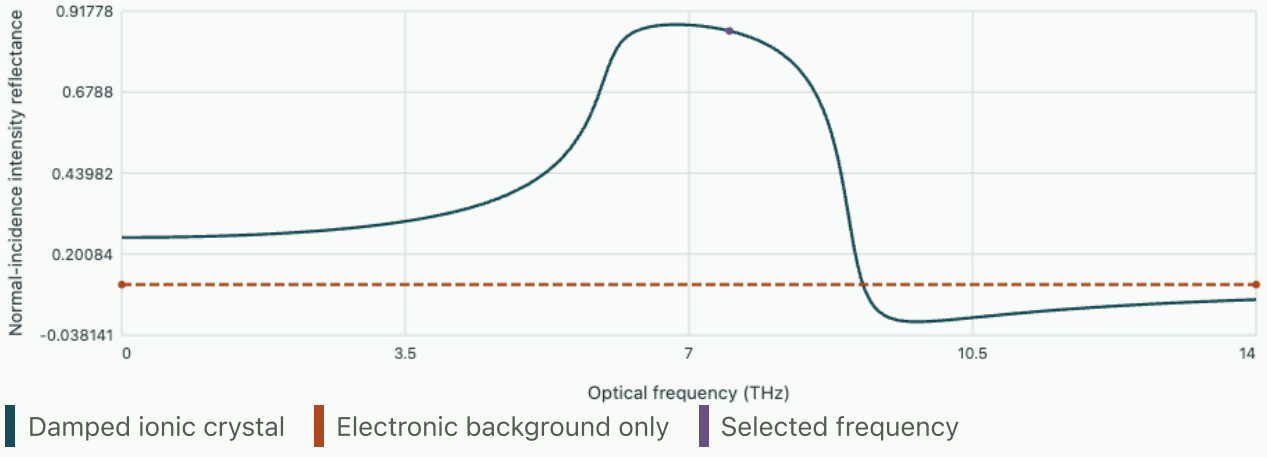

Lossless TO pole 6 THz and LO zero 9 THz. At 7.5 THz the chosen damping gives reflectance 0.8596.

Tutorial on recovering mode parameters

Optical data give $f_{\mathrm{TO}} = 8.0$ THz and $f_{\mathrm{LO}} = 10.0$ THz for a single weakly damped mode. With $\epsilon_\infty = 5.0$, estimate the static permittivity and dielectric increment.

▼ **Work through the solution**

$$\epsilon_s = 5(10/8)^2 = 7.8125.$$

The increment attributed to the optical mode is

$$7.8125 - 5 = 2.8125.$$

Angular and cyclic frequency ratios are equal, although an absolute oscillator strength S requires angular frequencies in the equation of motion.

At a fixed wavevector a probe sees displacement perpendicular to the wavevector. Why is the longitudinal restoring term absent in the electrostatic argument, and what additional coupling can still occur?

▼ **Work through the solution**

Its ionic polarisation has $\boldsymbol{k} \cdot \boldsymbol{P}_{\mathrm{ion}} = 0$, hence it creates no longitudinal bulk bound charge $-i\boldsymbol{k} \cdot \boldsymbol{P}_{\mathrm{ion}}$. There is no obligatory depolarising longitudinal field. A transverse propagating electromagnetic field can still drive the infrared-active vibration, as encoded in $\epsilon_r(\omega)$.

Further reading · MIT solid state physics, optical phonons and dielectric response.

Thermal radiance, étendue and detector calibration

Electromechanics, optics and measurement

A square millimetre of a heated object emits thermal radiation into an entire hemisphere. A detector sees only the rays accepted by its optics and only wavelengths to which it responds. Derive the equilibrium spectrum, follow it through an optical collection system and infer what temperature or emissivity a calibrated reading can actually determine.

Count optical states in a large equilibrium cavity

Take a cubic airless cavity of volume V in thermal equilibrium at temperature T. For its bulk mode density we can use periodic boundary conditions. Allowed wavevectors are separated by $2\pi/L$ in each coordinate. Count a spherical shell in wavevector space, divide by the volume per state and include two independent transverse polarisations. Substituting $k = 2\pi\nu/c$ gives

$$g(\nu)d\nu/V = 8\pi\nu^2 d\nu/c^3.$$

Using standing waves in an ideally reflecting large box gives the same leading volume density when directions and polarisations are counted consistently.

One oscillator mode of frequency ν can hold $m = 0, 1, 2, \ldots$ photons. Its thermal occupation has probabilities proportional to $e^{-mh\nu/k_BT}$. Summing the geometric series and taking its logarithmic derivative gives

$$\langle m\rangle = 1/(e^{h\nu/k_BT} - 1).$$

Multiplying by the density of modes and energy per photon gives spectral energy density u_ν. The zero-point contribution of a harmonic oscillator is not part of the thermal power exchanged between these black surfaces. At low frequency the result approaches an energy k_BT per mode; extending that classical limit to arbitrarily high frequency would make the integrated power diverge.

An isotropic distribution of energy travelling at speed c produces a normal flux through a small aperture weighted by the direction cosine. Per unit solid angle its radiance is $B_\nu = cu_\nu/(4\pi)$. Transform to equal wavelength intervals using $B_\lambda d\lambda = B_\nu|d\nu|$. The two useful spectral forms are

$$B_\nu = \frac{2h\nu^3}{c^2}\frac{1}{e^{h\nu/k_BT} - 1},$$

$$B_\lambda = \frac{2hc^2}{\lambda^5}\frac{1}{e^{hc/(\lambda k_BT)} - 1}.$$

The wavelength expression has units of watts per square metre per steradian per metre of wavelength. To express it per micrometre multiply by one millionth. At 400 K and 10 micrometres the result is 33.56 W per square metre per steradian per micrometre. At 295 K it is 9.14 in the same units. Equal wavelength and equal frequency bins carry different widths, so their distributions reach their maxima at different converted wavelengths. A spectrum's peak should always name its horizontal variable and spectral-density convention.

Distinguish emitter power from brightness in a chosen direction

For a black surface with angle-independent radiance into its outward hemisphere, integrate $\cos\theta\, d\Omega$. The result is π, not 2π. The emitted spectral power per unit area is πB_λ. Integrate over frequency using $x = h\nu/(k_B T)$. The remaining integral $\int_0^\infty x^3/(e^x - 1)\, dx = \pi^4/15$ follows by expanding $1/(e^x - 1)$ as a sum of exponentials and integrating term by term. This yields Stefan–Boltzmann exitance $M = \sigma T^4$ and identifies the constant $\sigma = 2\pi^5 k_B^4/(15 h^3 c^2)$.

A black patch of 1.0 square millimetre at 400 K emits about 1.45 mW into its entire outward hemisphere over all wavelengths. The same patch at 295 K emits about 0.429 mW. Replacing one by the other changes all-wavelength hemispherical power by about 1.02 mW. This is not the power a small infrared receiver sees through a narrow filter. The integrated emitter number is valuable for energy conservation, but by itself contains no receiver area, direction, passband or optical transmission.

The surface spectrum of a body in thermal equilibrium has directional spectral emissivity equal to its absorptivity in the corresponding reciprocal optical modes under the usual reciprocal-material conditions. An opaque diffuse grey surface with emissivity ϵ in an isotropically illuminated background can be modelled by $\epsilon B_\lambda(T) + (1 - \epsilon) B_\lambda(T_{\mathrm{bg}})$. It emits less than a black surface but also reflects background. A shiny object viewed in a warm room may consequently display a radiance temperature nearer the reflected scene than its own physical temperature. Specular surfaces require explicit directional knowledge of that scene.

Derive the geometric acceptance without counting the angle twice

An area element dA normal to the collection axis emits into $d\Omega$ at angle θ a spectral power $B_\lambda \cos\theta\, dA\, d\Omega$. For a uniformly filled cone of half-angle α and a uniform patch of area A, integrating gives $A\pi \sin^2\alpha$. This is its geometrical throughput in air when there is no vignetting. Optical étendue more generally uses $n^2 dA \cos\theta\, d\Omega$. A lossless passive optical relay in geometrical optics preserves this phase-space measure; correspondingly L/n^2 is conserved along a ray in transparent media under the usual stationary conditions.

A direct lens example makes the restriction tangible. An image relay enlarging a uniformly bright patch by a factor of two in each transverse direction forms an image four times the area. In a lossless paraxial design filling the original pupil, each image point then occupies one quarter the solid angle. A detector covering the whole enlarged image and accepting that cone can receive the original power. Placing a fixed small sensor at the image plane intercepts less area instead. Higher magnification has not created additional radiance, even if it makes a small detail easier to resolve.

For the same 1.0 square millimetre patch seen within ten degrees, the projected acceptance is 9.47×10^{-8} square metre steradian. A wavelength-independent combined transmission and detector absorption of 0.65 leaves effective acceptance 6.16×10^{-8} square metre steradian. Integrating the Planck curve from 8 to 12 micrometres gives source and background band radiances 133.74 and 35.43 W per square metre per steradian. Their difference yields only 6.05 microwatts in the receiver. Keeping a single midpoint radiance times four micrometres happens to be fairly close here, but there is no fixed universal midpoint correction across temperatures or filters.

Open the live experiment

Try the calculation

Source temperature (K) **400**

Ambient temperature (K) **295**

Collection half-angle (°) **10**

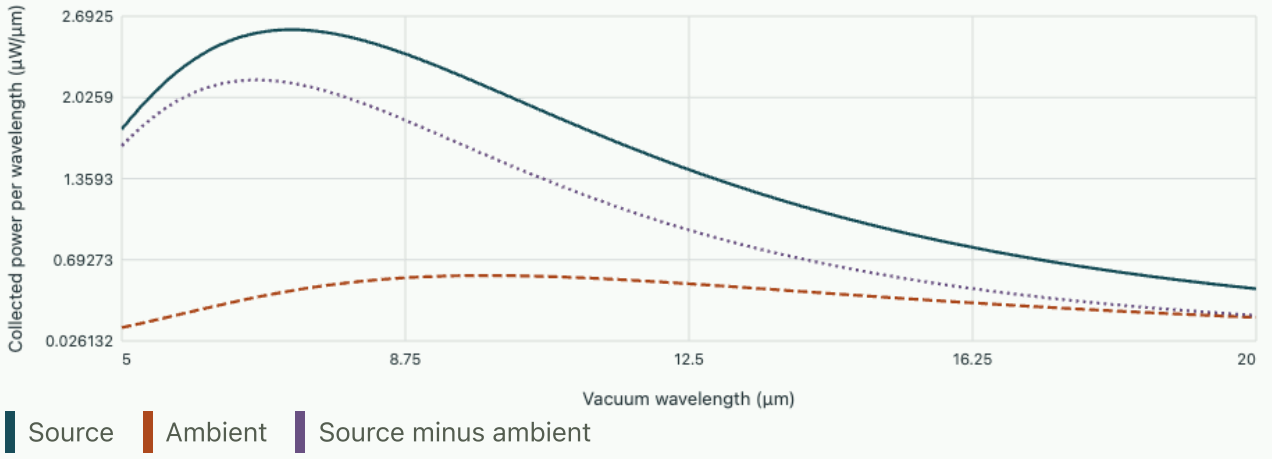

The net 8–12 µm signal is 6.053 µW for a 1 mm² area and 65% throughput. Ideal source peak wavelength is 7.24 µm.

Set source and ambient temperatures equal. The emitted and ambient curves remain finite while the differenced instrument reading vanishes. Increase the cone half-angle from ten to twenty degrees and predict the scale factor $\sin^2 20^\circ / \sin^2 10^\circ = 3.88$. The spectral shape should stay fixed if transmission is separable from angle. If real data shift spectral shape with aperture, inspect filter angle dependence, source nonuniformity and stray light. The integral in the readout is computed on the specified 8–12 micrometre band even though the graph shows wavelengths beyond it.

Convert received power to calibrated voltage or counts

A thermal detector may respond to absorbed power through a temperature rise and an electrical responsivity. A photon-counting detector measures absorbed photons, so its count rate contains a factor $\lambda/(hc)$ multiplying spectral power before integration. Two receivers with identical geometric acceptance and nominal quantum efficiency can respond differently to a changing source temperature if one output is proportional to energy and the other to photon number. Calibrate the actual electrical transfer and make the time interval explicit for accumulated counts.

For an ideal linear power detector let $V = V_0 + gP$, with offset V_0 and gain g. Two calibrated sources spanning the use range determine both. A black reference port at the same ambient temperature can also reject a common background in a chopped measurement. A filter, window or warm lens that absorbs thermal radiation also emits it. If these optics remain unchanged and stable, part of their own emission cancels in a reference difference. A reference traversing a different path generally leaves a residual input that is not captured by an electronic dark measurement.

Suppose effective source-minus-ambient powers of 2.00 and 8.00 microwatts yield 0.120 and 0.420 V. The fitted slope is 0.0500 V per microwatt with an offset of 0.020 V. A reading of 0.320 V implies 6.00 microwatts

after offset correction. Substituting the same reading directly into a single-point gain with zero assumed offset would bias the inferred temperature. Rechecking a third reference near six microwatts tests linearity, while repeated zero-input readings at realistic instrument temperature test offset drift.

Invert the radiance and propagate uncertainties

A narrowband blackbody radiance measurement can be inverted analytically by solving Planck's expression for T. In a broad passband use numerical quadrature across the calibrated spectral response and solve a scalar monotone equation when emissivity and background are known. The derivative at a single wavelength is particularly informative for estimating precision. With $x = hc/(\lambda k_B T)$,

$$\frac{\partial \log B_\lambda}{\partial \log T} = \frac{x}{1 - e^{-x}},$$

$$T = \frac{hc/(\lambda k_B)}{\log[1 + 2hc^2/(\lambda^5 B_\lambda)]}.$$

At 10 micrometres and 400 K the logarithmic sensitivity is 3.70. A one-percent relative radiance uncertainty therefore gives about 1.08 K temperature uncertainty if all other quantities are exact. An unknown five-percent emissivity error can dominate that estimate. A two-colour temperature ratio may reduce uncertainty from an emissivity that is constant across the two narrow bands, but cannot cancel an arbitrary spectral emissivity or a large reflected ambient component.

As an example consider a grey object at 400 K and ambient 295 K with emissivity 0.70. At ten micrometres its outgoing radiance is $0.70(33.56) + 0.30(9.14) = 26.24$ in the earlier spectral units. Inverting as a perfect blackbody gives a brightness temperature of about 375 K. The difference is a physical radiance–temperature degeneracy. Repeat at known ambient temperatures while keeping the object thermally stabilised to constrain the reflected fraction. If heating the background changes the actual object temperature, a second independent temperature sensor or a transient thermal model is necessary.

When detector heat capacity sets the measured signal

Consider a thermal receiver of heat capacity C_d coupled to its surroundings with thermal conductance G_d. For small excursions its temperature obeys $C_d\dot{\theta} + G_d\theta = P_{\mathrm{abs}}(t)$, where θ is the excess above its operating point. A voltage thermometer with responsivity r_T converts this into output. Its frequency-dependent power responsivity is $r_T/(G_d - i\omega C_d)$. The geometric and spectral calculations determine absorbed optical power, while this thermal equation determines how much of its modulation appears at the electrical output.

Take $C_d = 2.0$ microjoules per kelvin and $G_d = 0.10$ milliwatt per kelvin. The thermal time constant is 20 ms and ordinary corner frequency is 7.96 Hz. A sinusoidal absorbed-power amplitude of 6.0 microwatts would give a static temperature rise of 0.060 K. At 20 Hz the amplitude is reduced by $1/\sqrt{1 + (2\pi f\tau)^2} = 0.370$, leaving about 22.2 mK with a phase lag magnitude near 68 degrees. A receiver calibrated with a steady optical power is not automatically calibrated for this chopped experiment.

A square-wave chopper also changes the relation between time-domain peak differences and a lock-in's fundamental reading. For equal hot and reference dwell times, a square wave switching between powers P_1, P_0 has fundamental sinusoidal peak amplitude $2(P_1 - P_0)/\pi$, with different factors for rms and peak reports. The thermal transfer attenuates this component at the chopper frequency and the odd harmonics at

their respective frequencies. Substituting the full hot-minus-reference power directly into the first-harmonic response would overestimate the expected signal.

Measure a frequency sweep on a stable reference and test whether amplitude and phase share a single time constant. A geometrical error rescales the amplitude but ordinarily leaves this thermal phase unchanged. A second thermal mass or an electronic filter can introduce another pole. Lowering chop frequency increases thermal signal until the static limit but can expose low-frequency environmental drift. Choose the frequency by comparing the calibrated response with the measured noise spectrum.

Finally, a reference and an unknown source must illuminate the same active area if responsivity varies across the receiver. Imaging one source as a small spot and the other as a filled patch can create a calibration bias even when total spectral powers agree. Translating a reference spot across the detector or imaging the same aperture successively at both temperatures tests that spatial uniformity assumption.

If the sensor is used on a source much smaller than the calibrated reference aperture, record an image or a field-stop scan to determine its filling fraction. Otherwise even exact knowledge of filter transmission and thermodynamic temperature leaves an unknown geometrical scale factor. Aperture filling is part of the measurement equation, not an electronic gain setting.

Tutorial · Infer emissivity using two controlled backgrounds

An opaque diffuse sample at fixed temperature is observed at 10 micrometres against two isotropic backgrounds with calibrated spectral radiances 8.00 and 16.00 W per square metre per steradian per micrometre. The resulting outgoing readings are 25.00 and 27.00 in the same units. Infer emissivity and the sample's blackbody-equivalent radiance. Estimate the significance of independent reading errors of 0.10 units.

▼ **Separate the reflected and emitted contributions**

The increased outgoing radiance is 2.00 for an 8.00 increase of background. Thus reflectivity is 0.250 and emissivity 0.750. The first measurement gives intrinsic blackbody radiance $[25.00 - 0.25(8.00)]/0.75 = 30.67$. The second returns the same value. With 0.10 independent uncertainty on each outgoing reading and exact backgrounds, the slope uncertainty is $\sqrt{2}(0.10)/8 = 0.0177$, as is the absolute emissivity uncertainty. Larger background separation can improve this slope estimate if sample temperature stays fixed.

If the sample transmits radiation, the residual fraction does not all correspond to reflection. A controlled change behind the sample helps detect that additional path. A directional specular reflector exposed to a warm enclosure also does not necessarily see the same angular average used in the diffuse model. Without these checks, a formally precise fitted emissivity can describe the experimental surroundings instead of the material.

Tutorial · Design a numerical spectral-integration check

An engineer tabulates spectral radiance per micrometre on a wavelength mesh. Another rewrites the detector model on an equally spaced frequency mesh, obtaining a signal that differs by several percent after

independent mesh refinement. Explain what must transform, what stays dimensionless and how to compare a photon-counting channel with a power channel.

▼ **Transform the complete measurement integral**

The differential power $B_\lambda \, d\lambda$ must equal $B_\nu \, |d\nu|$, with $|d\lambda/d\nu| = c/\nu^2$. A dimensionless transmission curve becomes the same physical transmission at $\lambda = c/\nu$; it acquires no additional Jacobian. The integration bounds reverse when increasing wavelength is replaced by increasing frequency. A response specified per unit wavelength, or a calibrated gain integrating a particular finite bin, must instead be interpreted from how it was defined.

Evaluate a unit-transmission test using a very narrow band, then the full 8–12 micrometre band whose radiance values are given above. Refine each grid independently, treating sharp filter edges consistently and converting per-micrometre densities before using metres. For a photon counter divide each spectral power contribution by photon energy before summation. Replacing that energy by a single central-wavelength value is an approximation that should be checked against the full band and source temperature. Close agreement between two different meshes using the same incorrect density is not a valid units check.

NIST radiance-temperature calibrations · spectral radiance, sources and uncertainty

Counting, absolute entropy and the entropy of mixing

Statistical mechanics from microscopic counting to material response

An ideal gas expands into an evacuated bulb and its temperature can remain essentially unchanged. What changed if a thermometer reports the same value? Begin by specifying a microscopic count, recover the pressure and temperature of the gas, and test the meaning of distinguishability by removing a partition between two samples. The calculation also explains why the classical formula for absolute entropy needs a microscopic action scale.

Describe a count before taking its logarithm

A macrostate specifies only outcomes resolved by a chosen preparation and measurement, for example particle number, energy within a narrow interval and container volume. A microstate specifies detailed positions and momenta consistent with those outcomes, or a quantum state in the corresponding energy shell. The multiplicity is the number of allowed states under the stated resolution and constraints. The logarithm makes multiplicities of statistically independent systems additive, because the number of joint possibilities is their product. Boltzmann entropy is

$$S = k_B \log \Omega.$$

Before computing a gas, inspect an easier count. If eight distinguishable sites each have a variable taking values zero or one and exactly three sites are occupied, there are $\binom{8}{3} = 56$ configurations. Revealing which sites are occupied reduces the compatible set to one. This change of conditional information need not involve an actual heat flow in the original system. Thermodynamic entropy requires a physical specification of allowed states and operations; it cannot be inferred solely from how surprised a particular observer feels.

In equilibrium, equal a priori weighting in a narrow isolated energy shell is a postulate about the appropriate ensemble. It does not state that a finite trajectory visits every conceivable state in any chosen observation time. Long-lived constraints, such as a closed partition or a conserved component of angular momentum, change the accessible shell. A gas that has not exchanged particles between remote compartments can be assigned a separate particle number in each. Removing a constraint can enlarge the equilibrium count even while the volume of an exact fine-grained Hamiltonian distribution is preserved by microscopic dynamics. The thermodynamic description refers to a macroscopic resolved equilibrium state.

Count the spatial and momentum volumes of a monatomic gas

Take N indistinguishable noninteracting particles of mass m in volume V, with total kinetic energy at most E. Integrating over positions produces V^N. Allowed momenta lie inside a sphere in $3N$ dimensions of radius $\sqrt{2mE}$. Its volume follows from computing a multidimensional Gaussian integral both in Cartesian and radial coordinates. Divide by a microscopic cell h^{3N} to make the state count dimensionless and by $N!$ to prevent a permutation of indistinguishable particles from creating another physical state. The cumulative count is

$$\Phi(E, V, N) = \frac{V^N}{N! h^{3N}} \frac{(2\pi m E)^{3N/2}}{\Gamma(3N/2 + 1)}.$$

Using the shell count $\partial_E \Phi \, \delta E$ instead changes subextensive terms for a conventional narrow shell that still contains many states. The cumulative formula is especially convenient for this gas. Equalities at finite small N

depend on which entropy is chosen, so state that choice before claiming an exact finite-size temperature. The usual bulk relations follow from either in the thermodynamic limit for this unbounded, regular system.

For large N, use $\log N! \simeq N \log N - N$ and the analogous gamma-function asymptotics. After collecting the terms,

$$S \simeq N k_B \left[\log\left(\frac{V}{N} \left(\frac{4\pi m E}{3 N h^2} \right)^{3/2} \right) + \frac{5}{2} \right].$$

Doubling E, V, N at fixed ratios doubles the leading expression. Omitting $N!$ would leave a term growing like $N \log N$, violating this extensivity. The phase-space cell influences the entropy per particle and the chemical potential but drops out of classical fixed-particle derivatives for pressure and temperature. Quantum mechanics fixes the natural cell scale through state counting. It does not require physical particles to carry serial numbers while remaining otherwise identical.

Read temperature, pressure and chemical potential from one function

Write the entropy differential as $dS = dE/T + (p/T)dV - (\mu/T)dN$. Differentiating the count at fixed remaining variables gives

$$1/T = 3 N k_B/(2E),$$
$$p/T = N k_B/V.$$

Thus an energy exchange measurement and a force on a moving wall follow from derivatives of the same count. The thermal wavelength $\lambda_T = h/\sqrt{2\pi m k_B T}$ lets us rewrite the entropy and chemical potential as

$$\frac{S}{N k_B} = \frac{5}{2} - \log(n \lambda_T^3),$$
$$\mu = k_B T \log(n \lambda_T^3),$$
$$n = N/V.$$

The chemical potential uses the same single-particle energy zero as the energy expression. Shifting every particle energy by a constant shifts μ by that constant, leaving equilibrium criteria involving consistent references unchanged. An unresolved internal ground-state degeneracy g would multiply the state count by g^N in the classical dilute limit and replace $n\lambda_T^3$ by $n\lambda_T^3/g$. Real molecules and atoms can have internal states and additional degrees of freedom; never assign monatomic entropy to a molecular gas without justifying its internal contribution.

For ideal monatomic argon at 300 K and 1.00 bar, using $m = 39.948$ unified atomic mass units, $\lambda_T = 15.95$ pm and $n\lambda_T^3 = 9.79 \times 10^{-8}$. The result is approximately 154.97 J per mole kelvin. Its small degeneracy parameter supports the dilute classical approximation. Compressing isothermally by a factor of ten increases chemical potential by $k_B T \log 10$ and decreases entropy by $N k_B \log 10$. Extending the classical formula down to zero temperature at fixed density would eventually give a negative unbounded entropy, signalling the breakdown of the classical state occupation assumptions.

Why a partition between identical samples is different

Take two ideal gases initially separated at equal temperature and pressure. Their volumes are proportional to their particle numbers. For two distinguishable species each component expands to the full volume after the partition is removed. The internal kinetic energies are unchanged in the ideal model. Summing the two entropy increases and writing mole fractions $x_i = N_i/N$ gives

$$\Delta S_{\mathrm{mix}} = -N k_B \sum_i x_i \log x_i.$$

One mole containing fractions 0.30 and 0.70 gains 5.08 J/K upon ideal mixing from the pure components. At 300 K a reversible isothermal separation needs minimum work $T\Delta S_{\mathrm{mix}} = 1.52$ kJ for this mole if the ideal components and reference states are restored. The free expansion itself need not deliver that work. It is available in a reversible mixing path using selective barriers or equivalent controlled separation operations.

If both compartments initially contain identical gas at equal density, the final macrostate has the same density and same single species throughout. Combining two extensive entropies gives the final entropy to leading bulk order. There is no extensive mixing increment $N k_B \log 2$. Introducing artificial labels for the former left and right particles changes the physical question only if a real degree of freedom can be prepared, distinguished and selectively processed. A dye, isotope or internal spin state can serve as such a label under suitable experimental operations. Merely drawing colours on a computer representation of identical classical particles does not establish a membrane that could sort those labels.

A tempting argument says that making two different gases more and more similar must continuously reduce their ideal mixing entropy. In the ideal model the species occupation labels are either distinct or identical; the description does not specify a continuously degrading separation instrument. A model of partially distinguishable internal states or a noisy finite-resolution sorter is a different experiment. Its accessible work depends on the physical state preparation and measurement operations. Distinguish the mathematical idealisation from an actual membrane having imperfect chemical selectivity.

Check an entropy change using a reversible reference route

In an isolated free expansion of a dilute gas into a vacuum the external pressure vanishes and boundary work is zero. With negligible heat flow energy remains constant and for an ideal gas so does temperature. To compute the state-function entropy change choose a hypothetical reversible isothermal expansion between the same endpoints. Reversible heat input is $\int p\,dV = N k_B T \log(V_f/V_i)$, giving $N k_B \log(V_f/V_i)$. This reference heat is not a claim that heat actually crossed the boundary during free expansion.

For a more exact real-gas experiment, first check that the initial and final temperatures are indeed the same. Internal energy may depend on volume. From the fundamental differential and a Maxwell relation one obtains

$$(\partial U/\partial V)_T = T(\partial p/\partial T)_V - p.$$

A temperature change during an isolated free expansion thus carries information about intermolecular interactions. Measuring only final pressure and assuming constant temperature could misattribute an interaction-induced energy change to the ideal state count.

Use the existing energy-sharing experiment to see how multiplicity changes with a constraint

Recover the entropy cost of a number fluctuation

A gas can lose an extensive mixing entropy when a physically selective partition is imposed, yet the finite-size cost of constraining only a particle count is much smaller. To see this, divide a box into two equal geometric halves with an imaginary plane and let all particles travel freely. At fixed total particle number each classical ideal particle has probability one half to be on either side. The probability of exactly k particles on the left is $\binom{N}{k}/2^N$, independently of whether the particles have names in a bookkeeping algorithm. The binomial coefficient counts the spatial arrangements compatible with that number. The resulting variance of the left count is $N/4$.

The chance of exactly the mean count for even N is approximately $\sqrt{2/(\pi N)}$, from Stirling's formula. If a protocol constrains the left count to precisely $N/2$, the entropy difference between allowing all counts and that constrained ensemble is therefore about $(k_B/2)\log(\pi N/2)$. It grows only logarithmically with total size. In contrast, if the constraint identifies which half contains every member of two distinguishable species, the mixing entropy is extensive. The physical variables that the partition controls decide which count is appropriate.

For $N = 100$ the probability of exactly fifty on the left is about 0.0796. The associated entropy cost of conditioning is $2.53 k_B$. A macroscopic observer resolving the fraction only to a few percent can count a whole interval of nearby values as the same macrostate. As N grows the relative width shrinks like $1/\sqrt{N}$, placing almost all of the probability inside any fixed nonzero fractional interval about one half. This is how an apparently precise macroscopic density can coexist with an enormous number of fluctuating microscopic arrangements.

The far tails reveal a stronger distinction. For a fixed fraction $x = k/N$ away from one half, applying Stirling to the binomial gives probability proportional to $\exp[-N I(x)]/\sqrt{2\pi N x(1-x)}$, with

$$I(x) = x\log(2x) + (1-x)\log[2(1-x)].$$

This rate function is nonnegative, vanishes at one half and starts quadratically there. At $x = 0.7$ it is about 0.0823, so the leading exponential cost for a thousand particles is $e^{-82.3}$. A Gaussian near the peak gives a similar rough scale here but fails to capture the exact tail or the hard limits $0 \le k \le N$. A useful numerical exercise calculates binomial coefficients using log factorials, compares the central Gaussian approximation with the full rate function and checks that the probabilities sum to one. Evaluating the three factorials directly can overflow while the final probability is perfectly representable.

A reversible separation engine and a deliberately imperfect measurement

Imagine two ideal distinguishable gases separated initially at equal pressure and temperature. To obtain work while mixing, use species-selective membranes and allow reversible isothermal expansion of each partial gas against an appropriate pressure. The membrane traversed by one component must not transmit that component to the wrong side, and every infinitesimal motion must occur at negligible pressure imbalance for the component doing the work. The sum of the recoverable works is the sum of each component's ideal expansion work. Reversing the device separates the mixture at the same thermodynamic cost when dissipation can be made negligible.

For a one-mole mixture containing 0.20 mole of one species and 0.80 of another at 300 K, the work contribution of the minority component is $-0.20RT\log 0.20 = 0.803$ kJ and that of the majority is $-0.80RT\log 0.80 = 0.445$ kJ. Although the minority has fewer particles, each expands across a larger volume ratio and its work contribution is larger. The total is 1.25 kJ. A bulk piston doing no selective separation sees equal total initial pressures and no obvious macroscopic pressure difference. The selective mechanical operation matters to the interpretation of free energy.

A useful idealised sorter determines a binary label with symmetric error probability e, with equal prior fractions. The fraction of correctly sorted particles is $1 - e$, yet that single number by itself says little about its thermodynamic performance unless both output compositions are measured. For the specified symmetric device the posterior composition in each equal outlet is $(1 - e, e)$ in opposite order. The increase in mixing free energy relative to one unseparated equimolar reservoir at the same total temperature, density and amounts is

$$\Delta F = N k_B T[\log 2 + (1 - e)\log(1 - e) + e \log e].$$

This is a state-function comparison of outlet compositions, not an assertion that the actual apparatus consumes this amount or that measuring a label is free. At $e = 0.10$ the bracket is 0.368, about 53 percent of the complete separation value. At $e = 0.5$ it vanishes even though a nominal label was printed for every particle. If one outlet loses particles or the two error probabilities differ, compute actual outlet fractions and chemical potentials from a material balance instead of using this symmetric formula.

An experiment can close that balance without seeing individual molecular trajectories. Measure both outlet amounts, composition, pressure and temperature, and account separately for work or heat supplied to the sorter. Compare the Gibbs free-energy change at controlled pressure with the actual electrical work. A device consuming several times the reversible minimum can still separate effectively. Conversely, apparent work below the calculated minimum suggests an omitted material stream, a different final pressure, an energy supplied by an uncounted reservoir or uncertainty in composition. The mixing formula is only as meaningful as the thermodynamic endpoints used.

Follow isotope and quantum effects to the edge of the classical formula

The appearance of h in a classical-looking entropy sometimes suggests that the Sackur–Tetrode expression already describes any temperature so long as gas interactions are weak. These are two different approximations. Diluteness relative to the interaction range controls the importance of forces. Diluteness in phase space, measured by $n\lambda_T^3$, controls whether quantum exchange changes occupation weights. Atoms can interact weakly while exhibiting strong exchange statistics. Lowering temperature makes the thermal wavelength grow as inverse square root of temperature, quickly invalidating the classical count at a fixed density.

Consider two isotopes with one-particle masses in the ratio four to three, neglecting internal degeneracy differences and comparing at the same density and temperature. Their wavelengths have ratio $\sqrt{3/4}$. The heavier isotope has a translational molar entropy greater by $(3R/2)\log(4/3) = 3.59$ J per mole kelvin. This does not contradict the mass independence of the ideal mixing increment. The former compares two absolute entropy values of different substances; the latter is a change in each substance's accessible translational volume while its mass remains fixed.

The same mass scaling influences onset of quantum degeneracy. At fixed n, the temperature for a chosen value of $n\lambda_T^3$ varies inversely with mass. A heavier isotope reaches that value at a lower temperature. If the nuclei also have different spin degeneracies, distinguish the dimensionless density per internal component from the total density. Two isotopes can even have different Bose or Fermi exchange statistics, so a high-temperature entropy comparison need not extrapolate into similar low-temperature heat capacities.

A calorimetric consistency test illustrates the limits. Suppose an idealised solid has a measured low-temperature molar $C_p = AT^3$ up to 5 K, with $A = 0.20$ mJ per mole per kelvin to the fourth. Its entropy increment there is $AT^3/3 = 0.00833$ J per mole kelvin. Missing that contribution hardly changes a room-temperature entropy tabulated to a few tenths of a joule per mole kelvin. In contrast, omitting a latent heat of 1.0 kJ/mol at a transition near 100 K misses 10 J per mole kelvin. The low-temperature extrapolation and the phase transition enter the same entropy integral but can carry very different experimental leverage.

A finite thermometer measures a ratio of multiplicities

Temperature from an entropy derivative presupposes sufficiently smooth variation with energy. In a small discrete environment a two-level probe of gap Δ can instead measure an exact finite difference. If the whole isolated system has total energy E and the two probe states have equal degeneracy, their equilibrium odds are $\Omega_R(E-\Delta)/\Omega_R(E)$. Write the logarithm as $[S_R(E-\Delta) - S_R(E)]/k_B$. Interpreting it as $-\Delta/(k_B T_{\mathrm{probe}})$ defines an operational inverse temperature averaged over an interval of width Δ.

For a reservoir whose cumulative multiplicity is proportional to E^α, this gives $\beta_{\mathrm{probe}} = \alpha \log[E/(E-\Delta)]/\Delta$, while the local slope at E gives α/E. With $\Delta/E = 0.10$, the finite-gap value exceeds the endpoint slope by 5.36 percent. A second probe with one fifth that fractional gap reports only about a one-percent excess. Two well-calibrated probes can therefore disagree in a finite reservoir while each correctly reports its specified probability ratio.

The distinction vanishes for a macroscopically large heat bath whose entropy is locally almost linear over both gaps. It is experimentally testable in numerical energy-sharing models without calibrating a mechanical pressure. Count both occupancy ratios, establish whether the probe appreciably changes the reservoir's allowed states and compare against the exact discrete multiplicities before invoking a smooth temperature. If repeated measurements perturb the environment, the measured distribution is a sequence of new preparations and the equilibrium interpretation needs an explicit readout model.

Construct a thermodynamic cycle that isolates mixing entropy

The chemical-potential formula suggests a pressure measurement for testing mixing free energy. Let one side of a movable membrane transmit a solvent species freely and block a second dilute gas species, with common temperature maintained. In the ideal-gas case the permeable species equilibrates its partial pressure across the membrane. If on one side it is pure at pressure p_A, while on the other side it has mole fraction x at total pressure p_{mix}, equality gives

$$p_A = x p_{\mathrm{mix}}.$$

The remaining pressure load on the membrane is the partial pressure of the blocked species. This is a mechanical consequence of equality of the permeable component's chemical potential.

For an ideal mixture with $x = 0.80$ and total pressure 1.00 bar, the pure compartment must reach 0.80 bar for zero net flow of the permeable species. A square centimetre membrane then supports a net pressure difference of 0.20 bar, corresponding to 2.0 N, before any additional elastic support forces. A membrane sensitive only to the total pressure might move and alter the volumes until mechanical and chemical constraints are both fulfilled. The quoted force is a held-geometry measurement, not a prediction that a freely moving unsupported membrane can remain at that position.

A particularly important control switches the membrane for a nonselective porous partition. Both species can then flow and equilibrate their own chemical potentials; at final full equilibrium there is no maintained composition difference across these gas reservoirs. A pressure reading alone cannot establish whether selectivity, partial blockage, a small temperature difference or an external pump generated the observed imbalance. Measure both compositions and isolate slow leaks over the relevant time. An apparently stable jump on a short trace can be a long-lived transient.

The same thought experiment gives a check on how gas-phase mixing connects to osmotic pressure in dilute solutions. The permeable solvent's chemical potential is lowered by mixing at fixed total pressure; applying enough pressure to the solution can restore equality across a solvent-selective membrane. The ideal dilute result has a form proportional to solute number density times thermal energy, although a liquid solution requires its own solvent partial molar volume model. Equal-looking formulas arise from a chemical-potential balance, not because dissolved solute molecules necessarily act as an ideal gas physically bouncing against every membrane.

Finally, a cyclic device must restore membrane positions, reservoir compositions and temperatures as well as a piston reading. Ignoring the work to regenerate selectivity or restore separated gas streams can turn a one-use release of stored mixing free energy into an apparent heat engine operating on one bath. Write the mole count of each component in every reservoir at start and finish and integrate mechanical and chemical work consistently. The criterion for a true thermodynamic cycle is return of the complete controlled system to its initial state.

A final microscopic check concerns the dimensions of the count. If one changed the formal classical cell size from h^3 to $a h^3$ per particle, the gas entropy would shift by $-N k_B \log a$ at every fixed state. Pressure, temperature and the entropy change of a fixed-N expansion would be unaffected. Chemical potential would shift by $k_B T \log a$, and absolute comparisons between different particle species or chemical reactions would require a consistently calibrated microscopic count. This is why a piston experiment alone could never determine the natural action cell. Combining calorimetry with quantum level counting or fixed-reference chemical equilibria supplies information absent from fixed-N mechanical derivatives. The same bookkeeping distinction applies to any unobserved constant internal degeneracy: it can contribute entropy without producing a heat-capacity anomaly.

Connect ideal mixing to a separation-work measurement

Vary the mixture composition and the reservoir temperature. Compare the entropy with reversible work for separately identifiable ideal components.

Open this experiment online

Try the calculation

Mole fraction of component A **0.3**

Separation reservoir temperature (K) **300**

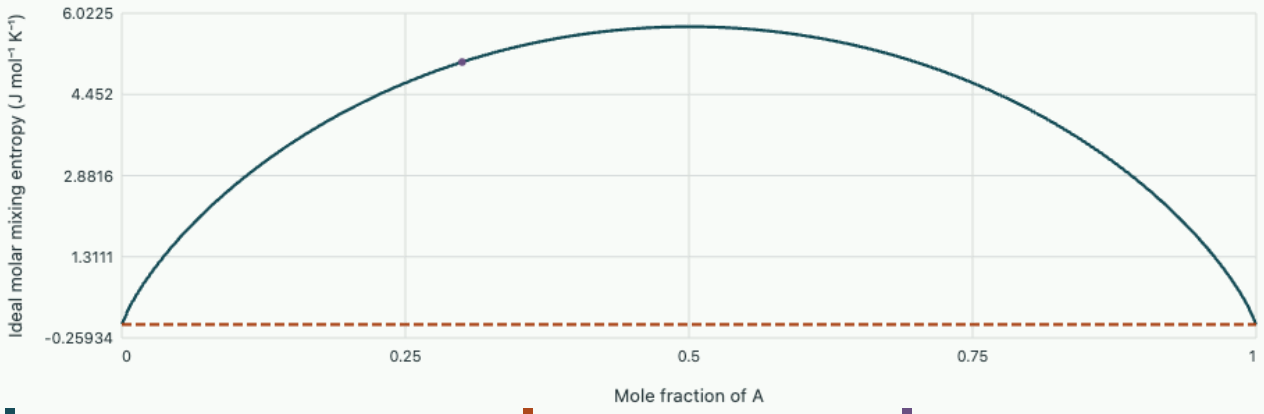

Two distinguishable ideal components | Identical gas on both sides | Selected composition

Ideal entropy increase 5.079 J mol⁻¹ K⁻¹. Reversible isothermal separation of the distinguishable mixture requires 1.5237 kJ/mol.

Tutorial · Measure an entropy difference by two routes

A 0.50 mol sample of ideal monatomic gas expands from 2.0 to 5.0 litres in a thermally insulated free expansion. Calculate entropy change and change of chemical potential at 290 K. Now identify what changes if the gas instead expands reversibly and adiabatically between the same volumes.

▼ **Hold the state variables appropriate to each route**

For free expansion temperature stays at 290 K and entropy increases by $0.50R\log(2.5) = 3.81$ J/K. At fixed temperature $\mu_f - \mu_i = k_B T \log(n_f/n_i) = -k_B T \log 2.5$, or minus 2.21 kJ/mol. The chemical-potential difference is a single-particle or molar quantity, while the entropy answer refers to the entire sample. Dividing or multiplying by the sample amount twice produces an inconsistent unit.

In a reversible adiabatic expansion $\Delta S = 0$ and the gas does boundary work. The final temperature becomes $290(2/5)^{2/3} \simeq 157$ K. These are not the same endpoints as in the free expansion, despite identical start and end volumes. Substituting the free-expansion entropy into an adiabatic isentropic relation would falsely equate insulation with reversibility.

Tutorial · An isotope mixture and a finite-size count

Two equal amounts of isotopically distinct ideal monatomic gases at the same temperature and density are mixed. Explain why their masses enter absolute chemical potentials but cancel from each component's ideal isothermal mixing increment. For just ten sites individually occupied at random by one of two distinguishable

isotope labels, compare the exact number of arrangements with five of each label to the leading extensive entropy.

▼ **Cancel fixed internal terms and keep the finite combinatorics**

Each component's thermal wavelength is unchanged between its own initial and final states. Taking the difference of its entropy therefore removes the mass and internal factors. Both species expand by a factor of two and the large-sample increase is $Nk_B \log 2$. This conclusion assumes independently identifiable species and an ideal mixture, and does not state that their absolute entropies are identical.

For ten sites with five labels of each kind there are $\binom{10}{5} = 252$ arrangements. Their logarithm is 5.529 in units of k_B, versus $10 \log 2 = 6.931$ from an extensive approximation neglecting the composition constraint's subextensive term. Summing over all possible label counts gives 2^{10}, recovering the latter count for an ensemble in which the composition itself fluctuates. A small-system discrepancy is an opportunity to ask which constraint the count imposes.

Tutorial · Infer what a calorimetric entropy determination misses

A crystal's heat capacity is measured from a low reference temperature up to a temperature above a reversible melting transition. How could this measurement be used to estimate entropy, and which missing physical contributions could spoil an apparent comparison with a dilute gas entropy?

▼ **Include the low-temperature reference and all phase changes**

Along a constant-pressure reversible route add $\int C_p(T) dT/T$ in each phase and add latent heat divided by transition temperature at every first-order change. To compare with gas one must also traverse vaporisation or another controlled path to the desired gas state and pressure. Missing a transition silently omits its finite entropy even if the heat capacities in both adjacent phases were measured perfectly.

Below the lowest measured temperature extrapolate using a physically justified low-temperature model and add any residual ground-state entropy consistent with the sample preparation. A unique ordered ground state has zero thermodynamic entropy in the ideal third-law limit. A crystal retaining configurational disorder on laboratory timescales can carry residual entropy associated with its accessible arrangements. Unknown isotopic composition or low-lying internal levels likewise change the absolute comparison. A discrepancy between a calorimetric integral and a monatomic ideal expression should first be traced to these explicit reference states and model assumptions.

Cambridge Statistical Physics · foundational postulates, entropy and statistical ensembles

Fließbach, Statistische Physik · statistical foundations, the ideal gas and entropy

Ensembles, thermodynamic potentials and measurable curvature

Statistical mechanics from microscopic counting to material response

A sample sealed inside an insulated container, a sample weakly coupled to a temperature bath and a surface exchanging particles with a reservoir need different probability laws. Derive the three laws from one composite-system count. Then use the curvature of their generating functions to predict noise, response and parameter uncertainty without assuming that microscopic quantities take their mean values in every run.

Expand the reservoir count rather than postulating a probability

Let a small system occupy a microstate i of energy E_i while a weakly coupled reservoir holds the remainder of a fixed total energy. The composite system is isolated. Conditional on i, the number of compatible reservoir states is $\Omega_R(E_{\text{total}} - E_i)$. This weight is the same for degenerate system microstates of equal energy. For a macroscopic reservoir, expand its entropy about an appropriate reference energy. The first derivative is $1/T$, and its second derivative at fixed other extensive parameters is $-1/(T^2 C_R)$.

$$S_R(E_R - \epsilon) = S_R(E_R) - \frac{\epsilon}{T} - \frac{\epsilon^2}{2T^2 C_R} + \cdots.$$

Exponentiating and taking the limit where the quadratic and further corrections are negligible yields $p_i = e^{-\beta E_i}/Z$, $\beta = 1/(k_B T)$, and $Z = \sum_i e^{-\beta E_i}$. A constant shift in all system energies multiplies both numerator and partition sum by the same factor. Relative probabilities are therefore independent of energy-zero convention. If a calculation sums over energy values rather than microstates, insert their degeneracies explicitly. These two equivalent ways of counting should never be applied simultaneously.

The derivation assumes weak coupling so that assigning energies separately makes sense, and a reservoir whose temperature does not respond significantly to typical exchanges. A small heat capacity in the bath can suppress sample fluctuations even at the same mean temperature. The relevant test concerns excursions explored by the sample, not simply whether the bath contains more particles. Near a transition or in an engineered small reservoir a canonical distribution may need a measurable correction.

Compare sample-energy fluctuations for finite and effectively infinite heat reservoirs

Differentiate the normalisation to obtain energy and its variance

Differentiate $\log Z$ with respect to β at fixed energy levels. The first derivative yields minus mean energy. A second derivative gives the mean square minus the squared mean. Changing differentiation from β to temperature gives the equilibrium heat capacity:

$$\langle E \rangle = -\partial_\beta \log Z,$$

$$\operatorname{var}(E) = \partial_\beta^2 \log Z = k_B T^2 C_V.$$

This proves nonnegative canonical heat capacity for a temperature-independent Hamiltonian. If the model Hamiltonian contains parameters that themselves vary with temperature after other degrees of freedom have been eliminated, the derivative also acts on those parameters. Inserting measured effective levels into the simple identity without that dependence need not reproduce the calorimetric capacity.

A two-level system has energies zero and Δ, with upper probability $p = 1/(1 + e^{\beta\Delta})$ for equal degeneracies. Its variance is $\Delta^2 p(1-p)$ and $C/k_B = (\beta\Delta)^2 p(1-p)$. At $\Delta/k_B T = 2$, $p = 0.1192$ and $C = 0.420 k_B$. In one thousand independent copies the expected upper count is 119 and its standard deviation about 10.2. An actual count of 130 is unremarkable. An experiment resolving only one sample at one time cannot identify the variance from its energy alone.

Higher derivatives of $\log Z$ generate cumulants. A Gaussian approximation discards cumulants beyond the second. For a small two-level population near low occupancy the histogram is skewed and the approximation can yield artificial negative occupancies. An exactly binomial count or a direct discrete partition sum handles this elementary case without assuming that the central-limit regime has arrived. Sample-to-sample correlations can also prevent variances from scaling as the sum of independent variances.

Why the Helmholtz free energy knows the same equilibrium as entropy

Group the canonical sum by energy. A shell with entropy $S(E)$ has weight proportional to $e^{S(E)/k_B - \beta E}$. For a macroscopic system with a dominant smooth interior maximum this exponent is stationary at $(\partial S/\partial E)_{V,N} = 1/T$. The leading free energy $-k_B T \log Z$ is then the minimum of $E - TS(E)$; the width of the peak gives the energy fluctuations. The exact finite-system integral is a Laplace transform. Replacing it by its stationary contribution is an approximation, often called the Legendre-transform limit.

An equivalent variational argument works for finite discrete states without a saddle point. Choose arbitrary trial probabilities q_i, their mean energy $U[q]$ and Shannon entropy $S[q] = -k_B \sum_i q_i \log q_i$. Subtract the canonical free energy and insert the canonical weights. Algebra gives

$$U[q] - TS[q] - F = k_B T \sum_i q_i \log(q_i/p_i) \geq 0.$$

The final inequality follows from convexity of the logarithmic divergence and becomes an equality only for the canonical distribution on the permitted support. For an isolated system one instead maximises entropy subject to its prepared energy constraint. Helmholtz free-energy minimisation at fixed temperature is thus a consequence of reservoir exchange and normalisation, not an additional force sending every microstate to the lowest energy. At finite temperature high-multiplicity excited states can dominate the equilibrium probability.

Now include volume work in the fundamental differential $dU = TdS - pdV + \mu dN$. Product differentiation gives $dF = -SdT - pdV + \mu dN$, $dG = -SdT + Vdp + \mu dN$, and $d\Omega = -SdT - pdV - Nd\mu$ for $\Omega = F - \mu N$. A potential's natural variables specify which gradient directly returns an observable. From $F(T, V, N)$ one obtains

$$p = -(\partial F/\partial V)_{T,N},$$
$$S = -(\partial F/\partial T)_{V,N};$$

their cross derivatives agree if the potential is smooth. At phase boundaries use phase branches and appropriate coexistence conditions instead of differentiating through a kink as if it were smooth.

Let the reservoir exchange particles as well as energy

With both exchanges allowed, expand reservoir entropy in E and N. Its differential includes $-\mu_R dN_R/T$. A system microstate with E, N therefore acquires weight $e^{-\beta(E-\mu N)}$. Summing normalises the grand canonical distribution. At fixed temperature and volume, derivatives with respect to $\beta\mu$ give mean particle number and number variance. In a homogeneous equilibrium region with suitable thermodynamic limit, $\Omega = -pV$.

$$\Xi = \sum_N e^{\beta\mu N} Z_N,$$

$$\langle N \rangle = \partial_{\beta\mu} \log \Xi,$$

$$\mathrm{var}(N) = k_B T \left(\frac{\partial \langle N \rangle}{\partial \mu} \right)_{T,V}.$$

A surface of M independent adsorption sites, each empty or holding one particle with binding energy $-\epsilon$, has $\Xi = [1 + e^{\beta(\mu+\epsilon)}]^M$. Its occupancy fraction is logistic and its variance $Mf(1-f)$. At $M = 400$ and half coverage, the mean occupancy is 200 with standard deviation 10. The slope with respect to chemical potential is largest there. Near saturation, a small number variance and an apparently quiet detector indicate poor sensitivity to further increases of μ, not necessarily high precision in inferring μ.

Compare the adsorption response slope with the independently predicted number variance

The grand formula links number fluctuations to isothermal compressibility in a bulk one-component fluid. Write $n = \langle N \rangle / V$ and use $dp = n\, d\mu$ at constant T. Then $(\partial n/\partial \mu)_T = n^2 \kappa_T$ and $\mathrm{var}\, N / \langle N \rangle = n k_B T \kappa_T$. A dilute classical ideal gas gives unity. These fluctuations refer to a region exchanging particles. If an observer counts all particles of a strictly sealed closed box, the number variance is zero even though its compressibility is not. The boundary condition of the measurement is part of the formula.

Infer a temperature and explain when it is identifiable

The logarithm of a measured population ratio is linear in β for nondegenerate known levels. With enough independent samples one can estimate temperature either from this ratio or by matching measured energy to its canonical expectation. If all observed states have nearly the same energy, changing β barely changes their probabilities. The available temperature information then becomes small, regardless of detector precision.

For independent canonical draws, differentiating a single log probability gives

$$\partial_\beta \log p_i = -(E_i - \langle E \rangle).$$

Its squared expectation is energy variance, the Fisher information per observation about inverse temperature. At n_s independent draws information adds. Equality between response and noise reappears as a statistical bound on temperature estimation. A long record with correlation time comparable to its length may contain few independent draws despite a large sample count. Recording more frequently without extending the observation window need not improve a temperature estimate.

A pressure reservoir selects enthalpy-weighted volume fluctuations

An isothermal piston in contact with a pressure reservoir supplies a fourth ensemble closely related to the first three. A configuration at volume V costs external work $p_{\mathrm{ext}}V$ in addition to its microscopic system energy. At fixed T, N, p_{ext}, integrate the canonical configurational partition over possible piston positions with weighting $e^{-\beta p_{\mathrm{ext}}V}$. For an ideal gas in a cylinder of fixed cross-sectional area the piston coordinate is proportional to V, so its geometric measure is a constant times dV. Mechanical kinetic factors independent of V cancel from its normalised distribution.

For $Z_N(V) \propto V^N$, the probability density for volume is proportional to $V^N e^{-\beta p_{\mathrm{ext}}V}$ on the positive real axis. Normalisation uses $\int_0^\infty V^N e^{-aV} dV = N!/a^{N+1}$. The maximum, mean and variance are

$$V_{\mathrm{mode}} = \frac{N}{\beta p_{\mathrm{ext}}},$$

$$\langle V \rangle = \frac{N+1}{\beta p_{\mathrm{ext}}},$$

$$\mathrm{var}(V) = \frac{N+1}{(\beta p_{\mathrm{ext}})^2}.$$

The finite-N offset between mode and mean comes from the skewed piston-volume distribution. The most likely volume obeys the usual ideal gas law evaluated using N, while the exact mean differs in this simple piston model. At $N = 99$ the difference is about one percent and relative volume fluctuations are ten percent. At thermodynamic sizes the fractional difference vanishes. When calibrating nanoscale pressure from a histogram, a mean, a peak and a mean of reciprocal volume should therefore not be treated as interchangeable.

Check the force equilibrium directly. The instantaneous gas pressure conditional on V is Nk_BT/V. The gamma distribution has $\langle 1/V \rangle = \beta p_{\mathrm{ext}}/N$, so its mean gas pressure equals the imposed external pressure exactly. Substituting the mean volume into the nonlinear pressure formula instead gives $Np_{\mathrm{ext}}/(N+1)$. This calculation illustrates a general statistical rule: mechanical balance concerns the mean actual force, not necessarily that force evaluated at the mean position.

If the piston lies horizontally and dry static friction is appreciable, configurations can remain stuck with nonzero force imbalance. A distribution fitted to the equilibrium piston formula would then return apparent biases in pressure or particle number. Reverse the pressure change and vary waiting times to expose hysteresis, or remove the mechanical sticking before interpreting volume fluctuations as an equilibrium susceptibility.

A matrix of fluctuations controls coupled responses

Many experiments permit exchange of more than one extensive quantity. Write a general ensemble with natural control parameters θ_a and weights proportional to $e^{\sum_a \theta_a X_a}$, with a fixed set of microscopic state values X_a. The first derivatives of the log normalisation give their means and the second derivatives give the full covariance matrix. Off-diagonal entries describe coupled responses. For the grand canonical ensemble choose $\theta_1 = -\beta$, $\theta_2 = \beta\mu$, $X_1 = E$, $X_2 = N$. One must keep the other natural parameter fixed while taking a partial derivative in this form.

A covariance matrix is positive semidefinite because the variance of any linear combination of the X_a is nonnegative. Consequently $\mathrm{var}(E)\mathrm{var}(N) \geq \mathrm{cov}(E,N)^2$. A measured matrix violating this inequality cannot describe these two observables on the same equilibrium sample space. Possible causes include mixing variances and covariances from different sample preparations, subtracting a poorly estimated instrumental background or an inconsistent energy reference when counting particles.

Suppose a volume exchanges particles each carrying a fixed energy e_0 plus other thermal energy. If the system energy reference shifts by a per particle, new energy is

$$E' = E + aN.$$

Its covariance with N changes by $a\,\mathrm{var}(N)$, although no physical heat changed. A useful combination for isolating fluctuations uncorrelated with particle transport is $E - bN$, with $b = \mathrm{cov}(E,N)/\mathrm{var}(N)$. Minimising the variance with respect to b gives residual variance $\mathrm{var}(E) - \mathrm{cov}(E,N)^2/\mathrm{var}(N)$, invariant under this reference shift. In a general nonlinear conditional distribution that residual is a best linear prediction variance, not automatically the exact variance conditional on a fixed integer N.

For a unit-normalised example, suppose energy variance is $25\epsilon^2$, number variance is 4 and their covariance is 8ϵ. The best linear energy increment per exchanged particle is 2ϵ, and the residual energy variance is $9\epsilon^2$. Replacing energy by $E - 3\epsilon N$ changes the covariance to -4ϵ, but recalculating the residual returns the same value. A data-analysis code can use this reference-invariance as a check on both numerical units and susceptibility derivatives.

The same covariance structure controls simultaneous inference. If two state observables are nearly proportional over all sampled configurations, the log-likelihood curvature has a nearly zero eigenvalue. The individual fields conjugate to them are then hard to determine even if one linear combination is tightly constrained. Preparing the system in a region with different accessible states or adding an independent observable can resolve the degeneracy. Increasing the sampling rate on exactly the same narrow manifold cannot create a missing independent direction.

When the maximum-probability replacement loses a phase

The passage from a Laplace transform to a Legendre transform assumes one dominant smooth contribution or handles multiple branches explicitly. Near a first-order transition two energy ranges can have comparable canonical weights. Replacing the energy distribution by a single Gaussian around its mean then produces a state that may be relatively improbable, lying in the trough between two peaks. The finite partition sum remains smooth, but its derivatives can vary rapidly over a narrowing range of temperature.

In a simple two-phase approximation take branch partition contributions proportional to $e^{-\beta F_1(T)}$ and $e^{-\beta F_2(T)}$. Their ratio is $e^{-\beta(F_2-F_1)}$. At coexistence the leading bulk free-energy densities agree. If their entropy densities differ, then a small temperature change produces an extensive branch free-energy difference proportional to volume times that temperature shift. The equal-weight finite-size window therefore narrows inversely with volume in this simplest picture. Interfacial states, boundaries and phase degeneracy can shift detailed crossing locations.

Consider an energy histogram with equal phase probabilities and sharply peaked branch energies differing by total ΔE. The mixture variance includes $p(1-p)(\Delta E)^2 = (\Delta E)^2/4$ in addition to within-phase fluctuations. As ΔE is extensive this gives an extensive heat-capacity peak per particle in a finite coexistence window, much larger than the smooth single-phase contribution. This does not imply an infinite heat capacity

throughout an interval of bulk temperature. The peak becomes higher and narrower and encodes latent heat in the thermodynamic limit.

An experimental or simulated time series can undersample one branch when interfacial barriers make switching rare. Estimating the variance from one branch then yields a completely different apparent heat capacity from a record containing a few switches. Compare heating and cooling preparations, identify branch-specific energy distributions and establish a reproducible equilibrium weight ratio before treating the union as a canonical histogram. In a macroscopic sample, calorimetric latent heat often provides a more reliable bulk quantity than a noisy attempt to observe whole-sample switching.

Use one equilibrium histogram to predict a nearby temperature

If a canonical run at inverse temperature β_0 samples energies E_i, a nearby expectation at β_1 can be expressed as a ratio of weighted averages with factors $w_i = e^{-(\beta_1-\beta_0)E_i}$. Multiplying numerator and denominator by the same constant allows one to subtract a reference energy to maintain numerical precision. Mathematically the ratio is exact when the original distribution covers every important configuration. Finite data can render an exact identity practically useless outside its overlap window.

For an approximately Gaussian energy histogram of standard deviation σ_E, reweighting shifts its influential energy region by roughly $-(\beta_1 - \beta_0)\sigma_E^2$. Expressed in original standard deviations the shift is $|\beta_1 - \beta_0|\sigma_E$. A value of 0.3 stays near the observed centre; a value of five asks for configurations five original widths away. At the latter value a smooth computed curve can be governed almost entirely by one observation at the extreme edge of the histogram.

Inspect the effective weight count $(\sum_i w_i)^2 / \sum_i w_i^2$ alongside the raw record length. This statistic does not account for time correlation or prove that missing configurations are irrelevant, but a collapse to a handful of weights is a clear failure of overlap. Simulate or measure directly at a second temperature displaced toward the target and compare its actual histogram with the reweighted prediction. This connects the formal equivalence of ensemble descriptions to a concrete question of measurement coverage and uncertainty.

A finite-spin thermodynamic cycle tests mixed derivatives

Consider independent two-state magnetic moments with energies $\mp\mu B$ in a prescribed induction B. Ignore interactions and orbital response. Per moment the canonical sum is $2\cosh x$, $x = \mu B/(k_B T)$. Its Helmholtz free energy is $-k_B T \log(2\cosh x)$ and the average magnetic moment is

$$m = -\partial F/\partial B = \mu \tanh x.$$

The entropy follows from $-\partial F/\partial T$, yielding $S/k_B = \log(2\cosh x) - x\tanh x$. These component equations define an exactly soluble controlled system in which thermal and magnetic responses are coupled.

Equality of the mixed free-energy derivatives predicts

$$(\partial S/\partial B)_T = (\partial m/\partial T)_B.$$

Both are negative for positive field in this model. Increasing field at fixed temperature aligns spins and reduces their entropy. If the spin system is instead isolated reversibly, its entropy is constant, which requires

x constant away from singular limiting cases. Reducing field then reduces temperature in proportion to field. This is the ideal paramagnetic demagnetisation argument, derived here from the canonical weights rather than an empirical cooling rule.

For an idealised system prepared at 3.0 K and 1.5 T, reducing the field quasistatically to 0.50 T under thermally isolated conditions gives 1.0 K. Its mean alignment fraction stays constant even as the Zeeman gap shrinks. A real specimen includes lattice heat capacity, residual interactions and inevitable heat leakage. If the lattice remains equilibrated with the moments, the total entropy includes both subsystems and the simple proportional temperature drop is modified. If spin–lattice exchange is too slow, a single temperature may fail to describe both during the field ramp.

A numerical test can use a rectangular loop in the T, B plane. Evaluate the change in F along one path using the integral of $-S\,dT - m\,dB$, then return along the other path. At successively finer quadrature the closed integral should approach zero. By contrast heat and magnetic work evaluated separately generally have nonzero loop integrals and cancel only in their appropriate state-function combination. A nonzero free-energy residual that does not shrink under refinement points toward a sign error, an inconsistent field-energy convention or different energy references on the two paths.

Finally use the weak-field limit to estimate experimental sensitivity. There $m \simeq \mu^2 B/(k_B T)$, so a thermometer inferred from magnetisation is sensitive to both the magnetic-field scale and the independent number of active moments. Saturation measurements can calibrate their total number, while a field reversal removes a constant magnetometer offset. A mixed-derivative test based on independently measured calorimetry and magnetisation is more discriminating than fitting a temperature parameter to a single magnetic trace that already assumed the predicted population law.

The spin example also gives a direct thermal-noise measurement of susceptibility. For one moment the variance of its projection is $\mu^2[1 - \tanh^2 x]$, and differentiating mean moment with respect to induction gives exactly this variance divided by thermal energy. At $x = 1$, magnetisation is 0.762 of saturation and variance is $0.420\mu^2$. For a thousand independent moments the total rms fluctuation is about 20.5μ, even though their mean total moment is about 762μ. An instrument averaging over many reversals may see the mean accurately and suppress the variance; replaying the same averaging operation on a simulated equilibrium time series makes the comparison meaningful. A direct susceptibility inferred from a very large field modulation would instead sample the curved magnetisation over a finite interval. Use a small symmetric modulation and verify its amplitude scaling to compare with the local fluctuation slope.

Count the configurations of two finite reservoirs

Change the reservoir size and compare the observed energy histogram with a canonical ensemble at the same mean.

Open this experiment online

Try the calculation

Reservoir modes 8

Excitation quanta 24

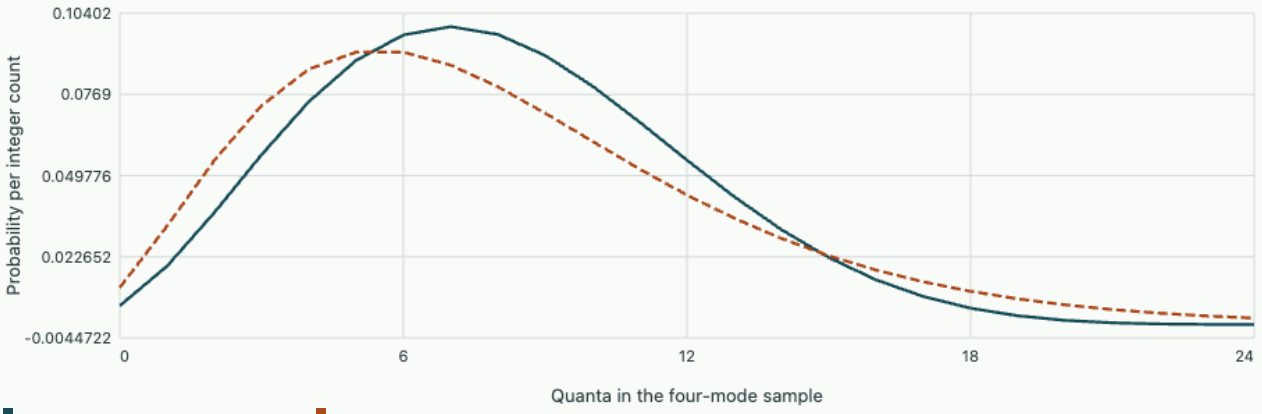

Isolated pair of solids Canonical sample

Mean 8 quanta; closed-pair variance 14.769. Canonical variance 24.

Tutorial · Can a finite bath explain unusually quiet energy data?

A sample has canonical heat capacity $60k_B$ at the mean temperature. Contact with a small reservoir reduces the measured equilibrium energy variance to three quarters of its canonical value. Approximate both entropies quadratically near their common equilibrium and infer the bath heat capacity. What other measurement could mimic the reduction?

▼ **Add the two inverse susceptibilities and check the observation**

The joint log probability has curvature $-[1/C_s + 1/C_R]/(k_B T^2)$. Hence the effective variance is $k_B T^2 C_s C_R/(C_s + C_R)$. Dividing by $k_B T^2 C_s$ and setting the ratio to $3/4$ gives

$$C_R = 3C_s = 180k_B.$$

The temperature at maximum probability still agrees for bath and sample at leading order. The suppressed fluctuations reflect their shared total-energy constraint.

Finite temporal averaging of a fast energy signal can also suppress measured variance. Vary detector bandwidth and check convergence of the measured area, then repeat with a different independently estimated reservoir capacity. A corrected variance scaling with the harmonic combination supports the finite-bath interpretation. A single reduced number cannot distinguish thermodynamic curvature from the detector response.

Tutorial · Determine a volume response from temperature data

For a fixed amount of a smooth single-phase material near one operating point, the thermal expansion coefficient is α, the isothermal compressibility is κ_T and the volume is V. Derive $C_p - C_V$ from a Maxwell relation. Evaluate it for $T = 300$ K, molar volume 1.0×10^{-5} cubic metres, $\alpha = 6.0 \times 10^{-5}$ per kelvin and $\kappa_T = 1.0 \times 10^{-11}$ per pascal.

▼ **Eliminate the unmeasured mixed entropy derivative**

Write $dS = (C_V/T)dT + (\partial p/\partial T)_V dV$, using the Helmholtz mixed derivative. On the constant-pressure path $dV = \alpha V dT$. Holding volume fixed in $dV/V = \alpha dT - \kappa_T dp$ gives

$$(\partial p/\partial T)_V = \alpha/\kappa_T.$$

Therefore $C_p - C_V = TV\alpha^2/\kappa_T = 1.08$ J per mole kelvin. Positive isothermal compressibility and positive temperature imply the difference is nonnegative even for negative expansion coefficient.

Tutorial · Infer two unknowns from occupancy measurements

At two distinct temperatures a set of dilute independent adsorption sites gives coverages 0.20 and 0.60. Why can these two readings fail to identify binding energy if the reservoir chemical potentials are not independently known? Express the unknown combination and suggest an experiment that would determine the binding energy.

▼ **Write an equation for each observation before counting parameters**

The log odds satisfy

$$\log[f/(1-f)] = (\mu + \epsilon)/(k_B T).$$

The measured log odds are minus 1.386 and plus 0.405. They determine the value of $\mu + \epsilon$ in each run multiplied by its known inverse thermal energy. With two unrelated reservoir potentials there are three unknown dimensional quantities and only two equations. Increasing coverage precision cannot resolve that degeneracy.

A separately calibrated reservoir supplies $\mu(T, p)$ relative to a consistent energy reference. Then several pressures at one temperature verify the model and several temperatures can test a constant binding energy against an additional binding entropy or site degeneracy. A reported equilibrium constant alone can always absorb a reference-state shift; the reference reservoir and site standard state are part of a physically interpretable adsorption energy.

Cambridge Statistical Physics · canonical and grand canonical ensembles

Rotational and vibrational populations in a molecular gas

Statistical mechanics from microscopic counting to material response

A pressure thermometer and a calorimeter do not count molecular degrees of freedom in the same way. A dilute diatomic gas can retain nearly ideal pressure while its heat capacity changes appreciably as discrete internal levels become populated. Build the molecular partition function, derive the rotor and oscillator contributions, and infer temperatures from line populations without treating spectral intensity as a direct state count.

When a molecule's partition function factorises

Assume a dilute gas of stable diatomic molecules with weak interparticle interactions. Approximate one molecule's Hamiltonian by independent translational, rotational, vibrational and electronic parts. This requires a bond sufficiently stiff that rotation does not greatly alter vibrational energy and a temperature well below electronic excitations or dissociation when the electronic factor is treated as constant. If energies add and allowed states form a product, Boltzmann factors multiply. Thus $q = q_{\mathrm{tr}} q_{\mathrm{rot}} q_{\mathrm{vib}} q_{\mathrm{el}}$, $Z_N \simeq q^N / N!$ in the dilute classical translational regime. The logarithm converts this product into additive free energies, energies and heat capacities.

The translational factor V/λ_T^3 provides the volume derivative and hence ideal pressure $N k_B T / V$. A volume-independent internal factor changes neither that derivative nor the associated translational equation of state. It does change entropy, chemical potential and calorimetry. The molar heat capacities satisfy $C_p - C_V = R$ for a fixed-composition ideal gas, even when both vary strongly with temperature. Dissociation, chemical equilibration or significant intermolecular interactions can change the conclusion for the measured sample.

One danger is double-counting states. Angular momentum orientations belong to rotational degeneracy, oscillator quanta to vibration, electronic spin and nuclear exchange symmetry may constrain one another. A constant electronic degeneracy contributes an additive entropy $N k_B \log g$ but no internal energy or heat capacity at fixed level energies. A large heat capacity is evidence for energetically separated states becoming populated, not just for a large number of exactly degenerate states.

Derive the quantum rigid-rotor contribution

For a molecule of moment of inertia I about its centre of mass, rotational Hamiltonian is $\hat{L}^2/(2I)$. Angular momentum eigenvalues are $\hbar^2 J(J+1)$, with $2J+1$ projections at fixed $J = 0, 1, \ldots$. Define rotational temperature $\theta_r = \hbar^2/(2 I k_B)$. For distinguishable nuclei and a common nuclear spin factor omitted, the rotational sum and populations are

$$q_r = \sum_{J=0}^{\infty} (2J+1) e^{-\theta_r J(J+1)/T},$$

$$p_J = (2J+1) e^{-\theta_r J(J+1)/T} / q_r.$$

At high temperature the populated range contains many adjacent J values. Replacing the sum with an integral and substituting $u = J(J+1)$ gives

$$q_r \simeq T/\theta_r.$$

Hence $U_r \simeq Nk_BT$, $C_{V,r} \simeq Nk_B$, the classical result for two independent rotational kinetic energies of a linear object. At low temperature only $J = 0$ and $J = 1$ matter. Their degeneracies and gap $2k_B\theta_r$ give $q_r \simeq 1 + 3e^{-2\theta_r/T}$ and an exponentially vanishing heat capacity. Keeping the degeneracy three is essential for both its magnitude and the temperature of its turn-on.

For an illustrative heteronuclear rotor with $\theta_r = 85$ K, direct summation at 20 K gives $q_r = 1.00061$ and $C_{V,r} = 0.0440k_B$ per molecule. At 90 K the corresponding values are 1.471 and 1.063. At 300 K the heat capacity is 1.002. A small overshoot above the classical rotational value is possible during the crossover. A model that simply switches from zero to one classical degree contribution at $T = \theta_r$ misses the entire smooth population redistribution and cannot predict line ratios.

The most populated J need not be zero, despite increasing energies, because the projection multiplicity increases with J. Viewing the numerator as a smooth function of J gives $J_{\mathrm{peak}} \simeq [\sqrt{2T/\theta_r} - 1]/2$, to be checked against neighbouring integer values. The estimate is unreliable if it predicts only one or two populated levels. Computing a short discrete sum costs little and reveals its own truncation error through the final omitted weights.

Derive oscillator energy and heat capacity from a geometric series

Near its bond-length minimum approximate the interaction by a spring. With reduced mass m_r and angular vibration frequency ω_v, quantum energies are $E_n = \hbar\omega_v(n + 1/2)$. Define $\theta_v = \hbar\omega_v/k_B$. Summing the resulting geometric series and differentiating gives

$$q_v = \frac{e^{-\theta_v/(2T)}}{1 - e^{-\theta_v/T}},$$

$$\frac{U_v}{N} = k_B\theta_v\left(\frac{1}{2} + \frac{1}{e^{\theta_v/T} - 1}\right),$$

$$\frac{C_{V,v}}{Nk_B} = \frac{x^2e^{-x}}{(1 - e^{-x})^2},$$

$$x = \theta_v/T.$$

The zero-point term contributes energy but no heat capacity at fixed frequency. For $x \ll 1$ the heat capacity approaches k_B per independent vibrational mode, comprising kinetic and potential contributions. For $x \gg 1$ excitation is too costly and the heat capacity behaves as $k_Bx^2e^{-x}$. A molecule can still vibrate quantum mechanically in its ground state while contributing virtually no vibrational heat capacity. A calorimeter measures sensitivity of mean energy to temperature and need not see the ground-state motion.

Take $\theta_v = 3000$ K. At $T = 300$ K the vibrational contribution is only $0.00454k_B$, while at 1200 K it reaches $0.609k_B$. Combining with the earlier rotor gives molar C_V approximately $2.507R$ at 300 K and $3.109R$ at 1200 K, assuming stable composition and frozen electronic structure. For acoustic measurements the ideal adiabatic ratio $\gamma = (C_V + R)/C_V$ then decreases from about 1.399 to 1.322. An acoustic drive faster than vibrational relaxation can measure a frequency-dependent effective ratio because internal populations fail to remain equilibrated through a compression cycle.

Compare this single oscillator with the frequency distribution in the lattice heat-capacity plot

Nuclear exchange and the symmetry number

For a homonuclear molecule, rotating its classical orientation by half a turn end for end does not create a second distinct arrangement. In a high-temperature rotational count a symmetry number two therefore divides the distinguishable-nuclei rotor partition function once the compatible spin states are accounted for. Since this factor is constant with temperature it shifts entropy and chemical potential, but not the classical rotational energy or heat capacity. At low temperature it is important to treat actual quantum parity, not divide every level by two.

A homonuclear diatomic rotational wavefunction acquires a factor $(-1)^J$ under exchange of nuclei. The complete state also includes nuclear spin, electronic and vibrational parts. For ordinary molecular hydrogen in its electronic and vibrational ground state the protons are spin-half fermions. Antisymmetry pairs even J with the single antisymmetric spin singlet and odd J with the three symmetric triplet spin states. In true spin-conversion equilibrium the rotor sum uses weights one or three in addition to $2J+1$. At temperatures large compared with rotational spacing, even and odd J sectors have nearly equal rotational sums and the equilibrium odd-to-even population ratio approaches three.

Conversion between spin species can be slow. A rapidly cooled sample may preserve its prepared fractions rather than attain spin equilibrium. Then compute separate parity-restricted partition functions for fixed species numbers, and average their thermal observables with the frozen proportions. Allowing the total odd fraction to relax yields a different heat capacity, including the energy liberated by conversion. Two experiments on chemically identical hydrogen can consequently report different heating curves because their preparation constrains different internal sectors.

Use intensities only after accounting for the instrument

An absorption line originating from level J has a strength depending on more than its population. Transition moment, polarisation, transition frequency, stimulated emission correction, instrument response and path length can all enter. Suppose line strengths and detector factors have been divided out to yield quantities proportional to level populations N_J. Then a straight-line fit of $\log[N_J/(2J+1)]$ against $J(J+1)$ has slope $-\theta_r/T$ for the simple heteronuclear model. Its intercept depends on total particle number, partition function and residual instrumental gain.

For $\theta_r = 10$ K, a measured corrected ratio $N_2/N_0 = 2.00$ implies $2 = 5e^{-60/T}$ and $T = 65.5$ K. Omitting the factor of five would make an upper population larger than the lower one appear incompatible with positive temperature. The inferred temperature assumes the rotational populations are equilibrated at one temperature; an optically pumped or rapidly expanding sample may have different rotational and translational distributions.

A plot over several J values can test whether one slope is adequate. Curvature may come from centrifugal distortion, line saturation, unresolved overlap, multiple temperatures or a wrong nuclear-spin degeneracy. Alternating residuals for even and odd J indicate a different physical problem from a smooth quadratic residual against $J(J+1)$. Repeat with lower optical depth to isolate saturation, and use independently measured transition frequencies to fix level energies before fitting populations.

Recover molecular translation separately from internal structure

In a low-density molecular beam, the translational speed distribution is often easier to measure than the internal populations. For an equilibrium gas leaving a small opening into vacuum, the molecular flux is weighted by the normal velocity. Fast molecules reach and cross the aperture more frequently, so the distribution of speeds in the beam differs from the Maxwell speed distribution inside the reservoir. The mean translational energy of molecules in an effusive beam is $2k_B T$, whereas the bulk mean is $3k_B T/2$. This selectivity occurs even when rotational and vibrational populations remain canonical.

To derive the difference, start with the bulk velocity density proportional to $e^{-mv^2/(2k_B T)}$. Write the outward normal speed as $v\cos\theta$, integrate over the outward hemisphere and include the velocity-space shell factor $v^2 dv$. The resulting beam speed density is proportional to $v^3 e^{-mv^2/(2k_B T)} dv$, while the original bulk speed density has v^2. Taking a ratio of Gaussian moments gives the stated energy. Comparing a beam energy to the wrong bulk prefactor would yield a temperature one third too high.

If the aperture is small compared with the mean free path and transition-detection efficiency is independent of speed after correction, internal state probabilities factor out of this flux weighting. A rotational population ratio can then report the reservoir temperature while the observed mean translational energy follows the effusive weighting. In a dense nozzle undergoing many collisions, expansion can cool translation and rotation differently. An inferred rotational and translational temperature disagreement is therefore not automatically bad calibration; first identify whether the beam is effusive, collisionally expanding or selectively detected.

As a bench example, a beam mean translational energy of 0.060 eV corresponds to approximately 348 K under effusion, using $k_B = 8.6173 \times 10^{-5}$ eV/K. Treating it as the bulk mean would give about 464 K. Compare against a population ratio from two corrected lines spanning $80k_B$ kelvin and equal degeneracies. At the two temperatures the predicted ratios are about 0.795 and 0.842. Resolving this difference demands a relative line calibration appreciably better than six percent even though the two erroneous temperature interpretations differ by more than one hundred kelvin.

An isotope mixture reveals which spectral changes affect entropy

Compare two isotopic forms with the same approximate electronic potential surface. Bond force constant and equilibrium bond length then remain nearly unchanged while reduced nuclear mass changes. The harmonic vibrational frequency scales as $m_r^{-1/2}$, while the rotational temperature scales as m_r^{-1}. Consequently vibration and rotation respond differently to the same isotope substitution. Their shifts can jointly test whether two sets of spectral features genuinely belong to the same molecular carrier.

Suppose the reduced mass rises by a factor 1.21. A vibrational line at 2000 inverse centimetres moves approximately to 1818 inverse centimetres, and a rotational constant initially 2.00 inverse centimetres moves to 1.65 inverse centimetres. The vibrational characteristic temperature follows from $hc\tilde{\nu}/k_B$, about 1.4388 K per inverse centimetre. At room temperature the two vibrational gaps still greatly exceed thermal energy, so the vibrational heat capacity can remain small even though the spectral isotope shift is easy to resolve.

The zero-point energies nevertheless differ. For this vibrational example their difference is $hc(2000 - 1818)/2$ per molecule, corresponding to about 131 kelvin in thermal units or 1.09 kJ/mol. Such offsets cancel within a fixed-species heat-capacity derivative but can matter to comparisons of chemical equilibrium or isotopic exchange. An analysis that discards all constant energy offsets independently in each

molecular species can corrupt a reaction free energy unless it restores a consistent reference across reactants and products.

For a thermal entropy contribution from this one oscillator use $S_v = k_B[(n+1)\log(n+1) - n\log n]$ with $n = 1/(e^x - 1)$. This follows by inserting U_v and $\log q_v$ into $S = k_B \log q + U/T$; the zero-point contributions cancel exactly. Thus the measured isotope shift can change vibrational entropy via occupancy even while the zero-point offset itself contributes no entropy. It is helpful to evaluate all three quantities—energy difference, entropy difference and heat-capacity difference—rather than presume that a line shift must have the same thermodynamic role in each.

An internal relaxation time appears in frequency-dependent calorimetry

The equilibrium partition function assumes internal mode populations re-equilibrate on times short compared with the measurement. Model one slower internal degree by an internal energy U_s relaxing toward its equilibrium value at the instantaneous translational temperature T, with $\dot{U}_s = [U_{s,\mathrm{eq}}(T) - U_s]/\tau_s$. For a small oscillation around T_0,

$$\delta U_{s,\mathrm{eq}} = C_s \delta T.$$

In the $e^{-i\omega t}$ convention, the complex internal heat capacity and the total measured capacity including an instantaneous background C_f are

$$C^*(\omega) = C_f + \frac{C_s}{1 - i\omega\tau_s}.$$

At low frequency the capacity is $C_f + C_s$. At high frequency the slow internal energy cannot follow and the apparent in-phase capacity approaches C_f. Its out-of-phase component peaks at $\omega\tau_s = 1$, where real and imaginary increments equal $C_s/2$. The structure resembles dielectric relaxation because both arise from one relaxing state variable driven by a controlled sinusoid. It follows here from an explicit energy-exchange model and its assumptions.

Take a molar fast contribution $2.50R$, slow equilibrium contribution $0.60R$ and relaxation time 2 ms. The full equilibrium value is $3.10R$. At about 79.6 Hz the real measured capacity is $2.80R$ and the imaginary increment magnitude $0.30R$. A heat pulse much shorter than two milliseconds initially raises the temperature as if the capacity were close to $2.50R$, then energy flows into the internal mode. A slow calorimetric ramp can be perfectly consistent with a faster acoustic measurement that infers a smaller active capacity.

A sample container connected to a thermal bath produces its own complex temperature response even if molecular populations are always equilibrated. Disentangle those effects by calibrating the empty cell or an inert reference gas across the same frequency range, changing pressure to alter collision-mediated internal-relaxation time, and checking the predicted equilibrium limit from independently known level spacings. A mechanical resonance in the apparatus can be separated by varying cell geometry while keeping gas composition fixed. Fitting one imaginary peak to a molecular relaxation model without these controls does not uniquely determine the microscopic relaxation time.

How a line ratio and a calorimeter constrain degeneracies differently

Two groups report a low-energy excited manifold with dimensionless population ratio $r = N_1/N_0$ at two temperatures. A single ratio determines only $r = (g_1/g_0)e^{-\Delta/(k_B T)}$, so changing the assigned degeneracy changes the inferred gap. Taking the log ratio at two known temperatures removes a common degeneracy factor. The slope of $\log r$ against $1/T$ is minus the gap in kelvin, and the intercept is the log degeneracy ratio if the line-strength calibration and equilibrium assumptions hold.

For illustrative values $r = 0.300$ at 100 K and $r = 0.949$ at 200 K, the log change is approximately 1.152. Dividing by the inverse-temperature change magnitude 0.005 per kelvin gives $\Delta/k_B \simeq 230$ K. Inserting the first ratio gives $g_1/g_0 \simeq 3.00$. The same two-manifold model predicts a dimensionless heat capacity $x^2 r/(1+r)^2$ per molecule. At 100 K it is about 0.94, a separate calorimetric prediction independent of the absolute number of detected photons.

A third temperature materially improves the test. An unseen nearby manifold or a line strength that changes with instrumental settings may fit two exact population ratios but produce curvature and excess heat capacity elsewhere. Spectroscopy resolves ordered energy differences while calorimetry sums energy variances over all accessible states. Their combination offers more information than adding further precision to one population measurement under the same ambiguous degeneracy assumption.

Compress a gas without pretending its heat capacity stays constant

A reversible adiabatic compression is often summarised by a fixed exponent γ. That power law assumes temperature-independent heat capacities along the path. Molecular populations give a direct case where the exponent changes. For an ideal gas at fixed composition,

$$dS = C_V(T)dT/T + Nk_B\,dV/V.$$

Along a reversible adiabatic path set the sum to zero and integrate. The volume ratio is determined by an integral of measured or calculated $C_V(T)/T$, not by a universal constant exponent across every temperature.

As an illustrative calculation approximate translation and fully active rotation by $C_V = 2.5Nk_B$ plus one harmonic vibrational mode with $\theta_v = 3000$ K. Heating from 300 to 1200 K at constant entropy gives a volume contraction factor from the fast degrees alone of $4^{2.5} = 32$. The vibrational entropy difference adds another factor $e^{\Delta S_v/(Nk_B)}$. At 300 K the oscillator entropy is only about $0.00050k_B$, while at 1200 K it is about $0.309k_B$. The combined estimate requires $V_i/V_f \simeq 43.6$, a substantially larger compression than the frozen-vibration approximation predicts.

This comparison assumes ideal gas throughout and no dissociation, conditions to be checked in a real large compression. It nevertheless identifies the mechanism: part of the compression work goes into levels that were effectively inactive initially. Using the final high-temperature heat capacity across the whole path would err in the opposite direction by assigning that energy storage to the low-temperature portion as well. The entropy integral correctly weights each temperature interval by its actual internal population response.

A numerically stable solution can avoid differentiating noisy calorimetry. Integrate the entropy increment from tabulated heat capacities in the log-temperature variable. Compare an adaptive integral with a refined straightforward grid, paying particular attention to the rotational or vibrational crossover scale. Alternatively compute endpoint entropies from partition functions when an independently justified spectral model exists.

Agreement between calorimetric integration and spectroscopy-derived endpoint entropy checks both assignments without demanding exact agreement of noisy pointwise derivatives.

Experimental pressure also checks the endpoint calculation independently. After equilibration, the ideal equation of state predicts $p_f/p_i = (T_f/T_i)(V_i/V_f)$, approximately 174 for this idealised path. Finite heat leakage or a pressure-dependent nonideality can change that value. A rapid adiabatic compression can be irreversible even if little heat leaves the chamber, and then the measured entropy can increase. Distinguish a reversible reference calculation from an actual rapid piston protocol before using a mismatch to infer missing molecular levels.

An independent isotropic acoustic test at a fixed initial temperature infers γ from the small-amplitude sound speed $c_s^2 = \gamma RT/M$, where M is molar mass. At fixed molecular mass a decrease of γ from 1.40 to 1.32 at the same temperature would lower sound speed by the square root of their ratio, roughly three percent. That comparison must truly hold temperature fixed: the large direct thermal dependence of sound speed can easily exceed the internal-mode effect. Frequency-resolved acoustics can see part of a vibrational mode freeze dynamically while equilibrium spectroscopy still predicts its population. Record both phase velocity and acoustic attenuation if the aim is to identify a relaxation time. Across the relaxation band, dispersion without the corresponding dissipative response should prompt a check of transducer phase calibration and acoustic boundary conditions.

Read the rotational level spacing from thermal energy storage

Sum the discrete levels of a linear rotor and compare the heat capacity with equipartition as temperature rises.

Open this experiment online

Try the calculation

Rotational temperature θrot (K) **8**

Chosen gas temperature (K) **20**

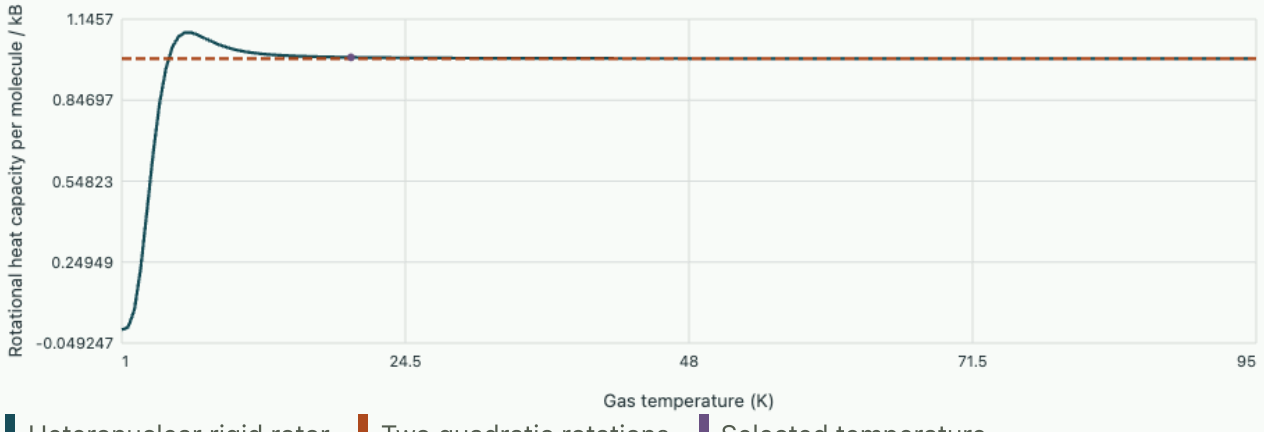

First excited rotational level has E_1/kB = 16 K. At the selected temperature Crot/kB = 1.0051. Vibrational and translational contributions are excluded.

Tutorial · Identify a rotor inertia from spectral spacing

For a heteronuclear linear rotor the lowest pure rotational transition has ordinary frequency 120 GHz. Infer θ_r, the moment of inertia and the expected $J = 1$ to $J = 0$ population ratio at 10 K. How could an isotope substitution test the interpretation?

▼ **Use the transition gap and keep angular and ordinary frequency distinct**

The gap is $2k_B\theta_r = h\nu$, so $\theta_r = 2.88$ K and $I = \hbar^2/(2k_B\theta_r) = 1.40 \times 10^{-46}$ kg square metres. The population ratio is

$$3e^{-5.76/10} = 1.69.$$

The upper manifold contains more molecules even though each of its levels has lower individual Boltzmann weight. Increasing reduced mass at nearly fixed bond length increases I, reducing the transition frequency approximately inversely with reduced mass. A vibrational frequency from the same spring would scale instead with inverse square root of reduced mass.

Tutorial · Infer an excitation temperature from calorimetry

A gas of linear molecules has molar $C_V = 2.91R$ at a temperature where translation and rotation contribute $2.50R$. Neglect electronic excitation and estimate the characteristic temperature of a single harmonic vibration in units of measurement temperature. Would a measured C_p be usable unchanged?

▼ **Invert the oscillator contribution and verify its feasible range**

The remaining $0.41R$ lies between zero and the one-mode classical limit R. Solve $x^2e^{-x}/(1-e^{-x})^2 = 0.41$ on positive x. At 3.0 it is 0.496 and at 3.5 it is 0.393, placing the root near 3.41. Thus θ_v is about 3.41 times measurement temperature. Use several temperatures to verify the characteristic value and frozen-composition model. For ideal gas at fixed composition, subtract R from C_p before applying this constant-volume balance. Latent processes, reaction and internal relaxation would require an extended model.

Tutorial · Distinguish equilibrium hydrogen from a frozen spin mixture

At low temperature only $J = 0$ in the even sector and $J = 1$ in the odd sector contribute. Use $\theta_r = 85$ K and $T = 30$ K. Estimate the equilibrium odd fraction and compare the energy stored in a sample frozen at a three-quarter odd fraction, measured relative to the common lowest even state.

▼ **Normalise both sectors and identify the energy released on conversion**

The relative weight of $J = 1$ is $3(2J+1)e^{-170/30} = 0.0311$; the even ground state has weight one. The equilibrium odd fraction is approximately 0.0302. A frozen sample at odd fraction 0.75 stores about

$0.75(170)k_B = 127.5k_B$ kelvin of rotational energy per molecule, versus about $5.1k_B$ kelvin for equilibrium. The excess expressed per mole is roughly $R(122.4\,\mathrm{K}) = 1.02$ kJ/mol.

The frozen low-temperature heat capacity can be small even while substantial energy remains trapped in the odd sector. A calorimeter monitoring a slowly converting sample measures time-dependent heat release, not simply the derivative of a stationary restricted partition function. Measure the spin fraction independently or vary equilibration time before comparing its heat curve with an equilibrium prediction.

Cambridge Statistical Physics · molecular partition functions and internal degrees of freedom

Bose condensation, dimensionality and a finite trap

Statistical mechanics from microscopic counting to material response

Cooling a gas of identical bosons can populate a single orbital macroscopically. The temperature at which that occurs depends on the low-energy density of states, not simply on the energy needed to put one particle in the first excited level. Compare a uniform box with a harmonic trap, estimate a condensate fraction and design an imaging test that separates occupation of a mode from a narrow but entirely thermal cloud.

Derive the occupation before approximating the density of states

A single noninteracting bosonic orbital of energy ϵ_j may contain $n_j = 0, 1, 2, \ldots$ particles. In contact with a reservoir of temperature T and particle chemical potential μ, its grand partition factor is the geometric series $\sum_n e^{-\beta n(\epsilon_j - \mu)}$. Convergence requires $\mu < \epsilon_0$, where ϵ_0 is the ground energy. Differentiation gives mean occupancy $1/[e^{\beta(\epsilon_j - \mu)} - 1]$. Shift the energy origin to $\epsilon_0 = 0$, so fugacity $z = e^{\beta\mu} < 1$. For a fixed mean particle number solve $\sum_j \langle n_j \rangle = N$.

As temperature falls, the distribution narrows toward low energies and μ generally rises toward the ground energy. It does not cross that energy in an equilibrium ideal finite gas. A large mean ground-state population corresponds to $1 - z$ of order $1/N_0$. Once N_0 is comparable to the total, replacing that single low-lying orbital by a continuum density of states can erase the very population of interest. Separate it before approximating the excited states.

A bosonic occupancy of one per quantum state is not an upper limit. At classical high temperature typical occupation of each relevant orbital is small and $n_j \simeq z e^{-\beta\epsilon_j}$. The statistical difference from a dilute classical gas becomes appreciable when the thermal de Broglie volume is comparable to the volume per particle. This criterion indicates growing exchange effects; the existence of a finite-temperature condensation transition additionally depends on dimension, dispersion and trapping geometry.

Count the three-dimensional excited particles

Use periodic boundary conditions in a large cubic box with side L. Allowed wavevectors occupy cells $(2\pi/L)^3$ in reciprocal space. For one spin species, counting the spherical shell and substituting $\epsilon = \hbar^2 k^2/(2m)$ gives an excited-state density proportional to $V\sqrt{\epsilon}$. Expand $1/(z^{-1}e^{x} - 1) = \sum_{\ell \geq 1} z^{\ell} e^{-\ell x}$ before evaluating the radial Gaussian integrals. In terms of the thermal wavelength and Bose functions $g_s(z) = \sum_{\ell \geq 1} z^{\ell}/\ell^{s}$, the result is

$$\frac{N_{\mathrm{ex}}}{V} = \lambda_T^{-3} g_{3/2}(z),$$

$$p = k_B T \lambda_T^{-3} g_{5/2}(z),$$

$$U_{\mathrm{ex}} = \frac{3}{2} p V.$$

Pressure follows from the grand potential or the kinetic virial identity in a uniform nonrelativistic gas. The ground state at zero momentum contributes no extensive kinetic energy or pressure in the thermodynamic continuum expression. It can nevertheless carry a macroscopic particle number. These formulas describe an ideal gas: interactions can add pressure and shift excitations even at very low temperature.

The series $g_{3/2}(1) = \zeta(3/2) = 2.6124$ converges. At a fixed temperature the excited population per volume therefore has a maximum. At fixed total density n, once this maximum falls below n, the remaining particles accumulate in the ground orbital. Equating the two numbers defines the bulk critical temperature, and the below-transition condensate fraction follows:

$$T_c = \frac{2\pi\hbar^2}{m k_B}\left[\frac{n}{\zeta(3/2)}\right]^{2/3},$$

$$\frac{N_0}{N} = 1 - (T/T_c)^{3/2}, \quad T < T_c.$$

For a single-state gas of atoms with mass 87 unified atomic mass units at density 10^{20} per cubic metre, $T_c \simeq 398$ nK in this ideal box model. At $0.70 T_c$ the predicted condensate fraction is 0.414. Doubling density raises T_c by a factor $2^{2/3} = 1.587$. If imaging calibration rescales density but the reported temperature is independent, such a shift can move an inferred reduced temperature significantly. A finite trap or several spin components must be modelled before turning these example numbers into a laboratory calibration.

Why a uniform two-dimensional ideal gas is different

For free quadratic dispersion in d dimensions the density of states per volume scales as $\epsilon^{d/2-1}$. As $\mu \to 0^-$, occupancy behaves like $k_B T/\epsilon$ at low energy. The excited-population integral near its lower limit therefore behaves as $\int_0 \epsilon^{d/2-2} d\epsilon$. It converges only for $d > 2$. A uniform three-dimensional ideal gas saturates; in two dimensions the integral diverges logarithmically and in one dimension even more strongly. In the infinite uniform two-dimensional ideal gas at fixed density a finite positive temperature does not require a macroscopic ground occupation.

The two-dimensional integral can be evaluated explicitly. Per unit area $n_{2D} = -\lambda_T^{-2}\log(1-z)$. At phase-space density $n_{2D}\lambda_T^2 = 5$ this gives

$$z = 1 - e^{-5} = 0.99326.$$

In the formal grand description the mean ground occupation would be about 147 for this chemical potential, large on an atomic scale but still independent of total area at fixed density in the thermodynamic limit. Enlarging a system is how one distinguishes a finite prominent occupation from an extensive one.

A strictly finite box has discrete energies, so even a nominally two-dimensional sample can show a finite-size crossover with a narrow central state. Gapped transverse excitations may also make a nominally planar cloud effectively three-dimensional at sufficiently high temperature. Interacting two-dimensional fluids can exhibit a vortex-mediated transition involving long-range phase correlations of a different kind. The ideal uniform density-of-states argument neither derives nor rules out that interacting transition.

Count a harmonic trap in phase space

In a three-dimensional harmonic trap with frequencies $\omega_x, \omega_y, \omega_z$, subtract the zero-point ground energy. Excited energies are $\hbar(\omega_x n_x + \omega_y n_y + \omega_z n_z)$ with nonnegative integers. When $k_B T$ exceeds all the level spacings by a large factor, approximate the number of states below energy ϵ by the volume of a positive octant tetrahedron in the three occupation-number coordinates. Its derivative yields

$$D_{\mathrm{trap}}(\epsilon) = \frac{\epsilon^2}{2\hbar^3\omega_x\omega_y\omega_z},$$

$$N_{\mathrm{ex}} = \left(\frac{k_B T}{\hbar\bar{\omega}}\right)^3 g_3(z),$$

$$\bar{\omega} = (\omega_x\omega_y\omega_z)^{1/3}.$$

Condensation now gives

$$k_B T_c = \hbar\bar{\omega}[N/\zeta(3)]^{1/3},$$

$$N_0/N = 1 - (T/T_c)^3.$$

In a trap with geometric mean frequency 100 Hz and $N = 10^5$, the ideal estimate is 210 nK. At 0.70 of this trap-specific critical temperature the condensate fraction is 0.657, appreciably greater than the 0.414 box result at the same reduced temperature. Mixing a trap density of states with a homogeneous critical formula cannot be repaired by inserting one arbitrarily chosen cloud volume.

For a harmonic trap in d dimensions the low-energy density behaves as ϵ^{d-1}. Excited-state saturation in this semiclassical trapped model already occurs for $d > 1$. The changed criterion arises because a higher-energy particle also explores more physical space in the trap. The thermodynamic limit must be specified: holding all trap frequencies fixed while adding arbitrarily many particles changes the central density. A conventional trapped limit sends the trap frequencies to zero with a suitable combination such as $N\bar{\omega}^3$ fixed.

Connect a narrow peak to occupation and spatial coherence

After an ideal noninteracting cloud is released, a particle initially at $\mathbf{r}_0$ and momentum $\mathbf{p}$ reaches $\mathbf{r}_0 + \mathbf{p}t/m$. At long expansion times the position distribution maps approximately to momentum distribution after accounting for initial cloud size. A condensate occupying one trap orbital can therefore produce an extended but narrow component in a momentum image superposed on a broader thermal distribution. Ballistic expansion length, resolution and interaction-driven expansion all affect the observed width.

For the 87-unit-mass example at 200 nK a one-axis thermal velocity scale is

$$\sqrt{k_B T/m} = 4.37\,\mathrm{mm/s}.$$

After 20 ms its thermal contribution to one-axis rms width is about 87 micrometres. A point-spread width of ten micrometres is a small correction for this broad part but can dominate a genuinely much narrower central component. Fit integrated particle number and the spatial response of each component. A Gaussian peak height alone increases when the same number of atoms occupies a narrower instrument function.

A more fundamental criterion uses the one-particle density matrix. Its eigenvalues are orbital occupations; a condensate has an eigenvalue proportional to the total number as the appropriate system size grows. Interference measurements of first-order coherence or occupation scaling over different atom numbers can test this picture. Condensation is not identical to zero viscosity or a critical flow speed. An ideal uniform quadratic gas has low-energy excitations with arbitrarily small energy per momentum and does not by itself predict the finite critical speed of an interacting superfluid.

Solve the discrete number equation before passing to a continuum

A small cold gas in a box is an excellent numerical test because the discrete energy levels, their degeneracies and their individual occupancies are known. For a periodic cube with ground energy zero, write its excited energies as $\epsilon_{\mathbf{j}} = \epsilon_L(j_x^2 + j_y^2 + j_z^2)$, where $\epsilon_L = 2\pi^2\hbar^2/(mL^2)$. The first shell has six states, the next twelve and the third eight. Replacing each shell by one level without its vector degeneracy artificially drives particles into the ground state.

For a specified finite number N, solve $N = \sum_{\mathbf{j}} [e^{\beta(\epsilon_{\mathbf{j}} - \mu)} - 1]^{-1}$ with a chemical potential strictly below zero. This is a monotone root problem: as μ goes from minus infinity toward zero the sum grows from zero without bound owing to the ground state. Use an energy cutoff well above the available thermal energy and estimate the omitted tail, increasing the cutoff to check convergence. Near zero, evaluate $e^x - 1$ with a numerically stable small-argument function. A naive evaluation of $e^x - 1$ can erase the tiny denominator representing a macroscopically occupied orbital.

A useful artificial example has $\beta\epsilon_L = 2$ and a mean ground occupation of 99. This fixes $z = 99/100 = 0.99$. Each of the six first-shell levels contains on average $1/[e^2/0.99 - 1] = 0.1548$ particles, contributing 0.929 in total. The twelve next-shell levels at $4k_BT$ contribute another 0.222 together. Their combined occupation is already greater than one even though every individually excited orbital is weakly occupied. A code can compare an explicit vector enumeration with a shell-degeneracy table and verify exact agreement for the same cutoff.

The finite grand-canonical solution fixes mean N, not the integer count in each run. A canonical simulation with exactly N can yield similar mean populations in an appropriate large system but very different ground-state number fluctuations. In the grand ensemble one isolated orbital has a geometric number distribution with variance $\langle n\rangle(1 + \langle n\rangle)$. A mean ground population of 99 then corresponds to a standard deviation close to the mean. A sealed sample of exactly 100 atoms cannot physically realise this unbounded distribution. Comparing its observed condensate fluctuations to the independent grand-canonical orbital formula without enforcing number conservation would be misleading.

The heat capacity sees the excited cloud, not a hidden latent heat

Below the bulk transition in a uniform ideal gas the fugacity tends to one. The energy expression becomes

$$U = (3/2)Vk_BT\lambda_T^{-3}\zeta(5/2),$$

which grows as $VT^{5/2}$. At fixed box size and fixed total atom number, the excited fraction adjusts with temperature while the ground state supplies or accepts the remaining particles. Differentiating gives

$$C_V = (15/4)Vk_B\lambda_T^{-3}\zeta(5/2).$$

Dividing by the total particle count gives a heat capacity per particle proportional to $(T/T_c)^{3/2}$.

At the critical temperature just below the transition, the coefficient per particle is $15\zeta(5/2)/(4\zeta(3/2)) \simeq 1.93k_B$. At half that temperature it is about $0.681k_B$. The energy and entropy are continuous at the uniform ideal transition; saturation of the excited population is not the same as a first-order phase change with a latent heat between two bulk densities. In the harmonic trap the different density of

states yields a different power and a different shape of the heat-capacity anomaly. Thermodynamic measurements can therefore help distinguish trap geometry when imaging alone is ambiguous.

A calorimetric estimate of temperature often assumes the gas can re-equilibrate after a controlled energy pulse. In a perfectly noninteracting isolated gas, occupation numbers of stationary orbitals do not redistribute by themselves. Weak collisions, a reservoir or an explicit thermalisation protocol must establish equilibrium. The ideal gas can be an excellent equation-of-state model while weak interactions remain necessary for reaching that state on laboratory timescales. If the waiting time after a pulse is shorter than the redistribution time, substituting the equilibrium heat capacity into $\Delta T = \Delta U / C$ may falsely imply anomalous heating.

One practical diagnostic is to vary the waiting time while monitoring both the broad-wing kinetic temperature and central fraction. Thermal equilibrium predicts a definite relation between those two quantities at fixed total number and geometry. A lagging central fraction together with apparently stable wings suggests a slow redistribution process or atom losses. Repeating with matched final number and with independently calibrated detection efficiency is more discriminating than fitting temperature from the central peak alone.

Detect bosonic bunching without assuming a condensate

Occupation can also be probed through intensity or number correlations. For a single thermal bosonic mode with geometric distribution, the normally ordered second-order correlation is

$$g^{(2)}(0) = \langle n(n-1) \rangle / \langle n \rangle^2 = 2.$$

A coherent-state Poisson number distribution gives one. A precisely fixed number n gives $1 - 1/n$. These statements concern the statistics of the selected mode and the preparation. They must not be read as a universal test in which every condensed cloud has the same value under every ensemble or detection geometry.

In a detector pixel collecting M independent identically populated thermal modes, the mean total number is $M\bar{n}$, variance $M\bar{n}(1+\bar{n})$. Substituting into the factorial moment gives

$$g^{(2)}(0) = 1 + 1/M.$$

Ten unresolved thermal modes therefore give 1.10 rather than 2, even though each separately bunches fully. Finite temporal averaging, velocity-bin width and spatial point spread can increase the effective number of collected modes. A near-unity measured correlation does not by itself demonstrate that the gas has developed a macroscopically coherent orbital.

Why use a factorial moment? For independent shot noise the ordinary second moment contains a self-count term $\langle n \rangle$. Subtracting it isolates pair coincidences between distinct particles. A detector that cannot distinguish simultaneous events may preferentially drop exactly those pairs, creating artificial suppression. Two independent detection channels after a calibrated splitter can estimate coincidences without a single-pixel dead-time bias, provided cross talk and timing jitter are characterised.

A useful exercise varies both integration width and total atom number. A correlation excess falling with broader detection windows supports collection of multiple unresolved thermal modes. A central mode whose occupancy remains a fixed fraction as the whole gas grows supplies different information. Neither test

replaces a correctly modelled spatial coherence measurement, yet together they constrain interpretations that can fit one mean momentum image.

How a number solver reveals its own near-critical sensitivity

At fixed temperature, derivative of total ideal boson occupation with respect to chemical potential is $\beta \sum_j n_j(1+n_j)$. A highly occupied lowest orbital can dominate the derivative even if the measured energy is primarily carried by the excited states. Hence a small error in chemical potential can produce a large error in number while barely changing calculated energy. A solver that stops when energy converges may report the wrong condensate fraction.

For a mean ground occupation 10^4, its chemical-potential distance below the lowest level is $k_B T \log(1+10^{-4})$, approximately $10^{-4} k_B T$. A shift in that tiny gap by ten percent changes the occupancy by roughly ten percent while leaving the apparent fugacity close to one to four significant digits. Track chemical potential relative to the exact ground energy and report particle-number residual directly. Conversely, a small absolute error in a reported fugacity near saturation is not by itself evidence for a large uncertainty in the empirically counted atom number.

Finally, compare continuum estimates on several finite boxes at the same density. The leading discrepancy often comes from low discrete modes rather than the high-energy cutoff. Refine the low-energy calculation and the tail independently: enumerate a set of low shells exactly and replace only the dense upper spectrum by an integral. Check for double counting at the matching cutoff. An apparently excellent bulk critical temperature at one size can otherwise result from two cancelling approximations with different size dependence.

Infer a box density from a harmonically confined thermal cloud

Far from degeneracy a harmonically confined gas has classical spatial density $n(\mathbf{r}) = n(0) e^{-V(\mathbf{r})/(k_B T)}$. For an isotropic trap $V = m\omega^2 r^2/2$, each coordinate is Gaussian with variance $s^2 = k_B T/(m\omega^2)$. Normalising its full three-dimensional density gives

$$n(0) = N/[(2\pi)^{3/2} s^3].$$

An absorption image generally measures a column integral, whose central value is $N/(2\pi s^2)$, with units per area. Substituting this column density into a volumetric condensation criterion is dimensionally incorrect.

For $N = 10^5$, mass 87 unified atomic units, ordinary trap frequency 100 Hz and temperature 400 nK, the classical one-axis size is about 9.85 micrometres. The inferred central volume density is 6.65×10^{18} per cubic metre. The phase-space density estimate $n(0)\lambda_T^3$ is about 0.17; the classical profile is therefore a useful leading reference rather than a precision model at all radii. A ten-percent systematic underestimate of width inflates the inferred central density by about 37 percent because width enters cubed.

Near the transition the local-density approximation replaces the constant chemical potential by $\mu - V(\mathbf{r})$ inside the uniform-gas density function. This preserves the global chemical potential while reducing local effective fugacity away from the centre. Integrating the resulting local density over position recovers the harmonic-trap density-of-states counting in the semiclassical regime. A central phase-space density

approaching the uniform saturation value can therefore coincide with a trap-wide critical atom number governed by $\zeta(3)$, not $\zeta(3/2)$. These statements use different spatial integrations and are consistent.

A practical reconstruction fits widths along multiple axes and includes known imaging point spread. For a Gaussian image convolved with a Gaussian response, variances add. A measured width of 10.0 micrometres and known response width 3.0 micrometres imply true width 9.54 micrometres, changing a cubed-density estimate by roughly fifteen percent. A narrow condensate component often has a non-Gaussian profile or a size comparable with the resolution, so the simple Gaussian deconvolution should not be applied indiscriminately to both components.

Compare inferred atom counts over increasing fields of view. Truncating broad thermal wings undercounts the thermal population and makes the apparent condensed fraction too large, particularly near onset. Independent normalisation from the complete cloud and fits to held-out outer regions help separate this effect from an actual change in the fraction occupying the lowest orbital.

There is a useful contrast with a single-occupancy adsorption site. Its occupation variance at mean fraction f is $f(1-f)$, whereas one thermal Bose orbital at mean occupation $\bar{n}$ exhibits $\bar{n}(1+\bar{n})$. At mean 0.50 these are respectively 0.25 and 0.75. The sign difference expresses whether multiple occupancy is excluded or enhanced. Vary the chemical potential in the single-site experiment and compare the response-to-noise identity before evaluating the Bose derivative on paper. Both obey the same grand-ensemble fluctuation formula, yet their number variances differ at the same mean because the allowed microscopic configurations differ. An observed sub-Poisson count can alternatively arise from detector dead time, and multiple unresolved thermal modes reduce bunching, so the ideal comparison is strongest for calibrated single-mode or single-site measurements.

Compare the measured single-site fluctuation law with the Bose orbital variance

Count the excited atoms in a trapped gas

Change atom number and confinement frequency. Compare the discrete level sum with the continuum estimate and relate the difference to the density of states.

Open this experiment online

Try the calculation

Atom number **20000**

Isotropic trap frequency (Hz) **100**

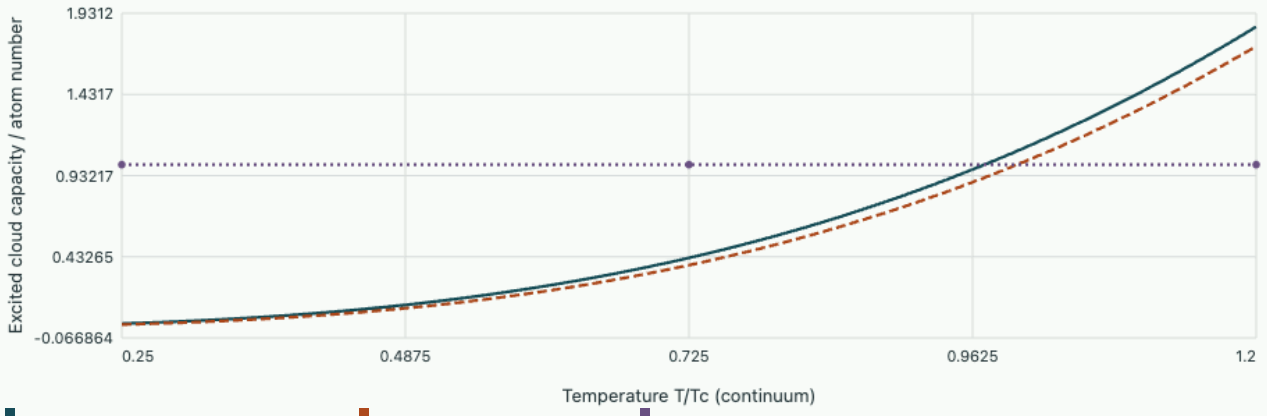

Continuum Tc 122.52 nK. The discrete excited levels reach N at 119.17 nK.

Tutorial · Check what changes during an adiabatic box expansion

Below the transition a dilute ideal condensed box gas expands slowly and reversibly with fixed N. Determine how temperature and condensate fraction vary with volume. Would an isothermal expansion give the same outcome?

▼ **Use the entropy and critical-temperature scalings together**

Below transition the bulk chemical potential tends to zero. Since $U = 3pV/2$, $S = (U + pV - \mu N)/T$ is proportional to $VT^{3/2}$ for the saturated excited gas. At constant entropy $T \propto V^{-2/3}$. The box critical temperature at fixed N has exactly the same volume dependence, so T/T_c and the leading condensate fraction stay constant. In an isothermal expansion T_c falls while T stays fixed, decreasing condensed fraction and eventually exhausting it. The result rests on ideal-gas equilibrium and reversible expansion; a sudden opening into a second container is not that adiabatic path.

Tutorial · Determine whether a trap can be treated as three-dimensional

An anisotropic harmonic trap has ordinary frequencies 40, 40 and 8000 Hz at a measured $T = 100$ nK. Compute the spacings expressed as temperatures and assess using the three-dimensional semiclassical trap formula. Propose a consistent calculation near the dimensional crossover.

▼ **Identify which quantum numbers are actually sampled**

Using $h/k_B = 4.799 \times 10^{-11}$ K s, the level spacings are 1.92, 1.92 and 384 nK. Many transverse levels in the soft directions are populated, whereas a single quantum along the tight axis costs nearly four thermal energies. A three-dimensional integral over all three quantum numbers would overcount accessible low-energy states. Keep the tight-axis sum discrete while approximating the two soft directions semiclassically. Separate the exact global ground orbital from that hybrid excited-state count before solving for chemical potential at fixed total number.

At a few hundred nanokelvin the next tight-axis branches start to contribute appreciably. Observe widths or occupation across a temperature scan instead of announcing a fixed dimensionality based solely on trap shape. The model should recover a two-dimensional low-temperature limit and cross over smoothly toward three-dimensional counting when all spacings become small relative to thermal energy.

Tutorial · When is a ground-state peak extensive?

Three otherwise comparable measurements at fixed estimated areal density and temperature use planar boxes of areas in the ratio one, two and four. They yield approximately the same central-state population of 150 atoms but total numbers 3000, 6000 and 12000. Compare with the ideal planar fugacity estimate above and explain what conclusion follows.

▼ **Examine scaling and the resolution of a single mode**

The fractions are 0.050, 0.025 and 0.0125, falling inversely with area. This is consistent with a highly occupied but nonextensive single level at fixed fugacity near 0.993. A homogeneous three-dimensional bulk condensate at fixed reduced temperature would instead have approximately constant fraction under the appropriate volume scaling. Absolute occupation larger than one is insufficient to distinguish the two regimes.

Confirm that the same experimentally defined central window isolates a single momentum mode as the box size changes. Allowed wavevector spacing shrinks with system size, so a fixed physical momentum aperture can sum an increasing number of levels. Correct for finite resolution, compare the complete momentum distributions and control interactions before treating the simple size dependence as a proof about the thermodynamic limit.

Cambridge Statistical Physics · ideal quantum gases and Bose condensation

Degenerate fermions, heat and magnetic response

Statistical mechanics from microscopic counting to material response

An ordinary metal at room temperature contains electrons with characteristic kinetic energies of several electronvolts, yet its electronic heat capacity is small compared with the classical estimate. Identify which electrons respond to heat, compression and a weak magnetic field. Then compare three independent ways to determine a density of states and investigate what a disagreement would mean.

Count states and fill them with fixed particle number

A spin-half noninteracting fermion orbital of specified spin can be empty or occupied once. Its grand partition factor is $1 + e^{-\beta(\epsilon-\mu)}$, giving mean occupation $f = 1/[e^{\beta(\epsilon-\mu)} + 1]$. At zero temperature the occupations form a step. In a uniform three-dimensional box, fill both spin states at each wavevector up to radius k_F. Counting reciprocal-space cells gives

$$n = \frac{k_F^3}{3\pi^2},$$

$$E_F = \frac{\hbar^2 k_F^2}{2m},$$

$$v_F = \frac{\hbar k_F}{m},$$

$$D(E) = \frac{V}{2\pi^2}\left(\frac{2m}{\hbar^2}\right)^{3/2}\sqrt{E}.$$

Here D counts both spin projections per unit energy for the whole sample. It obeys $D(E_F) = 3N/(2E_F)$. Distinguish this density from a density per volume or per spin, since factors of volume and two matter in susceptibility and specific heat. A lattice band may have a quite different dispersion and density of states; simply substituting the vacuum electron mass is a model to be tested for a given solid.

At $n = 8.5 \times 10^{28}$ electrons per cubic metre and the free electron mass, $E_F = 7.05$ eV, $T_F = 81{,}800$ K and $v_F = 1.57 \times 10^6$ m/s. Room temperature is only 0.00367 of this Fermi temperature. The gas is highly degenerate despite not being near absolute zero in everyday terms. These speeds are random occupied-state group velocities, with equal positive and negative contributions to net current in equilibrium. A metal with no applied field does not transport its entire electron density in one direction at that speed.

Integrating $ED(E)$ to the Fermi level gives

$$U_0 = 3NE_F/5.$$

Volume differentiation at fixed N, using $E_F \propto V^{-2/3}$, gives pressure $p_0 = 2nE_F/5$. Thus a nonzero kinetic pressure remains at zero temperature. It describes the electron subsystem and is balanced by other forces in a neutral solid. It is not the pressure one could read by attaching an ordinary fluid pressure gauge to a solid metal specimen.

Derive the low-temperature expansion without losing particle

conservation

For a smooth function $H(E)$, compare $\int H(E)f(E)\,dE$ with its zero-temperature integral up to μ. Write the difference as contributions above and below μ, substitute $x = |E-\mu|/(k_BT)$ and use $1 - f(\mu - k_BTx) = f(\mu + k_BTx)$. The difference then samples $H(\mu + k_BTx) - H(\mu - k_BTx)$, so only odd derivatives of the smooth function contribute. Integrating the leading term gives

$$\int_0^\infty H(E)f(E)\,dE = \int_0^\mu H(E)\,dE + \frac{\pi^2}{6}(k_BT)^2H'(\mu) + O(T^4).$$

This expansion assumes the Fermi level is many thermal energies above the bottom and that H varies smoothly through the thermal shell. Near a band edge or a very narrow density-of-states peak the assumption can fail even when temperature looks modest relative to some larger energy. Direct numerical quadrature retaining the correct band limits and a particle-number solution for μ is then preferable.

For fixed N, use $H = D$ and demand that the result remain its zero-temperature value. The chemical potential shift is $\delta\mu = -(\pi^2/6)(k_BT)^2D'(E_F)/D(E_F)$. For $D \propto \sqrt{E}$ it becomes

$$\mu \simeq E_F[1 - (\pi^2/12)(T/T_F)^2].$$

Next use $H = ED(E)$ in the energy integral. Its T^2 contribution from the shifted upper limit cancels the part involving E_FD', leaving

$$U - U_0 = \frac{\pi^2}{6}D(E_F)(k_BT)^2,$$

$$C_V = \frac{\pi^2}{3}k_B^2TD(E_F).$$

The typical number of thermally active state slots scales as thermal energy times $D(E_F)$. This produces a heat capacity linear in T, not proportional to the entire electron count times a temperature-independent classical factor. At 300 K in the example the electronic heat capacity is $0.0181Nk_B$. Its volume coefficient $C_V/(VT)$ is about 70.8 J per cubic metre per kelvin squared. A much larger lattice contribution may hide it at room temperature. At low temperature fit C/T against T^2 to separate an electronic intercept from a Debye T^3 lattice term where these asymptotic laws are valid.

Compare the conserved-number chemical potential with the Sommerfeld approximation

Apply a weak field to the spin populations

For a weak uniform magnetic induction B neglect orbital coupling temporarily. Spin moments of magnitude μ_B acquire energy offsets $\mp\mu_BB$ for their two alignments in the simple electron model. At fixed total density and zero temperature, one spin population gains states at its Fermi surface and the other loses an equal number to first order. The density of states per spin is $D(E_F)/2$; the difference in population is thus $D(E_F)\mu_BB$. Magnetisation is this difference times μ_B/V, giving the weak-field SI susceptibility

$$M = \frac{\mu_B^2 D(E_F)}{V} B,$$

$$\chi_P = \frac{\partial M}{\partial H} \simeq \mu_0 \mu_B^2 D(E_F)/V.$$

We have used $B \simeq \mu_0 H$ in a weakly magnetic macroscopic state and ignored demagnetising corrections. The result is approximately temperature-independent for a strongly degenerate smooth band. Independent localised moments would give a Curie susceptibility proportional to inverse temperature in their high-temperature weak-field regime. A small low-temperature upturn can therefore come from a small concentration of local moments even when most carriers remain itinerant.

For the earlier density $\chi_P \simeq 1.22 \times 10^{-5}$. A field of 1.0 T splits spin energies by $2\mu_B B = 0.116$ meV, vastly smaller than 7.05 eV. The fractional spin imbalance is $3\mu_B B/(2E_F) \simeq 1.23 \times 10^{-5}$. Tiny spin imbalance can produce a reproducible measurable magnetisation because the sample contains a macroscopic number of electrons. Linear susceptibility should be tested by field reversal. A hysteretic or saturating impurity background cannot be removed reliably by fitting one positive-field point.

At finite temperature the response replaces sharp $D(E_F)$ by $\int D(E)[-\partial f/\partial E]dE$ at the corresponding chemical potential. The negative occupation derivative integrates to approximately one and has width several thermal energies. This Fermi-window average is another reason the susceptibility and heat capacity emphasise nearby states rather than the entire filled band. Their precise kernels differ, so a rapidly varying band structure or interactions can alter their ratio.

Infer the same states through a density response

At zero temperature density varies with chemical potential as $\partial n/\partial \mu = D(E_F)/V$. Since $dp = n\, d\mu$ at fixed temperature, the electronic inverse compressibility $1/\kappa_T = n^2/(\partial n/\partial \mu) = 2nE_F/3$. In the example this corresponds to about 64 GPa. A compressibility measurement of the full solid includes lattice elasticity and charge-neutrality effects, so it cannot generally be substituted directly for an isolated electron-fluid result.

A confined electronic layer contacted through a gate can instead provide a thermodynamic density response through its quantum capacitance. If an induced charge increment per area is $e\, dn$, part of the applied electrochemical voltage changes chemical potential by $d\mu/e$. The corresponding capacitance per area is $C_Q/A = e^2(\partial n_{2D}/\partial \mu)$. In a simple series combination the geometric dielectric capacitance and quantum capacitance both reduce the measured gate capacitance. Comparing a measured total directly to $e^2 D$ without removing the geometrical part can create an artificial density-of-states suppression.

Three independent measurements of low-temperature C/T, Pauli spin response and a properly interpreted compressibility all probe the density of states in this noninteracting model. Their ratio then removes $D(E_F)$ from some combinations. Correlations can renormalise spin response differently from specific heat. Orbital diamagnetism, closed-shell ions and temperature-dependent impurity moments must also be separated from a measured total susceptibility. A disagreement is an opportunity to distinguish these separate ingredients, not immediate evidence that state counting itself has failed.

A nonlinear zero-temperature spin response beyond Pauli's slope

At low magnetic fields the shift of either spin population is linear and nearly invisible in total number. To see where the susceptibility approximation comes from, retain the exact zero-temperature band populations. For a parabolic three-dimensional band with field-independent mass, absorb the two equal spin degeneracies into a coefficient A such that $n_{\pm} = A[\mu \pm \mu_B B]_{+}^{3/2}$. The positive part enforces an empty spin band if its shifted edge lies above the chemical potential. The total $n_{+} + n_{-}$ must remain fixed in an isolated neutral sample, while $M = \mu_B(n_{+} - n_{-})$.

Expand both powers at small field. Linear terms cancel in the total density, so the chemical potential changes only at quadratic order. At fixed original Fermi energy the solution is

$$\mu = E_F - (\mu_B B)^2/(4E_F) + \cdots.$$

Population imbalance retains the linear term $3\mu_B B/(2E_F)$ as a fraction of total density. This derivation checks why simply leaving μ fixed gives the correct weak-field magnetisation slope yet fails in higher-order density calculations.

Complete spin polarisation in this model occurs when $\mu - \mu_B B$ reaches zero. All particles then occupy the lower spin band, with effective Fermi energy relative to that band's shifted bottom equal to $2^{2/3}E_F$. It follows that $\mu_B B_{\mathrm{pol}} = 2^{-1/3}E_F$. For a 7.05 eV metal this corresponds to roughly 9.7×10^4 tesla, far outside the regime of an ordinary static bench measurement and outside a naive description of unmodified material bands. The number is a scale comparison, not a proposed magnet design. A system with a much lower Fermi energy or an enhanced effective magnetic moment can show nonlinearity at far lower applied fields.

A practical diagnostic compares positive and negative fields after subtracting the holder signal. Pure weak Pauli response is odd and linear in field with little temperature dependence. Localised impurity spins may also be odd but exhibit curvature or saturation when their Zeeman energy competes with temperature. A carrier density changing with field can alter the inferred Pauli slope. Therefore pair a magnetic measurement with a density or transport measurement before attributing every nonlinearity to spin-band depletion.

Extract a two-dimensional density of states from a gate measurement

Take a uniform conducting layer gated by a planar metal through an insulating spacer. An increment in gate voltage partly changes electrostatic voltage across the insulator and partly shifts the electronic chemical potential. In a simple uniform single-sheet model the charge-density increment $dq = e\,dn_{2D}$ obeys $dV_g = dq/c_g + d\mu/e$, where lowercase c_g is geometric capacitance per area. Writing $c_q = e^2\,dn_{2D}/d\mu$ gives

$$1/c_{\mathrm{meas}} = 1/c_g + 1/c_q.$$

This derivation assumes a stable uniform equilibrium state and a calibrated voltage reference with negligible leakage.

Suppose independent dielectric calibration gives $c_g = 5.00\,\mu\mathrm{F\,cm^{-2}}$, while a low-frequency gate measurement yields

$$c_{\mathrm{meas}} = 3.00\,\mu\mathrm{F\,cm^{-2}}.$$

The inferred quantum capacitance is $c_q = 7.50\,\mu\mathrm{F\,cm^{-2}} = 0.0750$ F per square metre. Dividing by elementary charge squared gives 2.92×10^{36} states per joule per square metre, or 4.68×10^{17} states per eV per square metre. Confusing total sample capacitance with capacitance per area would produce a quantity that cannot be compared across different device sizes.

Sensitivity of the inverse step depends strongly on the ratio of the two measured values. If $c_q \gg c_g$, the measured value approaches the geometric one. Small calibration offsets then dominate their reciprocal difference. For example, a measured value only one percent below c_g implies $c_q \simeq 99 c_g$; an absolute error of one percentage point can remove any useful upper bound. In that limit the data support a lower bound on density of states rather than a precise central value.

Vary gate frequency to test whether traps or charge-relaxation delays contribute. A legitimate equilibrium compressibility response should be approached in the low-frequency limit under conditions where the electronic layer remains coupled to its reservoir. Interface states with slower charging can add a dispersive contribution that looks like extra density of states at one chosen frequency. Compare in-phase capacitance and loss, control for parasitic capacitance and measure several layer areas. Area-proportional capacitance and perimeter-dependent fringe capacitance can often be separated with devices of differing aspect ratio.

A differential heat pulse resolves the small electronic component

At low temperature a metal frequently has $C_{\mathrm{sample}} = \gamma T + \beta_D T^3$ over a limited range. A platform and thermometer add their own heat capacities. A measurement might apply a known energy pulse and fit a temperature relaxation through a calibrated weak thermal link. To infer instantaneous heat capacity from $Q/\Delta T$, the pulse must be small enough that C is nearly constant across the excursion and losses over the pulse negligible, or the energy balance must explicitly include them.

For a 10 mmol sample with $\gamma = 0.80$ mJ per mole per kelvin squared and $\beta_D = 0.10$ mJ per mole per kelvin to the fourth, molar heat capacity at 2 K is $1.60 + 0.80 = 2.40$ mJ per mole kelvin. The sample heat capacity is 24.0 microjoules per kelvin. A 1.2 microjoule pulse would raise it by about 0.050 K neglecting heat leakage and addenda. The electronic fraction here is two thirds. At 4 K the lattice contribution exceeds the electronic term, changing which temperatures most strongly constrain the intercept in a C/T fit.

An exact small-range energy calibration uses the integral $Q = \gamma(T_f^2 - T_i^2)/2 + \beta_D(T_f^4 - T_i^4)/4$ times the sample amount. This avoids confusing a finite pulse average with the capacity at the initial temperature. Measure the empty platform through the same range and apply the same thermal model. A mismatch in thermal contact may create a second relaxation time and imitate sample heat capacity. Two pulses of different sizes and a delayed thermometer provide an independent test of the assumed lumped-temperature approximation.

The density of states inferred from γ measures excitations that can equilibrate during the experiment. A superconducting transition, magnetic ordering, a low-lying impurity contribution or a gap around the Fermi level would change the temperature dependence qualitatively. Rather than forcing a linear plus cubic fit over such structure, identify a normal-state range or include a separately justified physical term. A visually straight line on a narrow interval can hide curvature large enough to alter the intercept materially.

How a Fermi window rounds a sharp density-of-states edge

A two-dimensional band with a constant density of states starting abruptly at energy zero provides an exactly soluble compressibility edge. Differentiating its number equation gives $dn/d\mu = D_0/[1 + e^{-\beta\mu}]$ if D_0 is per area. At zero temperature the response jumps from zero to D_0 as the chemical potential crosses the band edge; at finite temperature it rises as a logistic function. The width between ten and ninety percent equals $2k_BT\log 9$, approximately $4.39k_BT$.

At 10 K that intrinsic energy width is about 3.79 meV. At 30 K it becomes 11.36 meV. A gate experiment converts voltage to energy using a lever arm that may depend on the same quantum capacitance being measured. Calibrate that conversion before inferring temperature from edge width. Static disorder, inhomogeneous local band edges or electrical noise broaden the edge even at low temperature. Plot width against temperature and compare the full convolution model rather than interpreting a nonzero intercept as an electronic thermal floor without independent evidence.

This edge example shows exactly where a low-temperature expansion in smooth density derivatives fails. Every derivative of the constant band interior is zero, yet if the chemical potential lies within a few thermal energies of its abrupt edge there is a strong measurable density response crossover. Numerical integration of the full Fermi function or the exact analytic expression automatically retains this boundary. A calculation based solely on the local derivative of a density of states sampled within the band cannot recover it.

Separate Pauli spin response from independent local moments with a concentration test

A common experimental model of weakly magnetic metallic matter combines an approximately constant itinerant-electron susceptibility with a dilute Curie term C/T from paramagnetic impurities. At modest fields the local moment of independent spin-half sites with magnetic moment magnitude μ has average $m = \mu\tanh(\mu B/k_BT)$. Its small-field derivative gives volume susceptibility $\mu_0 n_i \mu^2/(k_BT)$. The conduction-electron spin response instead scales with the density of states at the Fermi level. Two subsystems exposed to the same bath can therefore have markedly different temperature dependences.

Suppose two volume-susceptibility readings in SI units are 2.20×10^{-5} at 10 K and 1.45×10^{-5} at 40 K. A fit to $\chi_0 + C/T$ gives $C = 1.00 \times 10^{-4}$ K and $\chi_0 = 1.20 \times 10^{-5}$. For independent moment magnitude μ_B, the implied impurity number density is $Ck_B/(\mu_0\mu_B^2) \simeq 1.28 \times 10^{25}$ per cubic metre. Relative to a carrier density of 8.5×10^{28} this is only about 150 parts per million. A small population can dominate the temperature-dependent change without representing most of the material.

Two points cannot verify the Curie form, so repeat across temperatures and fields. At sufficiently high B/T the independent moments approach saturation while the ideal degenerate electron subsystem remains approximately linear. At fixed B/T, plotting the temperature-dependent residual should approximately collapse for independent moments. Interacting or clustered impurities can display a shifted Curie–Weiss temperature, blocking or history dependence, and a distribution of moments can broaden the saturation shape. A successful two-point subtraction alone does not establish a unique background.

A related check uses calorimetry. A split two-level impurity contributes a field-dependent Schottky anomaly when its gap is several thermal energies. Changing field shifts this anomaly, whereas the leading electronic linear coefficient of a simple normal metal varies much less at modest fields. The integrated entropy of an isolated two-state population tends toward $n_i k_B \log 2$ when both orientations become equally populated. A

count inferred from susceptibility can therefore be tested by a separately integrated excess heat capacity, once lattice and electronic backgrounds are constrained.

Take care with units and signs when comparing papers or instruments. Some report molar susceptibility, others volume susceptibility, and cgs and SI magnetic conventions differ by factors involving 4π and volume conversion. Divide a volume susceptibility by number density only after identifying whether the number refers to all host atoms, carriers or the inferred localised impurities. Agreement to a few significant figures after an unverified unit conversion has little diagnostic value.

A final test estimates when local spins should start to saturate at a given field. At 5 tesla and 2 K a magnetic moment equal to one Bohr magneton has $\mu_B B/(k_B T) \simeq 1.68$. The independent-spin alignment is $\tanh(1.68) = 0.933$, already close to saturation. In the same field the example metal's relative itinerant spin imbalance remains only around 6×10^{-5}. A common high-field measurement can therefore discriminate two very different microscopic populations even if both looked nearly linear at much smaller fields. Remanent magnetisation on returning to zero would be inconsistent with both ideal noninteracting equilibrium models. It would motivate a check for ferromagnetic contamination, anisotropic barriers or hysteresis in the measurement apparatus before refining a fit to either simple susceptibility law.

Decide which electronic states contribute to heat capacity

Change temperature relative to Fermi energy and compare the low-temperature approximation with the statistical occupation.

Open this experiment online

Try the calculation

Reduced temperature T/TF

0.1

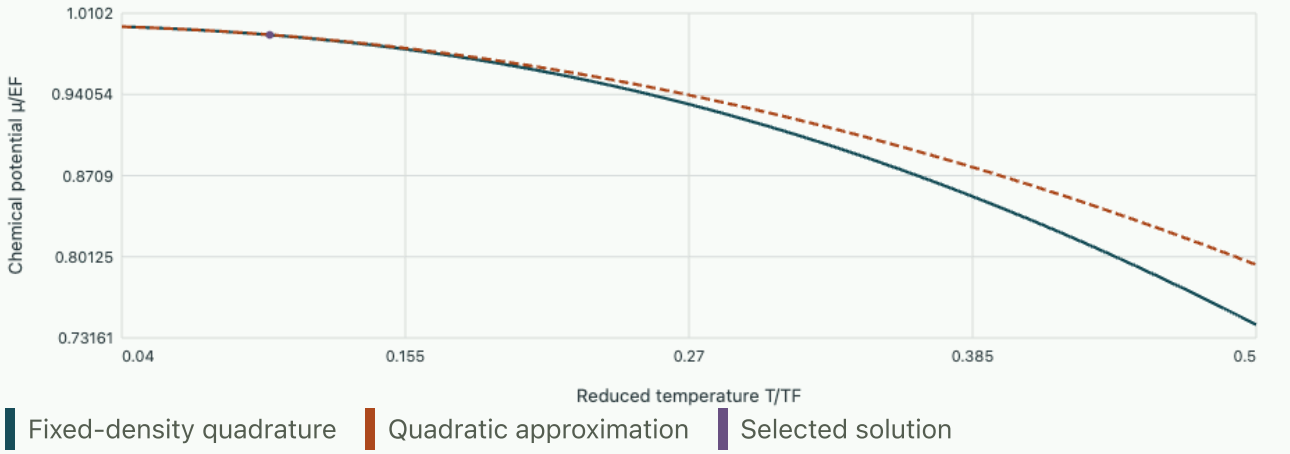

Fixed-number μ/EF = 0.9916, quadratic prediction 0.9918. Leading heat capacity $C_v/(Nk_\beta) = 0.4935$.

Tutorial · Recover a density of states from a calorimeter

A sample exhibits a low-temperature molar heat capacity $C/T = 0.80\,\mathrm{mJ\,mol^{-1}K^{-2}} + bT^2$. Use the intercept to estimate the total density of electronic states per mole per electronvolt. Compare with the prediction for one conduction electron per atom and $E_F = 7.0$ eV.

▼ **Extract the electronic coefficient and convert energy units**

In joules the total molar density is $D = 3(0.00080)/(\pi^2 k_B^2) = 1.28 \times 10^{42}$ states per joule per mole. Multiplying by the number of joules per electronvolt gives 2.05×10^{23} states per eV per mole. For one electron per atom the ideal value is $3N_A/(2E_F) = 1.29 \times 10^{23}$ states per eV per mole. The fitted effective density is therefore about 1.59 times larger. A molar heat capacity must be paired with a molar density of states; using a sample volume here without a molar volume would create a dimensional mismatch.

To decide whether the enhancement reflects a nonfree band, interactions or a wrong temperature window, examine the linearity of C/T versus T^2 over several restricted fit ranges. A low-lying impurity or additional vibrational mode can produce curvature that changes the intercept. Compare with independent band information and spin response after relevant backgrounds are subtracted.

Tutorial · Determine a weak-field magnetisation and its background

At 1.0 T the measured volume magnetisation is 7.0 A/m, independent of temperature over a low-temperature range. The sample's predicted Pauli susceptibility is 1.22×10^{-5}. Estimate the residual susceptibility if the magnetisation is linear and explain why its sign alone does not identify an orbital model.

▼ **Compare quantities in the same magnetic convention**

One tesla corresponds to $H \simeq 7.96 \times 10^5$ A/m in this weakly magnetic setting. The predicted Pauli magnetisation is 9.71 A/m. The residual is minus 2.71 A/m or $\chi_{\mathrm{residual}} \simeq -3.4 \times 10^{-6}$. Orbital diamagnetism of an ideal bulk free-electron gas supplies a negative contribution, but core diamagnetism, band orbital effects and sample-holder background can also contribute. Measure an empty holder and both field signs, consider specimen shape and calibrate the sample volume. A nearly temperature-independent negative residual does not uniquely identify any one mechanism.

Tutorial · Why a constant two-dimensional density of states behaves differently

For an ideal two-dimensional parabolic band of constant total density D_0 above energy zero, derive the finite-temperature chemical potential at fixed $N = D_0 E_F$. Compare its leading low-temperature change with the power-law shift in three dimensions.

▼ **Evaluate the number integral before expanding**

The integral gives

$$N = D_0 k_B T \log[1 + e^{\mu/(k_B T)}].$$

Solving yields

$$\mu = k_B T \log[e^{E_F/(k_B T)} - 1] = E_F + k_B T \log[1 - e^{-E_F/(k_B T)}].$$

Its deviation from E_F at low temperature is exponentially small, consistent with $D'(E_F) = 0$. A program that always imposes the three-dimensional negative quadratic shift would violate particle conservation in this band. The low-temperature heat capacity remains linear because D_0 itself is nonzero.

At high temperature the same exact expression approaches the classical logarithmic chemical potential. Evaluate the low-temperature numerical form as $E_F + k_B T \log(1 - e^{-E_F/k_B T})$ to avoid overflowing the large exponential. Here numerical stability and correct dimensional density of states are independent requirements; solving one does not fix the other.

Cambridge Statistical Physics · Fermi gases, the Fermi surface and low-temperature response

From a field to an observable spectrum

Separate orbital and spin responses, count magnetic states and compare a quantum prediction with the corresponding orbit. Revisit the assumptions behind a nondissipative screening current.

1. Cyclotron motion and quantised magnetic orbits
2. Superconductivity, magnetic screening and flux quantisation
3. Counting carriers and weighing a cyclotron orbit
4. Quantising the orbit and finding the zero Landau level
5. Identify mobile charge with a transverse field
6. Distinguish magnetic screening from electrical diffusion
7. How a magnetic field rotates the polarisation of transmitted light
8. Relativistic particle dynamics and the field tensor
9. Master equations, detailed balance and probability currents
10. Transport coefficients from equilibrium fluctuations

Cyclotron motion and quantised magnetic orbits

Condensed matter and band theory

Relate a magnetic-field measurement to band curvature, level spacing and a cross-section of the Fermi surface.

FROM FIRST PRINCIPLES

Find the orbit and energy spacing imposed by a magnetic field

Consider an electron in an isotropic parabolic band with effective mass m^*, no electric field and a uniform magnetic induction $B\hat{\mathbf{z}}$. The Lorentz force is perpendicular to velocity, so it changes the direction of the transverse velocity without changing its magnitude. Taking absolute values in the radial force balance gives

$$\frac{m^* v_\perp^2}{r_c} = e v_\perp B,$$

$$r_c = \frac{m^* v_\perp}{eB},$$

$$\omega_c = \frac{eB}{m^*}.$$

To obtain energy levels one also needs quantum mechanics. Choose $\mathbf{A} = (0, Bx, 0)$ and electron charge $-e$ with $e > 0$. The Hamiltonian $(\mathbf{p} + e\mathbf{A})^2/(2m^*)$ is independent of y, z, so their canonical momenta can be treated as fixed in a plane-wave state. Complete the square:

$$H = \frac{p_x^2}{2m^*} + \frac{m^* \omega_c^2}{2}\left(x + \frac{p_y}{eB}\right)^2 + \frac{p_z^2}{2m^*}.$$

This is a harmonic oscillator about a centre set by p_y. Its transverse levels have energies $\hbar\omega_c(n + 1/2)$; longitudinal translation remains free. Spin splitting, anisotropic band curvature and collision broadening require further information. The same frequency occurs in the classical orbit and the quantum level spacing for distinct reasons made explicit by these two constructions.

DERIVATION AND APPLICATIONS

Determine which mass a cyclotron resonance measures

An orbit on an anisotropic band need not be circular. Consider a weak uniform magnetic induction B along z, and motion in a single band with elliptical constant-energy contours in the xy plane. Neglect scattering over an orbit, electric fields, and spin splitting. Write the local dispersion as $E = \hbar^2 k_x^2/(2m_x) + \hbar^2 k_y^2/(2m_y)$, measured from its minimum. For charge $-e$, the semiclassical force law and the velocity components give

$$\dot{k}_x = -\frac{eB}{m_y} k_y,$$

$$\dot{k}_y = \frac{eB}{m_x} k_x.$$

Taking a second derivative cancels the other component. Both coordinates oscillate at the same fundamental angular frequency even though their amplitudes differ,

$$\ddot{k}_x = -\frac{e^2 B^2}{m_x m_y} k_x,$$

$$\omega_c = \frac{eB}{m_c},$$

$$m_c = \sqrt{m_x m_y}.$$

A geometric measurement leads to the same mass. The two semiaxes of the contour at energy E are $\sqrt{2 m_i E}/\hbar$. Its reciprocal-space area $A(E)$ and energy derivative are

$$A(E) = \frac{2\pi E}{\hbar^2} \sqrt{m_x m_y},$$

$$\frac{\hbar^2}{2\pi} \frac{dA}{dE} = m_c.$$

The derivative explains how a measurement of orbit area over a small energy range tests the resonance result. Two bands with the same area at one energy can have different derivatives there. For a circle, both principal masses and the cyclotron mass coincide. A single resonance frequency for this orientation measures their geometric mean and leaves their ratio undetermined.

WORKED EXAMPLE

Predict the resonance of an elliptical pocket

Change the inputs and investigate this derivation.

A layer has $m_x = 0.100 m_e$, $m_y = 0.900 m_e$ and is placed in $B = 1.50\,\mathrm{T}$. The cyclotron mass is $0.300 m_e$, so the ordinary frequency is

$$f_c = eB/(2\pi m_c) = 140\,\mathrm{GHz}.$$

The period is $7.14\,\mathrm{ps}$. The long reciprocal-space semiaxis lies along y and is three times the short one, even though both components complete their cycle in exactly the same time.

SECOND APPLICATION

Use two nearby contour areas as a mass measurement

Change the inputs and investigate this derivation.

In a different sample the orbit area rises by $0.200\,\mathrm{nm}^{-2}$ when energy rises by $20.0\,\mathrm{meV}$. Treat the measured slope as locally constant. In SI units $dA/dE = 2.00 \times 10^{17}/(3.204 \times 10^{-21})\,\mathrm{m}^{-2}\mathrm{J}^{-1}$. The area formula gives

$$m_c = 1.10 \times 10^{-31}\,\mathrm{kg} = 0.121 m_e.$$

An independent determination of $m_x = 0.0600 m_e$ would then imply $m_y = m_c^2/m_x = 0.245 m_e$ at the displayed precision.

Try it · Cyclotron mass

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Principal mass mx/me **0.1**

Principal mass my/me **0.9**

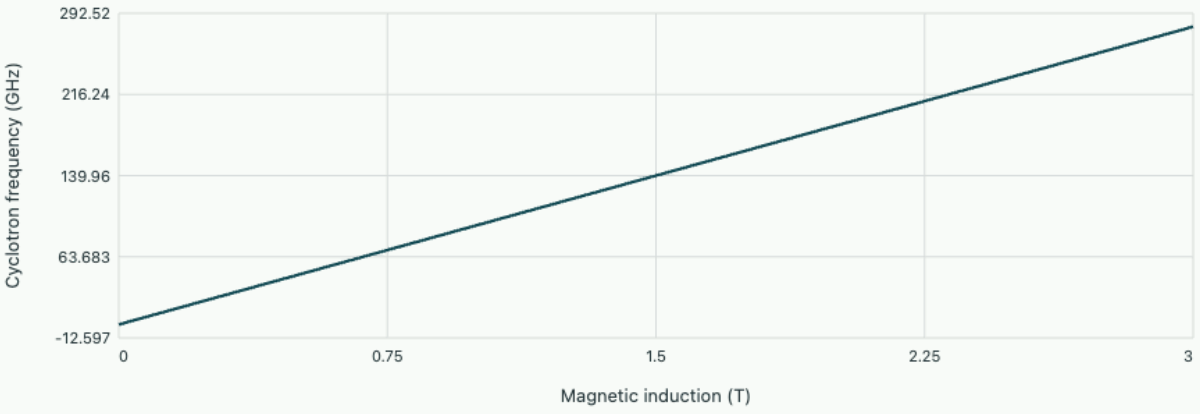

Elliptical parabolic pocket

Cyclotron mass / me = √(mx my) = 0.3. This geometric mean differs from the longitudinal effective mass.

Follow a wave packet in a band

A narrow wave packet in a well-isolated band moves with $\mathbf{v} = \hbar^{-1}\nabla_{\mathbf{k}} E$. In a static uniform magnetic field the semiclassical force law is

$$\hbar\dot{\mathbf{k}} = -e\,\mathbf{v} \times \mathbf{B}.$$

Thus $\nabla E \cdot \dot{\mathbf{k}} = 0$: the energy is constant because a magnetic field does no work. The wavevector component along $\mathbf{B}$ is also constant. The orbit in reciprocal space lies along the intersection of a constant-energy surface and a plane perpendicular to the field. Tilting the field samples different sections of a nonspherical Fermi surface.

For an isotropic parabolic band, motion transverse to the field is circular at $\omega_c = eB/m^*$. The real-space orbit radius is $v_\perp/\omega_c$. One can measure an effective mass from the resonant angular frequency without separately knowing the elastic scattering time. Whether the motion survives long enough to be resolved depends on $\omega_c\tau$. A very short scattering time smears the resonance and suppresses oscillatory magnetotransport.

Derive the quantised spacing

For a two-dimensional electron choose $A_y = Bx$ and $A_x = 0$. Since the Hamiltonian for charge $-e$ is $H = [p_x^2 + (p_y + eBx)^2]/(2m^*)$, a state of definite $p_y = \hbar k_y$ leaves a one-dimensional harmonic oscillator in x, centred at $-\hbar k_y/(eB)$. Its levels are $E_j = (j + 1/2)\hbar\omega_c$, neglecting spin. At fixed level, changing k_y moves the orbit centre without changing its energy. The centres allowed inside a rectangle of area A count to eBA/h per spin, up to edge corrections. In three dimensions motion along the field adds a dispersing kinetic-energy term.

Take $B = 5$ T and $m^* = 0.10 m_e$. The orbital spacing is about 5.79 meV, equivalent to $k_B T$ at 67 K. At 4 K the thermal scale is about 0.345 meV, smaller than the spacing. An observation still requires sufficiently weak disorder broadening and appropriate carrier filling. As the field changes, both the level spacing and the number of states carried by each level change.

Extract a reciprocal-space area from a measurement

For a closed semiclassical orbit of area $A_k(E)$, quantisation gives

$$A_k = 2\pi eB(j + \gamma)/\hbar,$$

where the phase offset γ depends on the orbit. Successive levels crossing the Fermi energy give oscillations periodic in reciprocal field with frequency $F = \hbar A_k(E_F)/(2\pi e)$, expressed in tesla. For $F = 100$ T, $A_k \simeq 9.55 \times 10^{17}$ m⁻². A circular section of that area has radius 5.51×10^8 m⁻¹. In a three-dimensional metal dominant oscillations come from extremal sections. Inferring a full Fermi volume from one frequency requires additional assumptions about shape and symmetry; measurements at several orientations can test them.

Count orbit centres and Hall filling

Landau level spacing determines the energy cost of changing orbital index. A second question is how many independent states fit into each level. Take a rectangular two-dimensional layer of area $L_x L_y$, magnetic field B normal to it and Landau gauge $A_y = Bx$. Translation symmetry along y gives wavevectors $k_y = 2\pi j/L_y$. In the transverse Hamiltonian the term $(\hbar k_y + eBx)^2/(2m^*)$ describes a harmonic oscillator with centre $x_0 = -\hbar k_y/(eB)$.

To place an orbit centre inside the sample, k_y must range over an interval of length $eBL_x/\hbar$. Dividing by the allowed spacing $2\pi/L_y$ gives $N_\phi = eBL_xL_y/h$ bulk states per spin-resolved orbital level, neglecting edge corrections. Each magnetic flux quantum h/e adds one state per level. For sheet density n_s define filling factor $\nu = n_s h/(eB)$. This counts occupied spin-resolved levels when it is an integer; if a residual spin degeneracy is unresolved, levels arrive in pairs.

In an ideal low-temperature integer quantum Hall state, with the chemical potential in a mobility gap, longitudinal dissipation can vanish and Hall conductance magnitude is $\nu e^2/h$. The classical expression $n_s e/B$ has the same value at exact integer filling, but alone cannot explain persistence of plateaus over a finite interval of field. Localization and edge transport are essential for that observation.

Locate an integer filling in a measurement

A two-dimensional electron system has $n_s = 3.00 \times 10^{15}$ m⁻² and effective mass $0.067 m_e$. The field for total filling $\nu = 2$ is

$$B = n_s h/(2e) = 6.20\,\mathrm{T}.$$

The number of states per area per spin-resolved level is $eB/h = 1.50 \times 10^{15}$ m⁻², half the measured sheet density. The orbital spacing at this field is

$$\hbar eB/m^* = 10.7\,\mathrm{meV}.$$

At 4.2 K the thermal energy is 0.362 meV, so thermal smearing is smaller than the orbital gap; disorder and spin structure still matter.

The ideal Hall resistance magnitude on the $\nu = 2$ plateau is $h/(2e^2) = 12.9$ kΩ. A weak current of 0.500 μA would produce a transverse voltage magnitude of 6.45 mV. Increasing the field at fixed density decreases filling and can move the Fermi level into another broadened level. A simultaneous longitudinal voltage measurement is indispensable: a Hall voltage of the right numerical size at a single field can also occur in a classical conductor. A broad flat Hall trace together with a longitudinal-resistance minimum gives stronger evidence for quantized transport.

Changing the sample area at constant sheet density and magnetic field increases the number of occupied particles and the number of orbit centres in the same proportion. Filling factor and the ideal Hall resistance therefore remain unchanged, even though total particle count can be very different. In contrast, a gate that varies sheet density at fixed field moves the layer through successive fillings. The expected density spacing per spin-resolved level is eB/h. A gate-calibrated scan can thus test Landau degeneracy independently of the cyclotron energy spacing inferred from optical or thermal observations.

When can a Landau-level spacing clear thermal broadening?

For a two-dimensional band with isotropic mass $0.12 m_e$ and negligible spin splitting, calculate $\hbar eB/m^*$. Compare the result with the thermal energy for two temperatures. This gives a scale for designing a sweep, rather than a sharp visibility criterion.

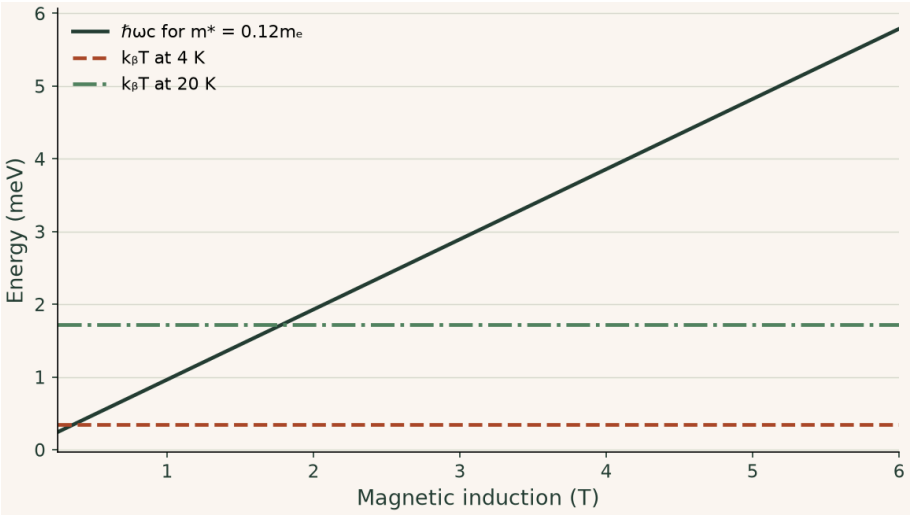

At 20 K the two energy scales match near 1.79 T. Lowering temperature to 4 K moves the equality to about 0.36 T. Level broadening from scattering and the measurement response can demand a stronger field than this scale alone suggests.

Try it · Landau scale

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Temperature (K) **20**

Effective mass / electron mass **0.067**

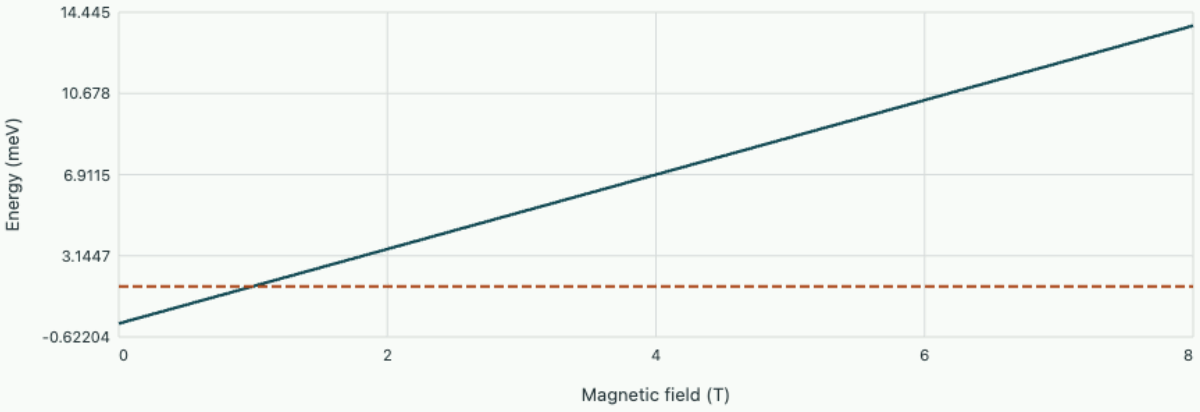

Cyclotron level spacing Thermal energy

Thermal and cyclotron scales match at 0.99745 T. Disorder broadening is a separate experimental scale.

▶ **Read values from the plot**

Superconductivity, magnetic screening and flux quantisation

Condensed matter and band theory

Separate persistent current from magnetic flux expulsion and derive two measurable length and flux scales of a superconducting state.

FROM FIRST PRINCIPLES

Connect a material current law to magnetic screening

A vanishing electrical resistance by itself allows an already existing current to persist. To predict the equilibrium field profile in a thick superconductor one needs additional physical input. In a uniform isotropic superconductor with no vortices the second London relation connects the curl of the supercurrent to the induction. Here n_s counts carriers of mass m and signed charge q. Combine this constitutive law with static Ampère's law:

$$\nabla \times \mathbf{j}_s = -\frac{n_s q^2}{m}\mathbf{B},$$
$$\nabla \times \mathbf{B} = \mu_0 \mathbf{j}_s.$$

Take a second curl of the field law and use $\nabla \cdot \mathbf{B} = 0$. The sign from the vector identity cancels that in the London relation:

$$\nabla^2 \mathbf{B} = \frac{\mathbf{B}}{\lambda_L^2},$$
$$\lambda_L = \sqrt{\frac{m}{\mu_0 n_s q^2}}.$$

Fill $x > 0$ with the material and impose a parallel surface field $B(0) = B_0$. With no field growth deep in the bulk, only one exponential survives:

$$B(x) = B_0 e^{-x/\lambda_L}.$$

Thus the single measured length constrains a ratio $m/(n_s q^2)$. This construction assumes linear local response and a field low enough for the vortex-free state and its prescribed surface condition to apply.

Background · London length

$$B(x) = B(0) e^{-x/\lambda_L}$$

This is a field profile in a thick homogeneous superconductor in a linear screening regime. A thin slab requires both boundary conditions.

DERIVATION AND APPLICATIONS

Determine how deeply a parallel field penetrates a finite slab

The exponential London solution assumes a sample extending far beyond a single exposed surface. A film has two surfaces whose screening currents can overlap. Consider an infinite slab of total thickness $2d$, occupying $-d < x < d$. Apply a static induction B_a parallel to its faces, along z, and assume a uniform London penetration length λ. Work in the vortex-free Meissner state with fields small enough for linear response. Translational invariance parallel to the slab gives the boundary problem

$$\frac{d^2 B_z}{dx^2} = \frac{B_z}{\lambda^2},$$
$$B_z(-d) = B_z(d) = B_a.$$

A general solution has a hyperbolic cosine and a hyperbolic sine. Reflection about the midplane makes the field even, which removes the sine term. The surface field fixes the remaining coefficient. Ampère's law then determines the screening current, whose sign reverses at the midplane,

$$B_z(x) = B_a \frac{\cosh(x/\lambda)}{\cosh(d/\lambda)},$$
$$J_y(x) = -\frac{B_a}{\mu_0 \lambda} \frac{\sinh(x/\lambda)}{\cosh(d/\lambda)}.$$

A probe at the centre and a probe integrating flux over the cross-section read different quantities. The latter reports the thickness average obtained by integrating the field profile,

$$\frac{\overline{B}}{B_a} = \frac{1}{2dB_a} \int_{-d}^{d} B_z(x)\, dx = \frac{\lambda}{d} \tanh(d/\lambda).$$

For a thick slab the mean fraction tends to λ/d, consistent with two active surface layers. For a thin slab it is $1 - (d/\lambda)^2/3$ to leading order. Thin samples can therefore exhibit weak parallel-field exclusion even when the material is fully superconducting throughout their volume.

WORKED EXAMPLE

Compare central and average field in a thick slab

Change the inputs and investigate this derivation.

A sample of thickness $600\,\mathrm{nm}$ with $\lambda = 100\,\mathrm{nm}$ has $d/\lambda = 3$. In $B_a = 2.00\,\mathrm{mT}$ its central field is

$$2.00/\cosh 3 = 0.199\,\mathrm{mT}.$$

The mean induction is

$$2.00\ \tanh(3)/3 = 0.663\,\mathrm{mT}.$$

The higher mean follows directly from the substantial contribution of material close to both exposed faces.

SECOND APPLICATION

Extract a penetration length using a local probe

Change the inputs and investigate this derivation.

A probe halfway through a $120\,\mathrm{nm}$ film measures $0.800B_a$. The central-field solution gives $d/\lambda = \operatorname{arcosh}(1.25) = 0.693$, so $\lambda = 86.6\,\mathrm{nm}$. The same film should transmit a mean fraction $\tanh(0.693)/0.693 = 0.866$. Comparing this value with an independent flux integral tests the assumptions of uniform thickness and a single penetration length.

Try it · Solid optics methods london slab

Vary the physical parameters in the derived relation. Use the units and assumptions in the readout to connect the curves to the two worked examples. A parameter study can test the forward prediction; the same relation can then be inverted to infer a parameter from a measurement.

Open this live calculation

Try the calculation

Film thickness (nm) 400

London length (nm) 100

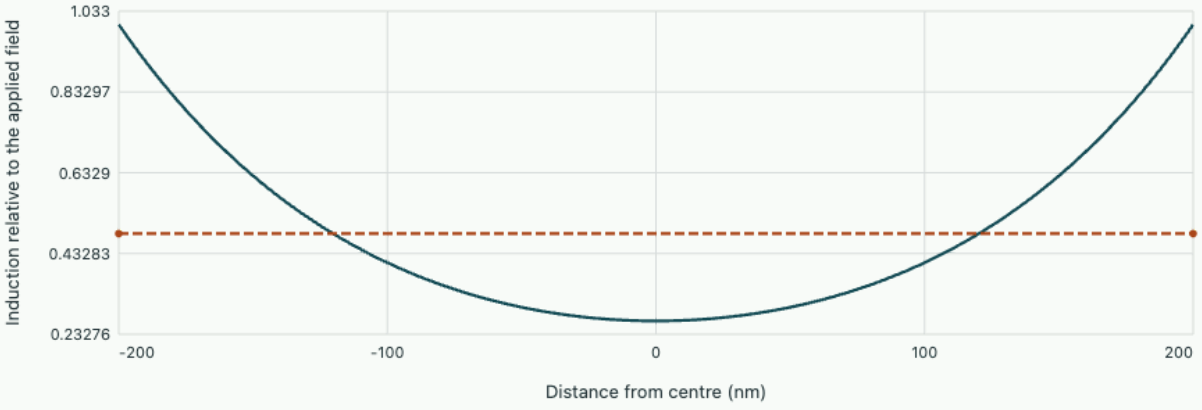

At the centre B/Ba = 0.2658, whereas the cross-section mean is 0.4820. These apply to a parallel field in a vortex-free slab.

Compare two magnetic histories

A normal conductor dissipates energy from a sustained current. A superconductor can support persistent current and, in an appropriate low-field equilibrium state, expels magnetic flux from its bulk. Imagine cooling a

specimen through its transition in a field and subsequently measuring its interior. Perfect conductivity alone would imply that a previously existing flux could remain frozen. Equilibrium flux expulsion supplies additional evidence about the state; sample shape, vortices and pinning affect what a real experiment detects. A resistance measurement and a magnetisation measurement therefore probe different properties.

Derive the London screening length

Let mobile condensate carriers have density n_s, mass m_s and charge q_s, using a consistent convention for particles or pairs. In a bulk region without vortices the London relation is

$$\nabla \times \mathbf{j}_s = -(n_s q_s^2/m_s)\mathbf{B}.$$

Combine this with magnetostatic Ampère's law $\nabla \times \mathbf{B} = \mu_0 \mathbf{j}_s$ and $\nabla \cdot \mathbf{B} = 0$. Taking a curl gives

$$\nabla^2 \mathbf{B} = \mathbf{B}/\lambda_L^2,$$

$$\lambda_L = \sqrt{m_s/(\mu_0 n_s q_s^2)}.$$

For a half-space $x > 0$ with surface field parallel to the surface and deep interior field zero,

$$B(x) = B(0)e^{-x/\lambda_L}.$$

At depth $3\lambda_L$, 4.98% of the surface amplitude remains. Taking $n_s = 2.0 \times 10^{28}$ m⁻³ electrons and electron mass in the simple carrier convention gives $\lambda_L \approx 37.6$ nm. Real effective masses, temperature-dependent superfluid density and nonlocality can change this value. Mixing pair density with single-electron charge would introduce a systematic factor error.

Quantise a circulating phase

A coherent condensate has a phase ϕ whose change around a closed path must be $2\pi n$ for integer n if the complex order parameter is single-valued and nonzero along the path. The gauge-invariant momentum is $m_s \mathbf{v}_s = \hbar \nabla \phi - q_s \mathbf{A}$. Integrating around a loop relates circulation and flux. If a path lies sufficiently deep in a thick ring that its current density is negligible, the velocity term vanishes and enclosed flux comes in units $\Phi_0 = h/|q_s|$. For paired carriers $|q_s| = 2e$ the unit is 2.068×10^{-15} Wb. In thin rings where the kinetic term matters, it is the full fluxoid combination that is quantised.

Use length scales to interpret the experiment

The penetration depth characterises field screening. A separate coherence length ξ measures the distance over which the condensate amplitude recovers after a perturbation. In bulk Ginzburg–Landau theory near the transition, the ratio λ/ξ distinguishes type I and type II at $1/\sqrt{2}$. In a type-II material between the lower and upper critical fields, flux can penetrate in quantised vortices with normal-like cores surrounded by screening currents. A specimen in this mixed state can retain superconducting regions while admitting magnetic flux. Stating field, temperature and cooling history is essential when interpreting the observed signal.

Locate the upper critical field from the linearized order parameter

The London equation assumes an established superconducting density. Near the loss of superconductivity at a high applied field, the order parameter itself becomes small and its spatial variation matters. For a bulk type-II material near its transition use Ginzburg–Landau free-energy density terms $\alpha|\psi|^2 + \beta|\psi|^4/2 + |(-i\hbar\nabla - q_s\mathbf{A})\psi|^2/(2m_s)$ with $\alpha < 0$, $\beta > 0$ and pair charge magnitude $|q_s| = 2e$. Very close to the upper critical field the density is small, field screening is weak and a uniform induction B is a suitable leading approximation.

Minimizing the free energy with respect to ψ^* and discarding the cubic term gives a charged-particle magnetic eigenproblem. The lowest transverse kinetic eigenvalue is the lowest Landau energy $\hbar|q_s|B/(2m_s)$, achieved with no longitudinal variation. A nontrivial small solution is possible at the onset when this cancels α. Defining $\xi^2 = \hbar^2/(2m_s|\alpha|)$ and $\Phi_0 = h/|q_s|$ gives

$$B_{c2} = \frac{\Phi_0}{2\pi\xi^2}.$$

Thus a coherence length inferred from upper critical induction measures the scale of condensate recovery. Penetration depth instead controls field screening in an already developed condensate. The GL derivation is controlled near the critical temperature and for orbital suppression. Strong paramagnetic pair breaking or pronounced anisotropy can complicate the interpretation of a measured high-field transition.

Connect critical induction to vortex spacing

At a specified temperature close to the transition a bulk specimen loses its superconducting response at $B_{c2} = 1.20$ T. The inferred coherence length is

$$\sqrt{2.068 \times 10^{-15}/(2\pi \times 1.20)} = 16.6\,\mathrm{nm}.$$

An independent penetration-depth measurement at the same temperature yields $\lambda = 120$ nm, giving $\lambda/\xi = 7.24$, above the type-II boundary. Using lengths measured at unrelated temperatures would weaken that check.

At mean induction 0.200 T within the mixed state each singly quantized vortex carries one flux quantum. Its share of sample area is

$$\Phi_0/B = 1.03 \times 10^{-14}\,\mathrm{m}^2.$$

For a triangular lattice that area also equals $\sqrt{3}\,a_v^2/2$, giving vortex spacing $a_v = 109$ nm. Increasing induction fourfold halves the spacing if the array remains in the mixed state and retains its geometry. Near B_{c2} cores overlap and an isolated-vortex picture becomes unreliable even though the linearized onset scale remains useful. Comparing a measured spacing with Φ_0/B tests flux counting independently of the inferred coherence length and of the magnetic field profile around any single vortex.

A useful unit check compares the superconducting flux quantum with the electron flux scale from ordinary Landau-level counting. A condensate pair has twice the charge magnitude of one electron, giving half the flux quantum for a phase winding. At 0.200 T a square micrometre should contain about $BA/\Phi_0 = 96.7$ singly

quantized vortices away from edges. Counting vortices in an image therefore gives a bulk average induction estimate independent of a local peak-field calibration. A surface barrier, pinning or an insufficiently large field of view can account for a local deviation from the average.

Imaging successively larger regions tests how well such local counts represent the bulk.

Counting carriers and weighing a cyclotron orbit

Condensed matter and band theory

A conical band has no constant quadratic band mass at its touching point. A carrier on a finite-energy orbit still has a measurable cyclotron mass. Counting states fixes the Fermi circle and determines how this mass changes when a gate adds charge.

Count states in a two-dimensional disc

Put a two-dimensional sheet of area $A = L_x L_y$ under periodic boundary conditions. Neighbouring allowed wavevectors differ by $2\pi/L_x$ and $2\pi/L_y$, so each state occupies reciprocal-space area $(2\pi)^2/A$. At zero temperature fill all conduction states up to radius k_F, measured from a Dirac point. There are two spin states and two inequivalent valleys in ideal monolayer graphene, giving total degeneracy $g = 4$. Density relative to charge neutrality is therefore

$$n = g\frac{\pi k_F^2}{(2\pi)^2} = \frac{k_F^2}{\pi},$$

$$E_F = \hbar v_F \sqrt{\pi n}.$$

For electron and hole densities use $|n|$ under the square root and give E_F the sign of n relative to the touching point. Differentiating the state count with respect to energy gives the density of states per area $D(E) = g|E|/[2\pi(\hbar v_F)^2]$. It tends to zero at an ideal neutral cone. Thermal broadening, disorder and a gap all change what a finite-resolution probe observes near neutrality.

Background · Cyclotron mass

$$m_c = (\hbar^2/2\pi)\,\partial A/\partial E$$

The area is that of a closed constant-energy orbit in reciprocal space normal to the magnetic field. This is a measured orbital mass rather than a universal free-particle parameter.

Let a perpendicular field bend the orbit

For a wavepacket in a band with weak magnetic field,

$$\hbar\dot{\boldsymbol{k}} = -e\,\boldsymbol{v} \times \boldsymbol{B},$$

where e denotes the positive elementary charge. An isotropic band has velocity normal to a circular constant-energy contour. The orbit is traversed with speed $|\dot{\boldsymbol{k}}| = ev_F B/\hbar$. Dividing its circumference $2\pi k_F$ by that speed gives the period. Express its angular frequency as eB/m_c to identify

$$m_c = \frac{\hbar k_F}{v_F} = \frac{|E_F|}{v_F^2}.$$

The mass measures the time needed to travel around a contour. It depends on the radius even when the magnitude of the group velocity stays fixed. For a general band area $\mathcal{A}(E)$, neighbouring contours separated perpendicular to themselves by $dk_\perp$ differ in energy by $\hbar v\, dk_\perp$. Thus $|d\mathcal{A}/dE| = \oint dl/(\hbar v)$, which gives $m_c = \hbar^2 |d\mathcal{A}/dE|/(2\pi)$ by integrating the same equation of motion.

Convert a gate density into an energy and a mass

At $n = 10^{12}\,\mathrm{cm}^{-2} = 10^{16}\,\mathrm{m}^{-2}$ and $v_F = 10^6\,\mathrm{m/s}$, one obtains $k_F = 1.772 \times 10^8\,\mathrm{m}^{-1}$, $E_F = 0.117\,\mathrm{eV}$ and $m_c = 1.87 \times 10^{-32}\,\mathrm{kg} = 0.0205 m_e$. Quadrupling the density doubles the Fermi energy and cyclotron mass. By contrast a single isotropic parabolic band $E = \hbar^2 k^2/(2m)$ has the same counting geometry but constant $m_c = m$. Transport mobility or disorder scattering time does not enter the state count, although scattering can obscure the cyclotron motion.

Change velocity and signed density. Raising the velocity at fixed density increases $|E_F|$ while reducing the mass because the orbit radius stays fixed and the wavepacket moves faster. Exactly at charge neutrality, a single classical Fermi orbit no longer describes the carriers. Finite-field quantum levels replace the continuum picture in the next chapter.

Open this experiment online

Try the calculation

Signed density (10^{12} cm^{-2}) 1

Cone velocity (10^6 m/s) 1

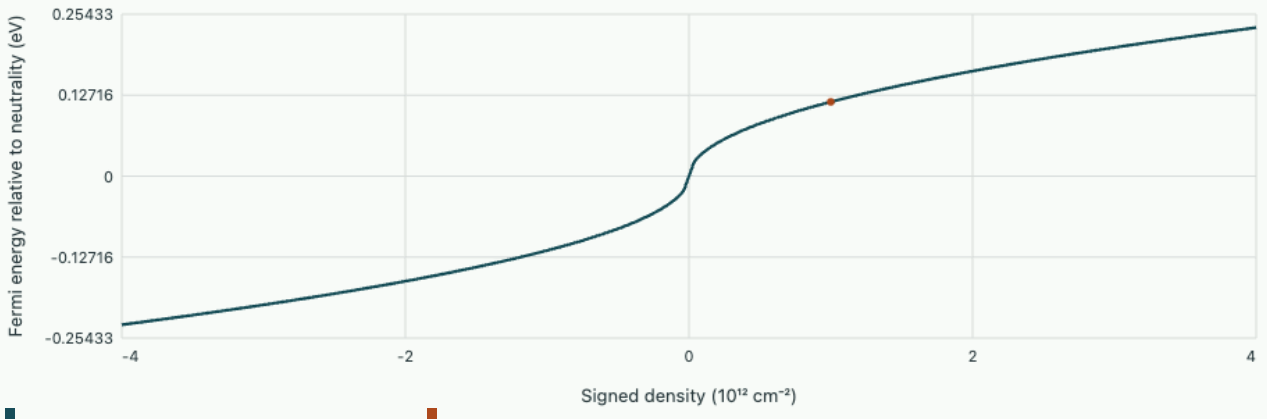

Fermi energy 0.1167 eV, wave number 0.1772 nm^{-1}, cyclotron mass 0.02052 electron masses.

Tutorial on a density sweep

Two cyclotron measurements on the same gapless sheet return masses $0.024 m_e$ and $0.036 m_e$. Assume the velocity is independent of density. Infer the ratio of absolute carrier densities and the change in the weak-field cyclotron frequency at fixed field.

▼ **Work through the solution**

Since m_c scales as $\sqrt{|n|}$, the second density is $(0.036/0.024)^2 = 2.25$ times the first. The frequency eB/m_c drops to two thirds of its first value. The masses alone cannot reveal whether either measurement used electrons or holes.

Find the weak-field cyclotron radius at the worked density $10^{12}\,\mathrm{cm}^{-2}$ and field $0.10\,\mathrm{T}$. How does raising the velocity at fixed density affect it?

▼ **Work through the solution**

The semiclassical radius is

$$r_c = v_F m_c/(eB) = \hbar k_F/(eB).$$

With $k_F = 1.772 \times 10^8\,\mathrm{m}^{-1}$,

$$r_c = 1.17\,\mu\mathrm{m}.$$

At fixed density and field the radius stays fixed when velocity changes. The orbit period changes because the speed around that circle changes.

A measurement doubles positive E_F at constant temperature sufficiently close to zero for the step filling approximation. How do state density per area and density of states per energy change?

▼ **Work through the solution**

For fixed v_F, the occupied density scales as E_F^2, so it rises fourfold. The derivative $D(E_F)$ scales as $|E_F|$ and doubles. Confusing the count with its derivative would give the wrong charge change inferred from spectroscopy.

Further reading · Novoselov and colleagues, measured energy dependence of the cyclotron mass.

Quantising the orbit and finding the zero Landau level

Condensed matter and band theory

In a magnetic field the two kinetic momentum components cannot be specified simultaneously. In a conical band their commutator turns the two sublattice equations into coupled oscillator equations. Their spectra differ by one state, leaving an unpaired level at zero energy.

Construct the oscillator from kinetic momentum

In one valley of a gapless monolayer use the long-wavelength Hamiltonian $H = v_F(\sigma_x \Pi_x + \sigma_y \Pi_y)$, with Pauli matrices acting on sublattices. For an electron of charge $-e$,

$$\boldsymbol{\Pi} = \boldsymbol{p} + e\boldsymbol{A}.$$

In a uniform perpendicular field $B > 0$, applying $[p_j, A_k] = -i\hbar \partial_j A_k$ gives

$$[\Pi_x, \Pi_y] = -ie\hbar B.$$

Introduce $a = (\Pi_x - i\Pi_y)/\sqrt{2e\hbar B}$. Its adjoint uses the opposite sign, giving $[a, a^\dagger] = 1$. The magnetic length $\ell_B = \sqrt{\hbar/(eB)}$ sets the spatial size of the oscillator.

$$H = v_F\sqrt{2e\hbar B}\begin{pmatrix} 0 & a \\ a^\dagger & 0 \end{pmatrix}, \qquad H^2 = 2e\hbar B v_F^2 \begin{pmatrix} aa^\dagger & 0 \\ 0 & a^\dagger a \end{pmatrix}.$$

In the B component the squared Hamiltonian has oscillator numbers $0, 1, 2, \ldots$. In the A component $aa^\dagger = a^\dagger a + 1$, so its sequence starts at one. For integer $n \geq 1$ combine $|n-1\rangle$ on A with $\pm|n\rangle$ on B in equal magnitudes. Acting with H confirms energies $E_{n,\pm} = \pm v_F\sqrt{2e\hbar B n}$. For $n = 0$, the B-only spinor with $|0\rangle$ is annihilated by a. Its energy vanishes. Valley and field conventions can exchange which sublattice carries that state.

Count independent centres at the same energy

Choose gauge $A_y = Bx$ and retain translation along y. Quantisation gives

$$k_y = 2\pi j / L_y.$$

The kinetic momentum $\Pi_y = \hbar k_y + eBx$ places the orbit centre at $x_0 = -\hbar k_y/(eB)$. Successive centres differ by $2\pi\ell_B^2/L_y$. Counting centres in a width L_x gives $L_x L_y/(2\pi\ell_B^2)$, or eB/h states per area for each spin and valley, apart from edge corrections. This degeneracy is independent of level index. Spin splitting, lattice effects and interactions are set aside here.

At $B = 1\,\mathrm{T}$, $v_F = 10^6\,\mathrm{m/s}$, $\ell_B = 25.7\,\mathrm{nm}$ and the first positive level is $36.3\,\mathrm{meV}$ above zero. The next sits at $51.3\,\mathrm{meV}$. The spacing between the first two positive levels is consequently only 15.0 meV. At 4 T all these energies double while the number of orbit centres per unit area quadruples. A parabolic band would instead give equally spaced levels at fixed field.

Identify the measurable dependence

Compare the first few positive levels as field changes. Dividing E_n by $\sqrt{Bn}$ should return one material-dependent coefficient. An optical measurement sees transitions between occupied and available states, subject to matrix-element selection rules. A measured transition energy therefore need not coincide with the energy of an individual level. At zero field this discrete-level construction loses its normalisable magnetic oscillator; the limiting energies on the graph describe the approach from finite field.

Open this experiment online

Try the calculation

Perpendicular field (T) 1

Cone velocity (10^6 m/s) 1

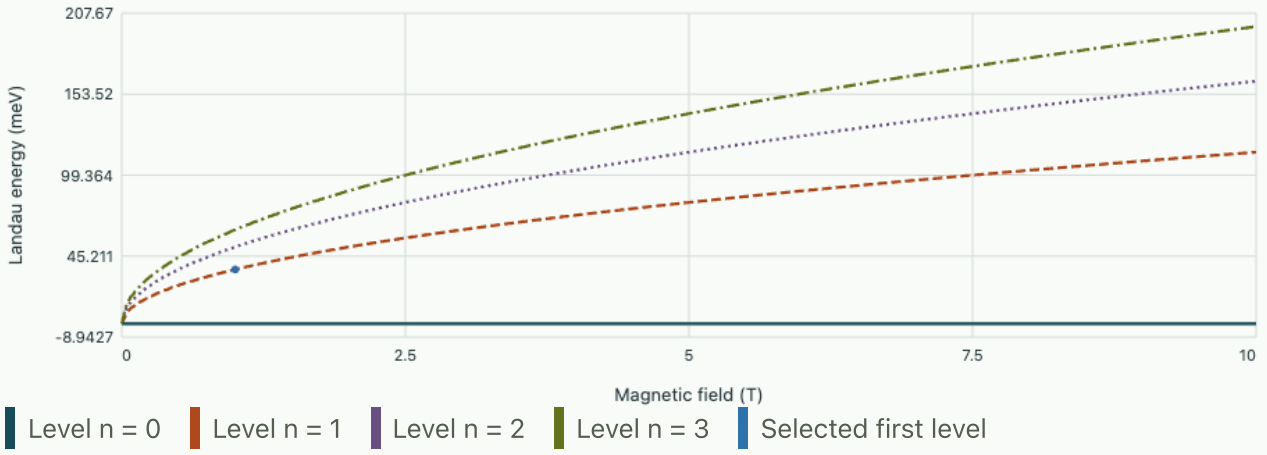

First level 36.28 meV, first-to-second spacing 15.03 meV. Each spin and valley has 2.418 × 10¹⁰ orbit centres per cm².

Tutorial on two spectral lines

At $B = 2\,\mathrm{T}$ an ideal sample has a zero-to-first level spacing of $48\,\mathrm{meV}$. Infer its velocity relative to $10^6\,\mathrm{m/s}$. Predict the first-to-second positive-level spacing at $8\,\mathrm{T}$.

▼ **Work through the solution**

The reference first level at 2 T is $36.3\sqrt{2} = 51.3\,\mathrm{meV}$, so $v_F \simeq 0.935 \times 10^6\,\mathrm{m/s}$. Raising the field by a factor of four doubles all energies. The first level at 8 T is 96 meV, and $E_2 - E_1 = 96(\sqrt{2} - 1) = 39.8\,\mathrm{meV}$. This inference assumes an ungapped cone without interaction shifts.

At $B = 2\,\mathrm{T}$, estimate the number of orbit centres in area $1\,\mu\mathrm{m}^2$ for a single level. Give both the per-flavour count and the full spin–valley count.

▼ **Work through the solution**

Use eBA/h, with $A = 10^{-12}\,\mathrm{m}^2$. This gives roughly 484 centres per flavour and 1934 across four flavours. The numbers are approximate because a finite sample admits edge corrections, whereas the area formula is a bulk degeneracy density.

At high positive level index show how the level spacing regains the semiclassical cyclotron frequency eB/m_c.

▼ **Work through the solution**

Write $E_n^2 = 2e\hbar B v_F^2 n$. At large n the difference $E_{n+1} - E_n \simeq dE/dn = e\hbar B v_F^2 / E_n$. Dividing by $\hbar$ gives eB/m_c with $m_c = E_n / v_F^2$. This connection uses many occupied levels, while the unpaired zero state is outside the semiclassical limit.

Further reading · Jiang and colleagues, infrared spectroscopy of Landau levels in monolayer graphene · Li and colleagues, spatially resolved graphene Landau levels.

Identify mobile charge with a transverse field

Fields and the response of matter

Derive a Hall-voltage measurement from transverse force balance, reverse current and field to reject offsets and explore what fails when two carrier types conduct.

Choose signs from a drawing

A rectangular conducting sample has its long axis along x, width w along y, and thickness t along z. A uniform magnetic field $B\hat{\mathbf{z}}$ is applied. An external circuit maintains conventional current $I > 0$ in $+x$. High-impedance transverse probes draw negligible current. Define the measured signal as $V_H = V(y = 0) - V(y = w)$. For uniform transverse field this convention gives

$$V_H = E_y w.$$

Reversing the voltmeter leads reverses the reported sign; writing a carrier sign without a lead convention is incomplete.

With one charge carrier species of number density n and signed charge q, the drift velocity satisfies

$$j_x = nqv_x = I/(wt).$$

If the transverse velocity vanishes, its magnetic force component is $-qv_xB$, opposed by qE_y. Thus $E_y = v_xB = j_xB/(nq)$. The resulting port measurement is

$$V_H = \frac{IB}{nqt}.$$

The width cancels for a uniform bar; changing the contact geometry or using a current source with a nonuniform injection profile can spoil that simplification near the terminals. The cancellation also exposes the sign. For electrons $q < 0$, positive conventional current means that the mobile charge drifts in $-x$. The inferred transverse field is negative for positive B in the chosen convention.

Reconstruct density and mobility from separate observations

Take $t = 0.50$ mm, $I = 5.0$ mA, $B = 0.20$ T and a transverse signal of -1.25 mV. The inferred signed Hall coefficient is

$$R_H = tV_H/(IB) = -6.25 \times 10^{-4}\ \mathrm{m^3/C}.$$

Using singly charged electrons gives

$$n = 1/(e|R_H|) \simeq 1.00 \times 10^{22}\ \mathrm{m^{-3}}.$$

A longitudinal voltage measurement, after correcting for probe separation, can yield the material conductivity. If $\sigma = 80$ S/m, the single-carrier mobility is

$$\mu = \sigma|R_H| = 0.050\ \mathrm{m^2/(V\,s)}.$$

These operations use different pieces of information. The longitudinal measurement alone gives only the product $n|q|\mu$. The transverse measurement supplies a sign and, in the one-carrier model, the density independently of mobility. This is an inverse experiment whose two observables have different sensitivities, not a more precise measurement of the same scalar.

One can derive the same result from a stationary Drude equation $0 = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}) - m\mathbf{v}/\tau$. The longitudinal drag gives

$$E_x = mv_x/(q\tau);$$

force balance in y gives the transverse field just obtained. In a homogeneous single-carrier sample, $E_y/E_x = q\tau B/m$, the signed Hall angle tangent. At fixed sample and field, doubling I doubles both voltages in this linear regime and leaves their ratio unchanged. Significant current-dependent apparent density can indicate heating or nonlinear contacts.

Isolate the transverse voltage from a misaligned contact

Practical side contacts seldom sit at exactly the same x. A fraction of the longitudinal voltage can leak into the transverse pair. Write the readout at a fixed current as $V_{\mathrm{m}}(B) = V_0 + \alpha V_{xx} + V_H(B)$. A background that is even in field cancels in $[V_{\mathrm{m}}(B) - V_{\mathrm{m}}(-B)]/2$. For example, readings of 3.05 mV and 5.55 mV at $+B$, $-B$ return -1.25 mV Hall signal and 4.30 mV of even background. Taking a single measured sign would have diagnosed the wrong carriers.

Open the live experiment

Try the calculation

Electron density (10^{22} m^{-3}) 1

Current (mA) 5

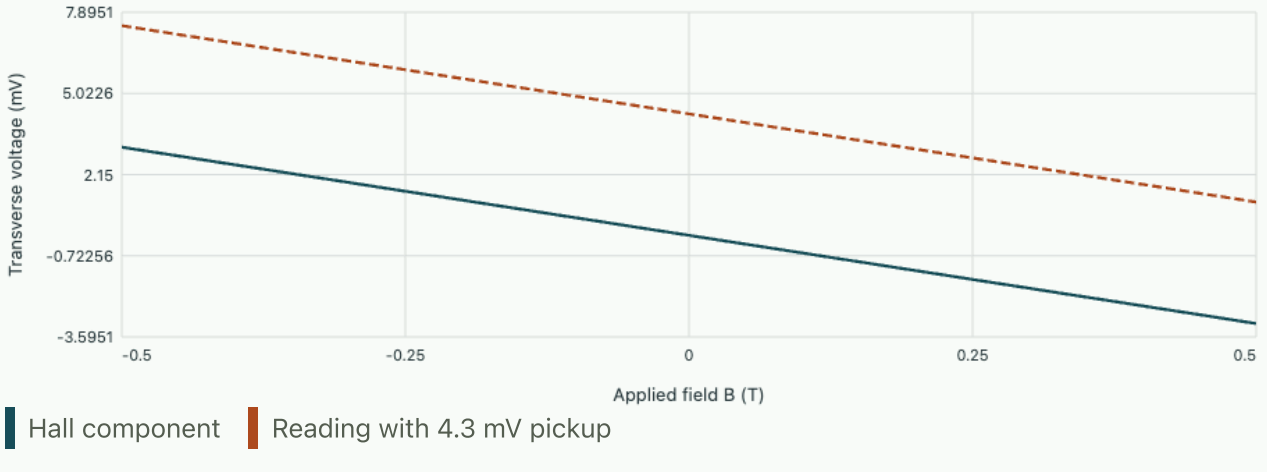

At +0.2 T the instrument reads 3.0517 mV; at −0.2 T it reads 5.5483 mV. Half-difference -1.2483 mV.

The plot deliberately superposes a constant 4.3 mV pickup on a signed Hall signal. Vary density and current independently before reading the pair at $B = \pm 0.2$ T. Large density weakens the signal; current increases it.

Verify that field reversal recovers the correct transverse voltage even when both instrument readings have the same sign. If the unwanted background changes between readings, reverse the field several times and check for repeatable drift.

With both current and field reversal, the double-odd combination $[V(+I,+B)-V(+I,-B)-V(-I,+B)+V(-I,-B)]/4$ isolates terms odd in both. A thermal emf that is independent of current cancels along with a stable contact offset. This symmetry decomposition is a physical assumption about the parasitics, so preserve all four readings and inspect the rejected channels too.

When a Hall density stops being a particle count

With electrons and holes, longitudinal conductivities add. In a weak magnetic field their transverse contributions also depend on squared mobilities. For positive mobilities μ_e, μ_h and electron and hole densities n, p, the two-carrier weak-field coefficient in the sign convention above is $R_H = (p\mu_h^2 - n\mu_e^2)/[e(p\mu_h + n\mu_e)^2]$. To see where the square appears, note that one factor of mobility converts the longitudinal field into drift; a second converts the resulting transverse magnetic force into a drift contribution. Summing species currents and imposing zero total transverse current gives the denominator management through the summed longitudinal conductivity.

Equal electron and hole densities with equal mobilities give zero Hall signal while current still flows readily. A single-species inversion would suggest unbounded density. At equal densities but unequal mobilities, the faster species can set the sign even though neither species is more numerous. Additional field dependence, temperature variation and independent structural information are needed before treating a Hall number as a carrier census.

Tutorial

A measurement doubles the width w at fixed thickness, total current and field in a uniform sample. Predict both E_y and V_H.

▼ **Work through the solution**

The current density halves, so v_x and E_y halve. The voltage integrates that halved field over twice the width and stays the same. If fixed current density were held instead, the transverse voltage would double.

A nominal thickness is overestimated by 5%. What bias does this cause in the single-carrier density inferred from a fixed transverse signal?

▼ **Work through the solution**

$n = |IB|/(et|V_H|)$, so using $1.05t$ gives $n_{\mathrm{fit}} = n/1.05$, approximately 4.8% low. This systematic error is common to all repeated field sweeps.

Why might a metal yield a smaller Hall voltage than a semiconductor with the same dimensions at the same current?

▼ **Work through the solution**

In the one-carrier regime the signal is inversely proportional to carrier density. A large density allows the imposed current at a very small drift speed, requiring a correspondingly small transverse balancing field. Conductivity alone cannot determine this comparison without mobility information.

Follow the transport model

Compare the one-carrier assumptions with the treatment of mobile carriers in David Tong's solid-state physics course. Connect an independently measured conductivity with the signed transverse response before interpreting an effective mass or scattering time.

Distinguish magnetic screening from electrical diffusion

Fields and the response of matter

Compare the frequency dependence, phase and thickness dependence of field penetration in a normal metal and a superconducting slab.

Perfect conductivity and magnetic expulsion are different statements

A normal conductor has current density $\mathbf{J} = \sigma\mathbf{E}$. Combining Faraday's and low-frequency Ampère laws in a homogeneous medium gives

$$\partial_t \mathbf{B} = (\mu\sigma)^{-1}\nabla^2\mathbf{B}.$$

A step in field diffuses inwards over a length of order $\sqrt{t/(\mu\sigma)}$. Under sustained sinusoidal drive the skin depth is $\delta = \sqrt{2/(\mu\sigma\omega)}$. As ω tends to zero, this skin depth tends to infinity. An ordinary nonmagnetic conductor in its final static equilibrium generally admits a uniform applied magnetic field.

The London description of an isotropic superconductor in a Meissner state makes a distinct, equilibrium constitutive assumption:

$$\nabla \times \mathbf{J}_s = -\mathbf{B}/(\mu_0\lambda_L^2),$$

with penetration length λ_L. In a homogeneous specimen carrying supercurrent and at frequencies where displacement current and normal losses may be neglected,

$$\nabla \times \mathbf{B} = \mu_0\mathbf{J}_s.$$

Curl once more and use $\nabla \cdot \mathbf{B} = 0$. The resulting field satisfies

$$\nabla^2\mathbf{B} = \mathbf{B}/\lambda_L^2.$$

Unlike the diffusive skin depth, the equilibrium length remains finite for a static applied field.

Taking the formal limit $\sigma \to \infty$ in a normal conductor suppresses the electric field for a finite current and can freeze a pre-existing flux during changes. That argument does not determine which static field a specimen adopts when cooled through its transition. Magnetic expulsion is extra thermodynamic physics. The comparison below assumes a weak enough field and a suitable state, omitting flux vortices, trapped flux and nonlinear or nonlocal effects.

Solve a slab with equal applied fields on both faces

Let a broad slab fill $-d/2 < x < d/2$, and impose identical tangential magnetic fields $B_a\hat{\mathbf{z}}$ on its two faces. Symmetry requires $B'(0) = 0$. Applying both boundaries to the London equation gives

$$B(x) = B_a\frac{\cosh(x/\lambda_L)}{\cosh[d/(2\lambda_L)]},$$

\frac{\overline B}{B_a}=\frac{2\lambda_L}{d}\tanh\frac{d}{2\lambda_L}.

Differentiating gives the screening current $J_y = -B_z'/\mu_0$. Currents flow in opposite directions at opposite faces; their net transport current across the thickness is zero. For a thick slab each face approaches a half-space with exponential decay. For $d \ll \lambda_L$, the two surface layers overlap. Expanding the average gives $1 - d^2/(12\lambda_L^2) + \cdots$, so even a superconductor can have a substantial mean internal field in a sufficiently thin specimen under this parallel-field geometry.

Take $d = 100$ nm and $\lambda_L = 50$ nm. The centre field is $B_a/\cosh 1 \simeq 0.6481 B_a$; the mean is $\tanh 1 \simeq 0.7616 B_a$. The distinction between centre value and volume average is already larger than many instrumental error bars. A measurement of magnetic moment in a finite non-slab geometry additionally requires its demagnetising field. Do not directly substitute the infinite-slab mean into an arbitrary sample-shape calibration.

Obtain the equilibrium length as an energy minimum

The same static profile follows from minimising $\int_{-d/2}^{d/2} [B^2 + \lambda_L^2 (B')^2]\, dx/(2\mu_0)$ with boundary values imposed. The first term is magnetic field energy; the second is the energy associated with the screening current in this effective description. Vary B by a function η vanishing at the faces. Integrating $B'\eta'$ by parts leaves an interior factor $B - \lambda_L^2 B''$. Its vanishing returns the London equation. This route makes the physical tradeoff clear: complete expulsion would require extremely sharp gradients and costly currents, while uniform penetration pays magnetic energy throughout the bulk.

The reduction assumes a fixed superconducting condensate amplitude. Near a critical temperature or critical field, allowing that amplitude to change introduces an additional coherence length and may alter both the field state and the validity of London theory. A numerical fit giving a stable value of λ_L at one temperature does not prove this approximation over an entire phase diagram.

Compare two experimental controls

Open the live experiment

Try the calculation

Slab thickness d (nm) 100

London length λ (nm) 50

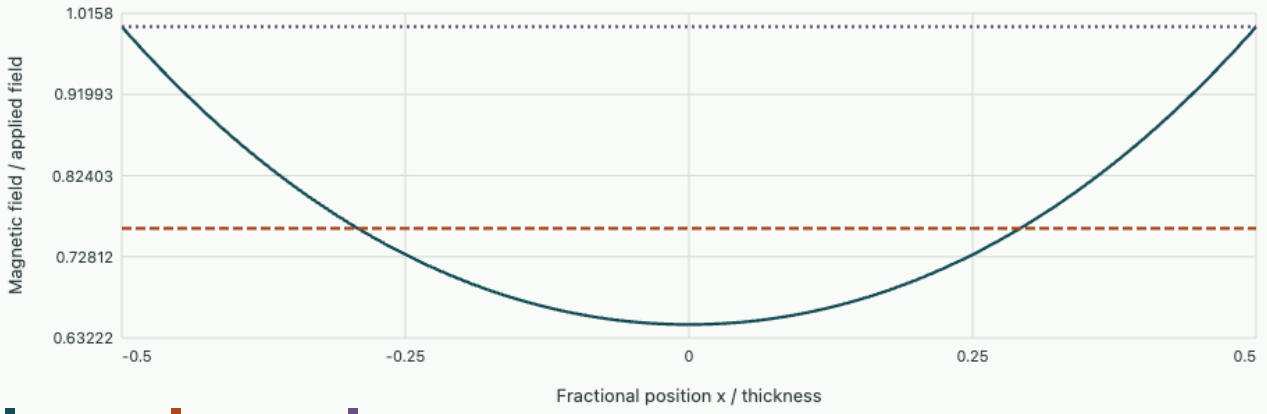

At d/λ = 2 the centre retains 64.805% of the applied field and the thickness average 76.159%.

Change the thickness and the penetration length independently. Find two different specimens with the same d/λ_L; their normalised profiles against fractional thickness match, even though their local physical gradients and current densities can differ. If all you can measure is a single normalised profile, geometry calibration is part of the inference of λ_L.

A frequency sweep supplies another distinction. In a normal sheet the complex field has both attenuation and phase lag with depth, and δ changes as $f^{-1/2}$. The static London profile above is real and determined by a material length. At higher frequencies superconductors can also exhibit complex response and dissipation, so measuring a single attenuation length in one frequency band cannot by itself establish static flux expulsion. Use the normal-film impedance experiment for a controlled comparison with its different rear boundary condition.

Tutorial

In a thick specimen, how far from one isolated face must a probe be to see 1% of the applied static field? How does the answer depend on field amplitude in this model?

▼ **Work through the solution**

The half-space profile is e^{-x/λ_L}. Set it to 0.01 to obtain

$$x = \lambda_L \ln 100 \simeq 4.605\lambda_L.$$

Linearity makes this scaled distance independent of B_a, until the assumed Meissner and weak-field conditions fail.

For a symmetric slab with $d = 4\lambda_L$, compute the centre and mean fractions.

▼ **Work through the solution**

The centre gives $1/\cosh 2 \simeq 0.266$. The mean is $\tanh 2/2 \simeq 0.482$. Both surface layers contribute substantially to the average even when the centre field is much smaller.

A phase-sensitive measurement detects a substantial depth-dependent phase lag in a slowly driven slab. Does the static London solution predict this?

▼ **Work through the solution**

The static real solution has no phase lag. Assess normal-fluid losses, vortex motion, measurement phase reference and the finite drive frequency. Comparing only magnitudes can conceal these distinct mechanisms.

Read the equilibrium model

The MIT applied-superconductivity course places London's laws alongside microscopic and order-parameter treatments. Use that ordering to identify the assumptions needed when taking a static limit.

How a magnetic field rotates the polarisation of transmitted light

Fields and the response of matter

Resolve a linearly polarised wave into circular eigenmodes, infer a refractive-index difference from an analyser and distinguish nonreciprocal rotation from a simple change of reference axes.

Derive rotation from propagation rather than drawing an arrow

Let monochromatic light propagate along $+z$ through a transparent magneto-optic sample, with an externally prepared axial magnetic field. Define fixed laboratory vectors $\mathbf{e}_\pm = (\hat{x} \pm i\hat{y})/\sqrt{2}$ and use phasors $e^{i(kz-\omega t)}$. These definitions specify the handedness algebraically. Calling a vector simply “right circular” without stating the viewing and time conventions can reverse a reported rotation sign.

An input linearly polarised in x is $(\mathbf{e}_+ + \mathbf{e}_-)/\sqrt{2}$. Suppose the medium supports wave numbers $k_\pm = 2\pi n_\pm/\lambda_0$ for the two bases, with negligible absorption. After a distance L, factor out the common phase $e^{i(k_+ + k_-)L/2}$. The remaining Cartesian phasor is $\hat{x}\cos\Delta - \hat{y}\sin\Delta$, where $\Delta = (k_+ - k_-)L/2$. Thus the polarisation angle in the fixed x, y plane is

$$\theta = \frac{\pi L}{\lambda_0}(n_- - n_+).$$

Equal transmissions of the two components preserve a linearly polarised output. Unequal absorption makes the circular amplitudes unequal and generally produces an ellipse. Rotating a linear analyser then finds a finite minimum, even for a perfectly coherent field. Residual intensity at the nominal extinction angle can also arise from stray light or imperfect polariser extinction, which must be measured independently.

Derive a magneto-optic eigenbasis

For axial magnetisation, a simple lossless transverse relative-permittivity block is $\left(\begin{smallmatrix} \epsilon & ig \\ -ig & \epsilon \end{smallmatrix}\right)$ with real g at the frequency of observation. Acting on $(1, i)$ gives $(\epsilon - g)(1, i)$; acting on $(1, -i)$ gives $(\epsilon + g)(1, -i)$. For relative permeability near one, $n_+^2 = \epsilon - g$, $n_-^2 = \epsilon + g$. At small g their difference $n_- - n_+ \simeq g/\sqrt{\epsilon}$. Reversing magnetisation reverses the sign of the off-diagonal coefficient and interchanges the two eigenvalues. Both the existence of circular modes and the sign-reversal test now follow from a constitutive model.

If $n_- - n_+$ is proportional to magnetic field over a specified range, write $\theta = VBL$, with Verdet constant V in rad/(T m). This is a measurable material and wavelength dependent coefficient. It is not determined by the mean refractive index alone. A strong nearby absorption line can give substantial frequency dependence and dichroism, invalidating a lossless narrow-band extrapolation.

Work through an analyser calibration

For $\lambda_0 = 600$ nm, $L = 0.20$ mm and measured index difference 5.0×10^{-4}, the rotation is $\pi(2.0 \times 10^{-4})(5.0 \times 10^{-4})/(6.0 \times 10^{-7}) = \pi/6$, or 30°. An ideal analyser at angle α transmits $I/I_{\text{in}} = \cos^2(\alpha - \theta)$. Its maxima occur at 30° and 210°; an intensity measurement determines a linear

polarisation angle modulo 180°. For a sample thick enough to rotate several half turns, an intensity maximum alone cannot give an unwrapped θ. Vary field or path length continuously from a known reference.

If the applied field is independently 0.40 T in that example, $V = \theta/(BL) \simeq 6.54 \times 10^3$ rad/(T m). The arithmetic is dimensionally correct, but the large value warrants a physical check of thickness, field calibration, units and material spectrum before reporting it as an ordinary glass coefficient. Degrees substituted for radians would make the inferred material coefficient too large by $180/\pi$.

Use field reversal to reject alignment errors

Open the live experiment

Try the calculation

Index difference (10^{-4}) **5**

Sample length (mm) **0.2**



Unwrapped single-pass rotation 30°. At a 45° analyser the two transmitted fractions are 0.93301 and 0.066987.

The experiment shows the predicted analyser traces for opposite index splitting, with the same input polariser. Increase sample thickness to see why tracking the maximum by one measurement can select the wrong branch. Reversing the splitting reverses the tilt in fixed laboratory axes. A fixed analyser offset shifts both inferred angles equally and cancels in half their difference, provided the branch has been unwrapped consistently.

For small rotation, an analyser exactly at a crossed position transmits intensity proportional to θ^2. Its response has no sign sensitivity close to zero. Set the analyser at 45° instead:

$$\cos^2(\pi/4 - \theta) = [1 + \sin(2\theta)]/2 \simeq 1/2 + \theta.$$

The difference between fields $+B$ and $-B$ then measures a signed signal linear in small rotation. This trades extinction of the background for sensitivity to sign. Balance the two orthogonal analyser outputs if source intensity fluctuates appreciably during a sweep.

Ask what happens on a return pass

For the ideal Faraday medium at fixed magnetisation, the circular eigenmodes relative to a fixed laboratory basis are determined by its transverse material tensor. On propagation back through the sample, the accumulated positive propagation path again distinguishes the same laboratory circular basis after accounting for the reflection boundary. With a suitable simple mirror at normal incidence, the resulting material rotation doubles on a round trip in fixed transverse coordinates. In a reciprocal optically active material, the eigenvalue assignment with respect to the laboratory basis reverses on reversing propagation, and the single-pass rotation cancels on the ideal return. Write both transformations in one fixed basis before comparing signs.

A real double-pass setup includes mirror phase shifts, stress birefringence of windows and possible ellipticity. Reversing magnetic field, removing the sample and recording both analyser quadratures provide separate checks. A nonzero rotation in one transmitted beam proves only that its polarisation changed; establishing nonreciprocity requires comparing a reversed propagation experiment with a stated polarisation convention.

Tutorial

Two analyser maxima occur at 47° and 39° for opposite weak fields. If the response is linear and the same branch is tracked, what field-induced rotation and static offset do they imply?

▼ **Work through the solution**

Half the difference gives 4 degrees of signed rotation for the first field relative to zero. The mean is 43 degrees, the combined input polarisation and analyser zero. Conversion to a Verdet constant additionally needs the applied field and path length.

In the 600 nm example, what thickness uncertainty corresponds to 1 degree of rotation uncertainty at fixed index difference?

▼ **Work through the solution**

Rotation is proportional to thickness. Relative uncertainty $1/30$ therefore gives about $0.20/30 = 0.0067$ mm, or 6.7 µm. Uncertainty in wavelength or field would add additional contributions.

An ideal analyser gives a minimum 4% of its maximum. Can this be explained solely by a rotation of an ideal linearly polarised monochromatic input in a lossless circularly birefringent sample?

▼ **Work through the solution**

Pure rotation preserves a linear state, so an ideal analyser should extinguish it completely. Differential loss, ellipticity, depolarisation, finite bandwidth or a detector background are possible alternatives. Use circular-analysis or full Stokes measurements and an empty-path baseline to distinguish them.

Compare with a full polarisation measurement

Use the Stokes reconstruction experiment to determine which intensities reveal a circular component. A fitted linear rotation alone cannot measure that missing component.

Relativistic particle dynamics and the field tensor

Waves, radiation and relativity

A circular track measures momentum per charge. A clock following the particle registers proper time, and a detector records how much kinetic energy it has gained. We will put these observations in one dynamical equation, solve that equation in uniform fields and exhibit a numerical pusher that can get the energy right while its orbit phase slowly slips.

What the track, clock and calorimeter measure

Consider a particle of invariant mass m, signed charge q and ordinary velocity $\mathbf{v}$. Inertial observers agree on the proper time between neighbouring events on its timelike worldline,

$$d\tau = dt\sqrt{1 - v^2/c^2}.$$

Adopt coordinates $x^\mu = (ct, x, y, z)$ and metric $g = \mathrm{diag}(1, -1, -1, -1)$. Dividing the displacement by proper time produces a vector of fixed Minkowski norm. Multiplying it by invariant mass gives four-momentum.

$$u^\mu = \gamma(c, \mathbf{v}),$$
$$p^\mu = m u^\mu = (\mathcal{E}/c, \mathbf{p}),$$
$$\mathcal{E}^2 - c^2|\mathbf{p}|^2 = m^2 c^4.$$

Energy includes the rest term. Momentum is $\gamma m \mathbf{v}$. Differentiating the last equation on a trajectory gives

$$d\mathcal{E}/dt = \mathbf{v} \cdot d\mathbf{p}/dt.$$

This identity replaces the nonrelativistic step of dotting $m\mathbf{a}$ with velocity. A magnetic field can rotate momentum with fixed magnitude and energy, whereas an electric field can change energy by doing work.

Suppose an electron begins at rest and gains 10.0 keV of kinetic energy across an electrostatic potential difference in vacuum. Its rest energy is 511 keV, giving $\gamma = 1.01957$ and $v/c = \sqrt{1 - 1/\gamma^2} = 0.1950$. The nonrelativistic estimate is 0.1978. In a modest energy measurement the discrepancy could lie below the experimental spread. For a precise orbit calculation the 1.96 percent change in relativistic cyclotron period is measurable independently of that discrepancy in speed.

A linear map that produces the Lorentz force

Seek a local equation linear in field and four-velocity and covariant under boosts. For the energy component to give $q\mathbf{E} \cdot \mathbf{v}$, and for the spatial components to give the electric force and the oriented cross product, construct the antisymmetric field tensor. With the stated metric and $\epsilon^{123} = 1$, set

$$F^{0i} = -E_i/c,$$

$$F^{i0} = E_i/c,$$

$$F^{ij} = -\epsilon^{ijk} B_k,$$

$$\frac{dp^\mu}{d\tau} = qF^{\mu\nu} u_\nu.$$

Lowering the index on u reverses the signs of its spatial components. The spatial product becomes $q\gamma(\mathbf{E} + \mathbf{v} \times \mathbf{B})$, and division by γ converts the proper-time derivative to $d\mathbf{p}/dt$. The time component similarly yields

$$d\mathcal{E}/dt = q\mathbf{E} \cdot \mathbf{v}.$$

Contract the force with u_μ. The symmetric product $u_\mu u_\nu$ annihilates the antisymmetric tensor, which proves that the evolution preserves the mass shell.

To recover Maxwell's kinematic constraints from potentials, put $A^\mu = (\phi/c, \mathbf{A})$ and $F^{\mu\nu} = \partial^\mu A^\nu - \partial^\nu A^\mu$. Then $F^{0i} = (\partial_t A_i + \partial_i \phi)/c = -E_i/c$, while the spatial-spatial components are the magnetic curl. Adding a four-gradient to A^μ leaves the tensor unchanged because ordinary derivatives commute. The particle action $S = -mc^2 \int d\tau - q \int A_\mu dx^\mu$ gives the same force after integrating variations by parts with fixed endpoints. Gauge changes add only an endpoint term to this action. A trajectory measured with local forces cannot distinguish such potential representatives.

Compare the potentials and field-based description in the field-tensor chapter

Constant force, variable coordinate acceleration

Put a particle initially at rest in a uniform static electric field along x. Let $a_0 = qE/m$ be signed. Integration of $dp_x/dt = qE$ gives $p_x = ma_0 t$, so the speed tends to c even while momentum increases linearly. Proper-time evolution is simplest in terms of rapidity η, where $u^0 = c\cosh\eta$, $u^x = c\sinh\eta$. The two tensor equations imply $d\eta/d\tau = a_0/c$.

$$t = \frac{c}{a_0}\sinh(a_0\tau/c),$$

$$x = \frac{c^2}{a_0}\{\cosh(a_0\tau/c) - 1\},$$

$$\frac{dv_x}{dt} = \frac{a_0}{\gamma^3}.$$

The initial time and position are zero. The work $qEx = mc^2(\cosh\eta - 1)$ equals kinetic energy, which checks both the integration constant and the dimensional factor in the displacement. For an electron use the direction of the electric force as positive to simplify signs. At 20.0 kV/m it gains 1.0 keV in 5.0 cm. The energy gain follows from the static potential even though the speed does not grow exactly as the square root of distance. The field region must last long enough for the particle to cross it and fringe fields can alter the entry trajectory.

Compare momentum, coordinate time and proper time in the charged-worldline experiment

Helical motion, radius and period

A uniform magnetic field along z keeps p_z, energy and $|\mathbf{p}_\perp|$ fixed. Write $p_x + ip_y$ as a complex number. The Lorentz force reduces to a first-order rotation equation and a single time integration gives the corresponding position. With positive B, positive charges starting along $+x$ initially turn towards $-y$.

$$\frac{d}{dt}(p_x + ip_y) = -i\frac{qB}{\gamma m}(p_x + ip_y),$$

$$R = \frac{p_\perp}{|q|B},$$

$$T = \frac{2\pi\gamma m}{|q|B}.$$

The helical pitch is $v_z T$. Radius gives transverse momentum and a calibrated period gives total energy through γ. A track seen only in projection along the field gives no direct evidence that longitudinal momentum vanishes. A second view or a timing measurement can resolve this ambiguity.

At 10.0 keV the electron has $pc = \sqrt{K(K + 2mc^2)} = 101.6$ keV. With entry perpendicular to a 2.0 mT field its radius is 0.169 m and its period 18.2 ns. Doubling the field halves both radius and period at fixed injection energy. Doubling momentum at fixed field doubles radius but changes period by the ratio of Lorentz factors. These are different tests of a track calibration. A beam with a spread of energies will dephase relative to a narrow timing gate even when every single particle remains in a perfectly uniform field.

A drift frame gives an exact crossed-field solution

For constant crossed fields with $E < cB$, transform to the frame moving at $\mathbf{v}_d = \mathbf{E} \times \mathbf{B}/B^2$. There $E' = 0$ and $B' = \sqrt{B^2 - E^2/c^2}$; calculate the transformed initial momentum and integrate a helix. Transform two successive equal-phase events back to the lab. Their perpendicular displacement in the electric-free frame is zero. If the motion has no component parallel to $\mathbf{B}$, the entire displacement there is zero, and the average lab cycle velocity is exactly $\mathbf{v}_d$. The detailed lab trace need not be circular and the instantaneous lab energy may oscillate.

Take $E = 30.0$ kV/m along y and $B = 0.50$ mT along z. The drift is along $+x$ at 6.00×10^7 m/s. A particle injected at that velocity crosses the uniform fields on a straight line because the forces cancel. An initially stationary positive particle instead begins by accelerating along $+y$. These observations can use the same field with different injected particles. If E exceeds cB, this electric-free frame is unavailable. Transform to a magnetic-free frame and integrate electric acceleration instead.

Continuous rotation and a discrete experiment

A naive forward step $\mathbf{p}_{n+1} = \mathbf{p}_n + q\Delta t\, \mathbf{v}_n \times \mathbf{B}$ adds a vector perpendicular to the previous momentum. Its squared norm increases by the squared magnitude of that kick. The apparent heating will persist in perfectly static uniform magnetic fields. For a purely magnetic update, evaluate the direction using the mean of old and new momenta. At fixed γ, the resulting implicit two-dimensional linear map can be solved algebraically and is an exact rotation by a modified angle.

$$\delta = 2\arctan\left(\frac{\omega_c \Delta t}{2}\right),$$

$$\omega_c = \frac{qB}{\gamma m}.$$

The sign of rotation follows the earlier complex equation. Dotting the update with the mean momentum gives zero change in squared norm, which verifies the energy property even for relatively large time steps. A common electromagnetic pusher surrounds the magnetic rotation by two electric half-kicks. In spatially varying fields, specify at what time and position each field is sampled, and check spatial interpolation and time convergence. For ultrarelativistic nearly cancelling electric and magnetic fields, preserving rotation in a pure field alone does not ensure cancellation is represented accurately.

Solve the constant tensor without transforming every time step

When uniform electric and magnetic fields are constant in spacetime, the Lorentz equation is a linear constant-coefficient equation for four-velocity as a function of proper time. Nonlinearity in ordinary time hides that simple structure, because the rate at which proper time accumulates changes along the orbit. Define the mixed tensor $M^{\mu}{}_{\nu} = (q/m)F^{\mu\alpha}g_{\alpha\nu}$. In a uniform field it does not vary along the trajectory. A formal solution is then

$$u(\tau) = e^{M\tau}u(0),$$

$$x(\tau) - x(0) = \int_0^{\tau} e^{Ms}u(0)\,ds.$$

The coordinate formula is written as an integral since M can be singular, as it is in a purely magnetic field. Blindly replacing this integral by $M^{-1}(e^{M\tau} - I)$ fails in exactly such simple cases. A power series for the integral or a block matrix exponential treats zero eigenvalues without an artificial inverse. Antisymmetry of the tensor implies $M^{\mathsf{T}}g + gM = 0$. Differentiating $e^{M^{\mathsf{T}}\tau}ge^{M\tau}$ shows that it stays equal to g, making preservation of $u^{\mu}u_{\mu} = c^2$ a structural property of the exact mapping.

For parallel constant fields along z, the matrix decomposes into a boost in the ct, z subspace and a rotation in x, y. Write $\alpha = qE/(mc)$, $\Omega = qB/m$, and denote the initial four-velocity components by u_0^{μ}. The proper-time solution is

$$u^0 = u_0^0\cosh\alpha\tau + u_0^z\sinh\alpha\tau,$$

$$u^z = u_0^z\cosh\alpha\tau + u_0^0\sinh\alpha\tau,$$

$$u^x + iu^y = (u_0^x + iu_0^y)e^{-i\Omega\tau}.$$

The longitudinal and transverse operations commute because they act in different subspaces. In lab time the transverse rotation slows as the electric field increases energy, although the proper-time rotation frequency stays constant. A particle therefore traces a helix whose transverse radius $p_{\perp}/(|q|B)$ is fixed while the pitch evolves. A single measured turn time without a tagged position in the accelerating region is insufficient for an energy reconstruction.

As a numerical test use an electron, quoting magnitudes, with $B = 1.0$ mT, $E = 10.0$ kV/m and initial purely transverse speed $0.10c$. One magnetic turn takes 35.72 ns in proper time and accumulates electric rapidity

magnitude $|\alpha|\tau \simeq 0.2096$. The final Lorentz factor is $\gamma_i \cosh(0.2096) = 1.0272$, giving final kinetic energy about 13.9 keV, compared with 2.57 keV initially. The lab-time turn duration is $\gamma_i \sinh(|\alpha|\tau)/|\alpha| \simeq 36.2$ ns. Reversing the electric field changes the direction of the axial displacement but leaves final energy and turn duration unchanged for this initially transverse particle. Reversing the initial longitudinal velocity would generally break that energy symmetry.

Infer uncertainty and test units in a momentum reconstruction

A track fitter commonly reports field, radius, pitch, time and a fitted covariance matrix. The dynamical relation turns these observed quantities into momentum. Let L_z denote a signed pitch per turn. With known signed direction of circulation and calibrated absolute field magnitude,

$$|\mathbf{p}| = |q| B \sqrt{R^2 + (L_z/2\pi)^2}.$$

For small errors, logarithmic differentiation gives weights for fractional radius and pitch errors:

$$d\log p = d\log B + w_R\, d\log R + w_L\, d\log |L_z|,$$

$$w_R = \frac{R^2}{R^2 + (L_z/2\pi)^2},$$

$$w_L = 1 - w_R.$$

For $R = 0.20$ m and $L_z = 0.30$ m the radius accounts for 94.6 percent of the squared transverse-plus-longitudinal geometric factor. Independent relative uncertainties of 0.20 percent in field, 0.50 percent in radius and 2.0 percent in pitch yield about 0.525 percent uncertainty in total momentum. A common detector scale controlling both radius and pitch would enter with correlated signs and should not be counted as two independent errors. When the pitch tends to zero, use the original derivatives with respect to L_z; a reported fractional error on a quantity consistent with zero is ill defined.

Differentiating $K = \sqrt{m^2c^4 + p^2c^2} - mc^2$ gives

$$d\log K = (\gamma + 1) d\log p / \gamma.$$

The coefficient approaches two at low energy and one at very high energy. At the approximately 3.70 keV energy of the tutorial the above momentum error would imply about a 1.05 percent energy error. A kinetic-energy estimate from a small difference of nearly equal numbers can also lose numerical precision at much lower energies. Rationalise it as $K = p^2c^2/(\sqrt{m^2c^4 + p^2c^2} + mc^2)$ when working near rest.

Several independent checks use this reconstruction without assuming the track fitter is correct. The turn time inferred from the fitted γ can be compared with a separate timing channel. Reverse the magnetic field to reverse the direction of curvature while keeping the speed and energy approximately fixed. Translate the beam within a mapped uniform field and look for changing fitted radius that correlates with position, an indication of geometric or field-map bias. Finally plot the residual of the magnetic force law separately for transverse magnitude and angle. A time integrator may conserve momentum magnitude for symmetry reasons even when its field interpolation biases the phase or shifts the guiding centre systematically.

Advance a charge using relativistic momentum

Increase the impulse while following speed, energy and proper time. Identify the range where a Newtonian momentum update becomes inaccurate.

Open this experiment online

Try the calculation

Electric field magnitude (kV/m) **100**

Lab time for a reading (ns) **17**

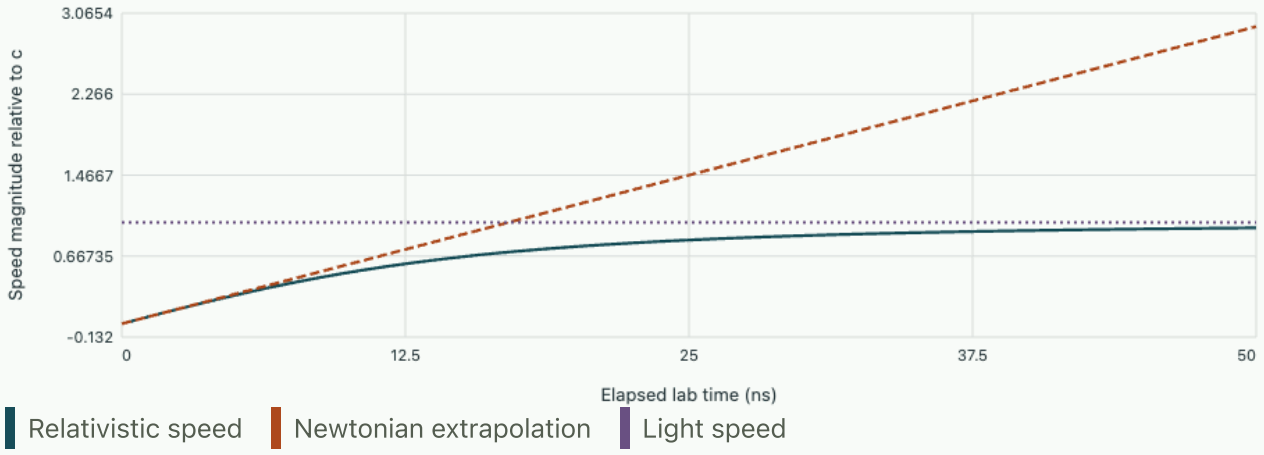

At 17 ns laboratory time the electron speed is 0.706 c. Its clock has advanced 15 ns and it has gained 211 keV.

Worked timing problem · infer an accelerating voltage without a radius

A tagged electron bunch completes ten nearly circular turns in a calibrated 5.0 mT magnetic field. Two time stamps indicate an elapsed 73.0 ns. The field is perpendicular to the orbit. Assuming complete turns and negligible electrical work during the timed interval, estimate the injection energy. What timing precision would distinguish a one-percent change in that energy?

▼ **Use the relativistic period sensitivity rather than the energy fraction directly**

The nonrelativistic turn period is

$$2\pi m/(eB) = 7.145\,\mathrm{ns}.$$

The measured period is 7.30 ns, giving $\gamma \simeq 1.0217$ and $K = (\gamma - 1)mc^2 \simeq 11.1$ keV. The relative energy sensitivity is much larger than the relative period change when γ is close to one. A one-percent change in K corresponds to $\Delta\gamma \simeq 2.17 \times 10^{-4}$, or relative period change 2.12×10^{-4}. Over the ten turns this gives approximately 15.5 ps.

To resolve that change spectroscopically, uncertainties in the elapsed time and field scale must together be significantly smaller than the equivalent change. A constant pickup cable delay cancels if the same electronics register both ends of the interval, while a drift of its timebase scale does not. Timing over more turns improves the total accumulated signal but also increases exposure to field inhomogeneity, energy loss and dephasing among particles with different injection energies. A narrower timing histogram does not automatically eliminate a systematic field calibration.

Finally verify the integer turn count independently. Missing a pickup crossing would cause a much larger change in inferred energy than the desired resolution. The measured wave packet may broaden so strongly that later crossings overlap. Recording the full pickup waveform and tracking ordered crossings, along with an independent rough radius estimate, can constrain the turn assignment before fitting the small relativistic shift. A perfect circular image with no time labels cannot supply that missing count.

Why momentum and kinetic energy do not gain fixed fractions together

In an electrostatic accelerating gap, the kinetic energy increment for a particle passing between two equipotential surfaces is its charge times the potential drop, regardless of injection energy. This statement follows by integrating the Lorentz power law over the actual path. A magnetic guiding field can alter the path and transit time while doing no direct work. A time-dependent electric field, a missed equipotential endpoint or appreciable radiation changes the simple inference and must be addressed separately.

For an electron entering with 10 keV and acquiring an additional 10 keV, kinetic energy doubles. Its initial momentum times c is about 101.6 keV and its final value about 144.4 keV. The radius in the same subsequent perpendicular magnetic field increases by a factor approximately 1.42. Doubling the radius would instead require doubling momentum, which demands a substantially greater energy increment at this low initial energy. Comparing radii before and after a known voltage supplies an independent check on the sign and magnitude of a particle's charge-to-mass model.

A measured energy gain below the nominal voltage does not by itself diagnose radiation. The potential actually traversed, finite angular acceptance, entrance or exit fringe fields, collision losses and biases in reconstructing longitudinal momentum can all matter. Repeat with reversed field polarity or several injection energies and compare the kinetic-energy difference rather than the two radii directly.

Tutorial · Reconstruct a track with two measurements

A particle known to be an electron or positron crosses a uniform $+z$ field of 1.0 mT. The projected radius is 0.200 m and the pitch magnitude is 0.300 m. Recover the transverse and longitudinal momentum, energy and period. What else must the detector record to determine the sign of charge?

▼ **Reconstruct both momentum components**

The transverse momentum is $|q|BR = 3.20 \times 10^{-23}$ kg m/s. During a cycle the parallel displacement is $p_z T/(\gamma m) = 2\pi p_z/(|q|B)$, so $|p_z| = 7.65 \times 10^{-24}$ kg m/s. The quadrature total is 3.29×10^{-23} kg m/s, yielding $p/(mc) = 0.1206$, $\gamma = 1.00724$, and approximately 3.70 keV kinetic energy. The period is about 36.0 ns.

Neither radius nor pitch magnitude carries the signed circulation. Record the direction of traversal relative to the calibrated direction of magnetic field. Reversing both charge and the direction of traversal yields the same unlabelled transverse ring. If timing or two ordered detectors supply a direction, the sign follows from the first Lorentz kick.

Tutorial · Can a constant-energy circle be trusted?

A pusher uses 20 uniform steps per exact cyclotron period in a uniform magnetic field. It reproduces initial kinetic energy to machine precision for 50 turns. Estimate the phase drift from the midpoint rotation formula. Design a refinement diagnostic that distinguishes a time-discretisation bias from a field calibration error.

▼ **Compare the exact and discrete angles**

The exact unsigned angle per step is

$$h = 2\pi/20 = 0.31416.$$

The discrete one is

$$2\arctan(h/2) = 0.31161.$$

The accumulated deficit over 1,000 steps is about 2.55 rad, roughly 0.41 turn. The method can keep every track point on a closed momentum circle yet miss the pickup electrode by a large fraction of a revolution. Coordinate radius may depend on how positions were staggered, so perform the angular comparison on momentum first.

Double the number of steps at fixed physical end time and take the difference in unwrapped phases. The leading angular error $h^3/12$ per step becomes a quadratic global error, so halving the step should reduce it by about four. A calibration error in $qB/(\gamma m)$ persists under time refinement. Track both the expected frequency and the convergence of the time series; comparing the coordinates at a single late time can confuse complete turns with near agreement.

Cambridge Electromagnetism · four-vectors, fields and charged matter

Master equations, detailed balance and probability currents

Statistical mechanics from microscopic counting to material response

A two-state detector can switch at random while displaying reproducible occupation fractions and relaxation times. A three-state detector can maintain exactly constant occupancies while carrying a persistent directed current. Derive the probability dynamics, infer transition rates from timed records and decide what a stationary histogram can and cannot reveal about equilibrium.

From short-time transition probabilities to a generator

Take discrete states i resolved by an instrument. In a short time dt, conditional on current state i, let the probability of one transition $i \to j$ be $w_{ij}dt + o(dt)$. The probability of remaining in i is one minus the sum of departures to this order. Assume time-independent rates and the Markov property, meaning the conditional law of future transitions is determined by the current state. Sum gains and losses to derive

$$\dot{p}_i = \sum_{j \neq i} (p_j w_{ji} - p_i w_{ij}),$$

$$\sum_i \dot{p}_i = 0.$$

Each gain in one state is a loss in another, so total probability is conserved. In a column-vector convention $\dot{\mathbf{p}} = L\mathbf{p}$, off-diagonal entry $L_{ij} = w_{ji}$ and columns sum to zero. In a row-vector convention transpose the matrix. This is not cosmetic in a numerical calculation: using a matrix assembled for column probabilities while evolving rows can lose normalisation or give an incorrect stationary distribution.

For a given state the exit rate is

$$r_i = \sum_{j \neq i} w_{ij}.$$

Its survival probability obeys $\dot{S}_i = -r_i S_i$, so $S_i(t) = e^{-r_i t}$. Conditional on exiting, the destination has probability w_{ij}/r_i. This permits direct event-driven simulation: draw an exponential waiting time, choose a destination with these probabilities, advance the clock and repeat. Sampling one transition per arbitrary fixed time step generally distorts waiting-time statistics unless a controlled approximation or exact finite-time propagator is used.

Determine two rates from one occupancy and one time

For a two-state system name the rates $a = w_{01}$, $b = w_{10}$. Eliminating $p_0 = 1 - p_1$ gives

$$\dot{p}_1 = a - (a + b)p_1.$$

Its stationary value $p_* = a/(a + b)$ is independent of the common rate scale, and deviations decay with time constant $\tau = 1/(a + b)$. Thus one stationary fraction gives a ratio, one relaxation time gives a sum, and the two together identify both rates.

A location sensor records an upper-state fraction 0.30 and relaxation time 0.20 seconds. Then $a = 1.50$ and $b = 3.50$ per second. Mean waiting times are $1/a = 0.667$ s in the lower state and $1/b = 0.286$ s in the

upper one. Neither equals the population relaxation time of 0.20 s. A transition rate inferred by taking the reciprocal of the average of these two waiting times estimates a different quantity.

Define the indicator $X = 1$ in state one and zero otherwise. Conditional probabilities from the same linear solution give

$$\langle [X(t) - p_*][X(0) - p_*] \rangle = p_*(1 - p_*) e^{-|t|/\tau}.$$

Consequently a normalised equilibrium autocorrelation determines τ without needing a deliberate preparation away from stationarity. A bandlimited detector may miss short visits, reducing counts and distorting waiting-time distributions even if the mean occupation looks plausible. Compare time sampling with both exit times as well as the relaxation time, or model the measured signal with an explicit observation process.

Detailed balance is stronger than stationarity

Let π_i denote a stationary distribution. The net probability current on an oriented edge is $J_{ij} = p_i w_{ij} - p_j w_{ji}$. Stationarity requires only that currents sum to zero at each vertex. Detailed balance requires $\pi_i w_{ij} = \pi_j w_{ji}$ on every edge individually. For a system in equilibrium with one heat bath, with states even under time reversal and energies E_i, thermal rates satisfy $w_{ij}/w_{ji} = e^{-\beta(E_j - E_i)}$ for nondegenerate states. If states collect multiple microstates, degeneracies or coarse-grained free energies enter their equilibrium weights.

For the previous two-state process at equilibrium with equal state degeneracies, $p_*/(1 - p_*) = 3/7$ implies

$$(E_1 - E_0)/(k_B T) = \log(7/3) = 0.847.$$

Multiplying both rates by ten preserves the energy inference and speeds relaxation tenfold. Thermodynamic weights therefore do not determine the kinetic barrier. A physical sample may change its rate scale with solvent viscosity or transition geometry without changing the equilibrium ratio.

A convenient test on a connected graph with reciprocal edges is the product of rate ratios around each closed cycle. With detailed balance all stationary factors telescope to one. Conversely, if the product is one on every independent loop, consistent state weights can be reconstructed from any reference state by integrating ratios along a path. Path independence makes these weights a detailed-balance stationary distribution after normalisation. A missing reverse edge or unobserved channel requires additional care; a detector that cannot resolve two routes can hide directed currents.

A steady nonequilibrium system with the same state histogram

Consider three equally named sites arranged in a ring, with clockwise rate k_+ and anticlockwise rate k_- on every link. Symmetry gives stationary occupation $1/3$ at each site for all positive rates. Each oriented clockwise edge carries current $J = (k_+ - k_-)/3$. The cycle affinity in dimensionless units is $\mathcal{A} = \log[(k_+/k_-)^3]$. For $k_+ = 4$, $k_- = 1$ per second, each edge transports net probability one per second, and $\mathcal{A} = 3 \log 4 = 4.159$. A perfectly flat histogram conceals the directed operation.

For Markov jumps with resolved reversible edges and their conventional time reversals, the total stochastic entropy production rate is

$$\dot{S}_{\mathrm{tot}} = \frac{k_B}{2}\sum_{ij}(p_i w_{ij} - p_j w_{ji})\log\frac{p_i w_{ij}}{p_j w_{ji}} \geq 0.$$

Each term is nonnegative under the product $(x-y)\log(x/y)$; the factor of one half prevents counting both directions separately. For the stationary ring this gives $k_B J\mathcal{A} = 4.159 k_B$ per second. An actual heat dissipation rate needs a physical local detailed-balance model identifying environments, work and entropy flows. A collection of counted transitions without that calibration can still yield evidence for broken time-reversal symmetry at its level of resolution.

Stationary Shannon entropy $-k_B\sum p_i \log p_i$ is constant in this ring. The positive total entropy production is compatible with that constancy because an environment sustains the directed current and carries compensating entropy. Equating production with just the derivative of stationary state entropy would predict zero for every driven stationary process.

Why a distribution relaxes toward a stationary law

For a finite irreducible time-independent chain let $\pi_i > 0$ be its unique stationary distribution. The relative entropy $D(p\|\pi) = \sum p_i \log(p_i/\pi_i)$ is nonnegative. To see its contraction over a finite observation interval, write the propagated probability as a stochastic mixture of earlier probabilities and apply the log-sum inequality to each output state. Summing gives

$$D(Pp\|P\pi) \leq D(p\|\pi).$$

Since $P\pi = \pi$, the relative entropy to stationarity cannot grow.

For a detailed-balance chain, substituting the master equation into this derivative gives $-k_B dD/dt = \dot{S}_{\mathrm{tot}}$ in the simple closed heat-bath model. For a driven chain the system may remain at π with constant relative entropy while producing positive steady entropy. Relaxation toward stationarity and dissipation required to maintain that stationary state are distinct measurable quantities.

Compare relaxation inferred from a time trace with the thermodynamic irreversibility experiment

Infer generator entries from residence times

An ideal continuously observed Markov record contains successive states and dwell lengths. Let T_i be accumulated time in state i and N_{ij} the number of resolved $i \to j$ jumps. Its log likelihood contains $\sum_{ij} N_{ij}\log w_{ij} - \sum_i T_i \sum_{j\neq i} w_{ij}$. Differentiation gives

$$\widehat{w}_{ij} = N_{ij}/T_i.$$

Dividing by total experiment time instead estimates the gross stationary flux $\pi_i w_{ij}$, not a conditional transition rate.

For a nearly stationary two-state record with $T_0 = 70$ s, $T_1 = 30$ s and 105 transitions in each direction the estimates are 1.50 and 3.50 per second. Counting uncertainty for each reasonably large Poisson-like count is roughly the inverse square root of the count in relative terms; precise interval estimates condition correctly on

exposures and the observation process. Equality of opposite transition counts in a long two-state record is essentially forced by alternation. It cannot by itself establish a thermal energy gap or exclude multiple hidden driven pathways.

Turn a barrier diagram into two independently measurable rates

A thermally activated transition between two metastable configurations often has a forward rate $a = \nu_A e^{-(G^\ddagger - G_A)/(k_B T)}$ and reverse rate $b = \nu_B e^{-(G^\ddagger - G_B)/(k_B T)}$, where $G^\ddagger$ denotes a dividing-surface barrier under a chosen model. If the same transition-state construction applies in both directions with consistent partition factors, their ratio reproduces the equilibrium free-energy difference. The overall barrier controls how rapidly equilibrium is approached. Changing a barrier while leaving both wells unchanged can alter both kinetics by many orders of magnitude without changing stationary occupancy.

As an illustration let the higher-energy state cost $2k_BT$, and put a transition configuration $10k_BT$ above the lower state. With a common attempt rate of 10^6 per second, the upward rate is 45.4 per second and the downward rate 335 per second. The stationary upper fraction is approximately 0.119 and the relaxation time 2.63 ms. Lower the barrier by $2k_BT$ without changing the two state energies. Both rates increase by $e^2 = 7.39$, while occupancy is unchanged. A detector recording only the equilibrium level histogram could miss the intervention entirely.

For a temperature scan, a straight line in the log rate against inverse temperature is evidence for an approximately constant activation energy and prefactor over that interval. A free-energy barrier includes entropic as well as enthalpic contributions; its numerical value can vary with temperature even if the microscopic potential does not. A change of dominant pathway can bend the apparent Arrhenius plot. Measure forward and reverse rates independently, then compare their ratio with an independently determined equilibrium population. A barrier fit explaining one rate but violating the equilibrium ratio usually indicates inconsistent state definitions, missed short visits or an unmodelled pathway.

A useful control is a weak static bias that lowers one well relative to the other. The equilibrium log-odds shift measures the difference between the state couplings to the bias. The individual log-rate shifts additionally report the related dividing-surface response. If the state-occupation ratio moves as expected but only one rate changes appreciably, the effective transition configuration may be much closer in bias response to the other state. Such statements depend on the activated single-path interpretation; the response slopes themselves are more directly observed than an inferred barrier shape.

Use uniformisation to propagate a generator without negative probabilities

The formal finite-time propagator of a time-independent jump process is $P(t) = e^{Lt}$. A forward Euler update $p(t + \Delta t) = (I + L\Delta t)p(t)$ conserves total probability if the generator is assembled correctly, but a large step can yield negative diagonal probabilities. An exact algebraic representation starts with a rate q at least as large as every total exit rate. Then $K = I + L/q$ is a stochastic matrix in the same column convention. Write $L = q(K - I)$; since K and the identity commute,

$$e^{Lt} = e^{-qt} \sum_{n=0}^{\infty} \frac{(qt)^n}{n!} K^n.$$

This can be read as drawing a Poisson number of opportunities at rate q, followed by n transitions of K. At a lightly connected state some opportunities become self transitions. The construction leaves real waiting-time statistics correct while eliminating negative transition probabilities. Its truncation error in total probability is bounded by the omitted Poisson tail because every matrix power is stochastic.

For a two-state model with rates two and six per second choose $q = 6$ per second. Starting in the lower state, the first opportunity after some waiting time moves upward with probability one third and otherwise does nothing. In the upper state the next opportunity moves downward with probability one. At 0.25 s the Poisson mean is 1.5. Retaining terms through order eight leaves a tail below 3×10^{-5}. Compare the resulting upper probability against $(2/8)[1 - e^{-8(0.25)}] = 0.2162$. A small discrepancy due solely to truncation should shrink with the computable tail probability.

The representation also suggests a measured-data sanity check. A fitted generator with negative off-diagonal rates cannot yield a valid K at any positive q. A snapshot transition matrix can itself be nonnegative and stochastic but need not arise from a time-homogeneous continuous-time generator. In the two-state case its nontrivial eigenvalue must be $e^{-(a+b)t} > 0$. A measured negative eigenvalue rules out that simple generator model under perfect sampling, though finite counts and hidden periodic driving can distort empirical estimates. Taking a matrix logarithm blindly may return complex or negative rates rather than a physical kinetic explanation.

Population relaxation and a first-arrival experiment answer different questions

A first-passage measurement stops its clock when a target is first reached, even if that target would subsequently be left many times. To calculate its mean, treat the target as absorbing. For any other state i, wait on average $1/r_i$, then add the mean remaining time from the next state weighted by its conditional transition probability. The backward equations are

$$-1 = \sum_j w_{ij}(t_j - t_i),$$

with $t_{\mathrm{target}} = 0$. Their solution depends on possible returns before the first successful passage. It generally differs from an inverse equilibrium occupation or from a slowest relaxation time.

Consider a three-site linear chain $0 \leftrightarrow 1 \leftrightarrow 2$ with equal rate k on each adjacent directed edge, starting at zero and stopping at two. From zero one must reach one after mean $1/k$. From one the mean wait is $1/(2k)$, then there is probability one half of returning to zero and one half of finishing. Solving $t_0 = 1/k + t_1$, $t_1 = 1/(2k) + t_0/2$ gives

$$t_0 = 3/k.$$

Ignoring returns would yield $1.5/k$, exactly half the correct first-passage mean.

The unmodified closed chain has equal stationary weights, but that fact alone says nothing about the numerical value of $3/k$. Speeding every link by ten leaves occupancies unchanged and reduces passage times tenfold. A detector that registers only visits longer than a threshold makes a brief first arrival invisible and can report even longer apparent passage times. Calibrate detector dead time and distinguish entrance to the target from persistence within it.

For a larger network, solve the sparse backward linear system and compare against many event-driven paths. Check boundary initial conditions by starting simulations exactly at the absorbing target: the first-passage time must then be zero under that convention. Some experiments instead ask for a first return after leaving, which requires a different initial condition. A zero-lag event in one definition is a physically different measurement in the other.

Estimate a cycle affinity with finite counting statistics

A clockwise ring trajectory contains gross counts in both directions on each edge. Estimating each conditional rate as count divided by origin-state exposure, the product of clockwise-to-anticlockwise rates around a complete closed loop cancels the exposures. The estimated cycle affinity is the sum of log forward counts minus log reverse counts, assuming every required count is nonzero and rates were stationary during the record. This is a useful cross-check independent of how the total time was partitioned among states.

If all three edges register roughly 120 clockwise and 80 anticlockwise events, the estimated affinity is

$$3\log(1.5) = 1.216.$$

For large independent Poisson-like gross counts the approximate variance of this log product is $3(1/120 + 1/80) = 0.0625$, giving standard error 0.25 as a first estimate. Actual path constraints couple counts; uncertainty can be assessed from independent records or a complete stochastic likelihood. The point estimate alone does not tell whether an observed directionality is statistically distinguishable from zero.

One missing reverse event makes the unregularised logarithm infinite, but a finite record cannot establish an infinite physical affinity. Report a confidence bound based on the exposure and upper limit on the unseen reverse rate. Extend the acquisition or change sampling speed to exclude missed short-lived reverse states. Numerical addition of an arbitrary count can help regularise a plotting routine but is not a substitute for a specified prior or measurement model when claiming physical entropy production.

A driven occupation loop can carry charge without changing concentrations

Place three sites in a cycle and drive their transition rates by two controls, for example energies and barriers. Even if each instantaneous frozen generator satisfies detailed balance with a well-defined equilibrium distribution, the physical occupations need not equal that moving equilibrium distribution at each moment. A slow lag can produce a small edge current at every instant and a finite integrated current over a cycle. The current through an edge is time-integrated transfer; it can be nonzero while all occupations return to their initial values after a complete drive period.

A simple piecewise illustration avoids subtleties of adiabatic geometric pumping. In stage one, connect only sites zero and one while strongly favouring one. In stage two connect only one and two while favouring two. In stage three connect only two and zero while favouring zero. Begin with nearly all probability at zero and wait long enough for each enabled pair to relax. Almost the entire population travels once around the loop per cycle. No site accumulates probability over repeated complete cycles, but each edge registers about one forward crossing net per cycle. External control changes the energy and barrier landscape and supplies the required work.

If each programmed transfer succeeds with probability 0.95, it is tempting to announce a current of 0.95^3 per cycle. This is the probability of one fully successful trajectory starting on the prescribed site in a simplified no-return preparation, not generally the long-run edge current of a periodically driven ensemble. Failed transfers leave probability on other sites in later stages. The correct periodic distribution is the unit-eigenvalue vector of the product of the three finite-time propagators. Compute edge currents during each stage using its evolving occupations and integrate, then compare a complete cycle after transients have decayed.

This construction supplies a statistical check on experimental counting. Over a complete cycle the integrated net inflow at a site equals its change in occupancy probability. In a periodically stationary state that change is zero, so all three consistently oriented integrated edge currents must agree. A disagreement between measured edges can reveal a missed transition channel, a start-up transient or incorrect gating. Agreement does not prove absence of drive or detailed balance. A torque-free stationary ring and a three-stage external pump can share nearly the same time-averaged site histogram yet differ clearly in time-resolved edge currents.

Finally compare the work and heat bookkeeping with the master-equation entropy-production expression. During an external change of a site's energy without a transition, the state's energy changes by work. During a transition at frozen controls, assigning heat requires the energy difference and the relevant bath. Lumping an entire driven cycle into a single time-averaged generator can conceal those distinctions and underestimate dissipation.

The ordered protocol can also be run backwards. If the three stages are reversed while preparations are equilibrated in each corresponding periodic state, the preferential current direction should reverse in the ideal transfer picture. A fixed instrumental preference for counting neighbouring sites in one direction would not generally reverse. This distinguishes controlled transport from a direction-dependent detection efficiency. Do not confuse this observed current reversal with exact microscopic time reversal of each individual stochastic trajectory: the reversed experiment has its own ensemble of random jumps. Comparing path probabilities, rather than pairing individual paths as though their noise histories were reproducible, is the appropriate starting point for entropy-production and fluctuation-relation tests. The distinction becomes particularly visible when a stage transfers almost all the probability in one direction while a few rare trajectories travel against the applied protocol.

Distinguish stationarity from detailed balance

Keep the three populations identical while varying the directed transition rates. Identify the observable that reveals time-reversal asymmetry.

Open this experiment online

Try the calculation

Clockwise bias 0.5

Total jump rate (s⁻¹) 8

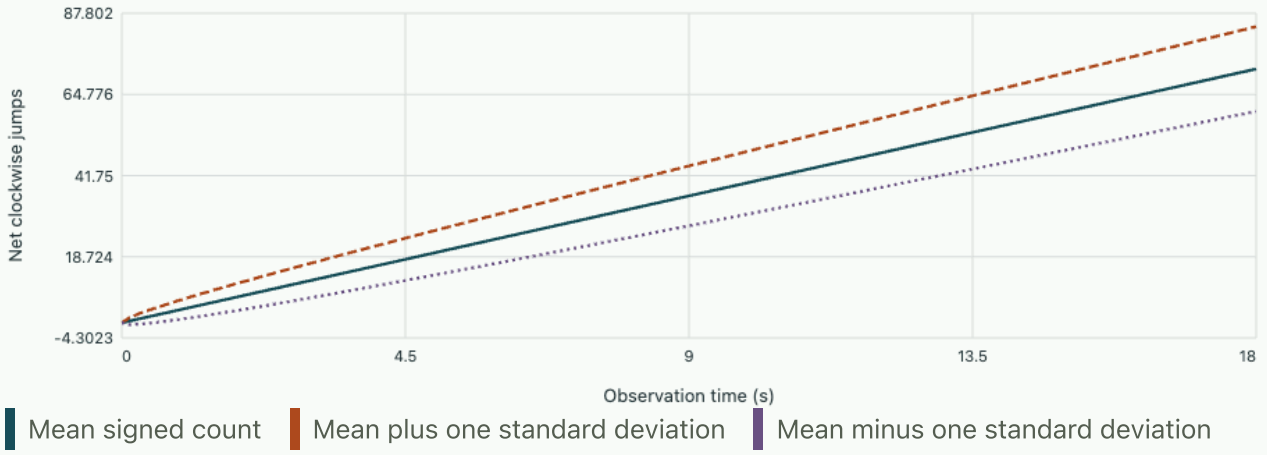

Every stationary population is 1/3. Net clockwise edge current 1.3333 s⁻¹. Entropy production 4.3944 kB/s.

Tutorial · Recover a generator from snapshots rather than full events

At sampling interval 0.10 s, a two-state record yields $P_{01} = 0.12$ and $P_{10} = 0.28$. Infer the continuous-time rates, and compare with directly dividing each transition probability by 0.10 s.

▼ **Use the exact finite-interval propagator**

The exact transition probabilities are $P_{01} = p_*[1 - e^{-(a+b)\Delta t}]$ and $P_{10} = (1 - p_*)[1 - e^{-(a+b)\Delta t}]$. Their sum is 0.40, so $a + b = -\log(0.60)/0.10 = 5.108$ per second. The stationary upper fraction is

$$0.12/0.40 = 0.30.$$

Thus $a = 1.532$ and $b = 3.576$ per second. Dividing by the interval gives only 1.2 and 2.8 because return transitions inside a sampling interval can erase evidence of the intermediate visit.

Tutorial · Can measured rates arise from one equilibrium potential?

For a three-state triangle suppose clockwise to counterclockwise conditional rate ratios are 2, 3 and 0.25 around the three distinct edges. Could these be equilibrium weights in a common bath? If a single uncertain ratio were corrected, what would it have to be on the third link?

▼ **Apply the closed-loop consistency condition**

The product is $2 \cdot 3 \cdot 0.25 = 1.5$, giving nonzero loop affinity $\log 1.5 = 0.405$. No consistent detailed-balance set of three state weights exists for these exact ratios. The third ratio would have to be $1/6$, giving a telescoping unit product. If counts are noisy, propagate uncertainty in the log ratios or compare a constrained detailed-balance likelihood against an unconstrained model. A value barely inconsistent with unity within one sparse record is not enough to establish driving.

Tutorial · Detect a hidden state with the same observed level

A detector reports two visible levels, but its lower level combines two unresolved substates. The measured lower-state survival probability is $0.5e^{-t} + 0.5e^{-9t}$ with time in seconds. Calculate the initial hazard and its long-time limit. Can a two-state Markov fit reproduce both?

▼ **Differentiate survival before interpreting the tail**

The hazard is minus the derivative of $\log S$. At zero time it is the initial weighted mean of rates, 5 per second. At long times the rapidly departing component has almost disappeared and the hazard tends to 1 per second. A single Markov exit rate would give a constant hazard for all dwell ages. The observed model cannot reproduce both limits with one rate.

A changing apparent hazard may also result from limited detector temporal resolution or variable preparation of the substates. Repeat with shorter acquisition intervals and condition the dwell distribution on the preceding visible transition or an independent preparation marker. A fitted sum of exponentials identifies a need for a richer observation or state model; it does not uniquely identify microscopic structures absent in the measurements.

Cambridge Kinetic Theory · stochastic processes and approach to equilibrium

Transport coefficients from equilibrium fluctuations

Statistical mechanics from microscopic counting to material response

How can a fluid in equilibrium reveal the diffusion constant that would govern a spreading drop, or the viscosity that would resist an applied shear? Derive an exact connection between displacements and velocity correlations, generalise it to stress and heat currents and test how a finite record can bias each number. The assumptions behind convergence matter as much as the prefactor.

Integrate the velocity correlation before taking a long-time limit

Observe one Cartesian coordinate of a tagged particle in a homogeneous equilibrium fluid with zero mean drift. The displacement over time t is the integral of its velocity. Square that integral, average over the equilibrium ensemble and use stationarity to replace the two-time expectation by $C_v(t_1 - t_2)$. For a scalar component in a time-reversal-symmetric equilibrium state this correlation is even. Integrate over the two triangles of the time square:

$$\langle [x(t) - x(0)]^2 \rangle = 2 \int_0^t (t - s) C_v(s)\, ds,$$

$$\frac{1}{2} \frac{d}{dt} \langle \Delta x^2(t) \rangle = \int_0^t C_v(s)\, ds.$$

If the correlation is integrable and boundary effects remain absent, the long-time slope yields a diffusion coefficient $D = \int_0^\infty C_v(s) ds$. In d isotropic dimensions one may use $D = (1/d) \int \langle \mathbf{v}(s) \cdot \mathbf{v}(0) \rangle ds$ or the mean-square-displacement slope divided by $2d$. Using the full vector dot product without division by dimension gives an overestimate. Directional diffusion in an anisotropic material requires tensor components rather than a forced scalar average.

The derivation is kinematic and does not assume a Markovian motion. A fluid with velocity memory or an oscillating solid still obeys the finite-time relation. Those systems need not exhibit the same long-time proportionality to time. A bound oscillator can have persistent or alternating velocity correlations and a bounded displacement. An observation window in which the mean square grows over a fraction of an oscillation is insufficient to claim diffusion.

A solvable benchmark from ballistic to diffusive motion

For one coordinate take the underdamped Langevin equation $m\dot{v} = -\zeta v + \xi$, with zero-mean white noise satisfying $\langle \xi(t) \xi(t') \rangle = 2 k_B T \zeta\, \delta(t - t')$. In equilibrium $\langle v^2 \rangle = k_B T / m$. Condition on $v(0)$ and integrate for positive time. The later random force has no correlation with the initial velocity, giving $C_v = (k_B T / m) e^{-t/\tau_v}$, where $\tau_v = m/\zeta$. Its area and the resulting displacement are

$$D = \frac{k_B T}{\zeta},$$

$$\langle \Delta x^2(t) \rangle = 2D[t - \tau_v (1 - e^{-t/\tau_v})].$$

At $t \ll \tau_v$, the constant and linear terms inside the brackets cancel and displacement is ballistic:

$$\langle \Delta x^2 \rangle \simeq (k_B T / m) t^2.$$

At $t \gg \tau_v$ the slope approaches $2D$. Fitting a line through the origin at times only a few relaxation times overestimates the influence of the negative intercept in this long-time expression and underestimates D. Comparisons at multiple lags distinguish crossover curvature from a true anomalous power law.

For an inert sphere of radius 0.50 micrometres and density 2000 kg per cubic metre in water with $\eta = 1.00$ mPa s at 298 K, its bare mass is 1.05×10^{-15} kg and Stokes drag 9.42×10^{-9} kg/s. The estimate $\tau_v = 0.111$ microseconds is far faster than ordinary video tracking. $D = 0.437$ square micrometres per second and the one-coordinate rms displacement over 10 ms is 0.0935 micrometres. Hydrodynamic added mass and memory affect very short-time behaviour; the simple exponential is an analytical reference, not a claim that a water suspension remains Markovian down to arbitrary times.

Relate a weak imposed force to equilibrium motion

In the same soluble model a small constant force F turns the mean velocity into $F(1 - e^{-t/\tau_v})/\zeta$, so the mobility $b = \lim \langle v \rangle / F = 1/\zeta$. Together with equilibrium diffusion this gives the Einstein relation $D = k_B T b$. A broader equilibrium argument considers a dilute suspension in a weak conservative potential $U(x)$. Its current is

$$j = -D\,\partial_x n + bFn,$$

with $F = -\partial_x U$. The Boltzmann distribution $n \propto e^{-U/k_B T}$ must carry zero current; substitution gives the same relation without assuming exponential velocity relaxation.

Linear response connects the full weak-force impulse response to an equilibrium correlation. For a perturbation coupled as $-F(t)x$, classical equilibrium response of velocity has causal kernel $C_v(t)/(k_B T)$ for $t > 0$ under the usual microreversibility assumptions. Integrating the response kernel for a constant weak force gives

$$b = (1/k_B T) \int_0^\infty C_v(t) dt,$$

which recovers the kinematic diffusion relation. Strong forcing, an active bath or a nonequilibrium steady state need not obey this simple equilibrium identity.

A direct check compares independently measured weak sedimentation or calibrated drift with freely diffusing trajectories. For the previous bead a force 10 fN would produce a mean drift about 1.06 micrometres per second in the ideal Stokes regime. During a 10 ms frame its mean drift is about 0.011 micrometres, much smaller than rms Brownian motion, so many frames or many particles are needed. A single short jagged record can appear to move against the force. Use equal positive and negative applied forces to cancel constant stage drift in the estimated mobility.

Work through diffusion and mobility in the random-walk chapter

Stress and heat current carry their own normalisations

In a homogeneous isotropic laboratory fluid, define P_{xy} as a spatially averaged off-diagonal microscopic stress, with units Pa. Its equilibrium mean is zero in an isotropic state without imposed shear. A weak velocity

gradient couples to momentum transport. The Green–Kubo expression for dynamic shear viscosity is

$$\eta = \frac{V}{k_B T} \int_0^\infty \langle P_{xy}(t) P_{xy}(0) \rangle \, dt.$$

If a program outputs an extensive stress moment $J_{xy} = V P_{xy}$, the prefactor is $1/(V k_B T)$. Multiplying an already extensive current by volume again is a common orders-of-magnitude error. Check dimensions: pressure squared times time times volume divided by energy gives pascal seconds. Kinetic and interaction contributions to microscopic stress must both be included according to the model; a calculation with ideal kinetic flux alone need not reproduce the viscosity of a dense fluid.

Similarly for a component of total heat current $J_{Q,x}$ with units energy times length per time, the thermal conductivity in an isotropic equilibrium single-component system is

$$\kappa = \frac{1}{V k_B T^2} \int_0^\infty \langle J_{Q,x}(t) J_{Q,x}(0) \rangle \, dt.$$

Here $k_B T^2$ means k_B times T^2. If a heat-flux density is used instead, the volume prefactor moves as for stress. Heat current must exclude energy carried by a net mass flow as appropriate to the chosen thermodynamic constraint. In mixtures several conserved species currents can couple, and the measured conductivity at zero mass flux can differ from one unprojected current coefficient. The meaning of the measured flux and the imposed condition must be specified before comparing coefficients.

For an illustrative volume 10^{-24} cubic metres at 300 K, assume one stress autocorrelation $C_P(t) = (2.0 \times 10^6 \ \mathrm{Pa})^2 e^{-t/(2.0 \ \mathrm{ps})}$. Its area is 8.0 square pascal seconds. Multiplication by $V/(k_B T)$ gives viscosity 1.93×10^{-3} Pa s. Doubling the simulated volume of an otherwise similar homogeneous fluid generally reduces the variance of its volume-averaged stress, keeping the physical transport coefficient intensive. Holding the same numerical stress variance while merely changing the prefactor would describe a different equilibrium fluctuation model.

Find a correlation integral's useful time window

An estimate at lag s averages over pairs of samples separated by that lag. In a record of duration T_{rec}, the number of available pairs decreases with lag. Overlapping time origins strongly correlate the estimates, so plotting one error bar per raw pair exaggerates precision. Integrating an estimated correlation to the maximum possible lag frequently accumulates late random noise rather than useful signal. Conversely stopping too early omits a physical tail.

For a positive single-exponential reference, integrating to three correlation times retains $1 - e^{-3} = 95.0$ percent of the area. Five times retain 99.3 percent. A correlation containing a small, much slower component can look essentially flat on the main decay scale while contributing appreciably to the area. Divide a long run into blocks, compare running integrals and check several window lengths. Report the variation between independent blocks and any explicit tail model used beyond the measured interval.

At short lags, finite exposure averages correlated positions or velocities. Uncorrelated localisation errors in position increase nonzero-lag displacement variance by a constant $2 s_x^2$ per coordinate. Numerically differentiating those same noisy positions creates large artificial velocity variance and negative correlations between adjacent increments. An apparent negative velocity lobe can therefore be physical backscattering or

a measurement artefact. Comparing a raw displacement fit with the velocity integral after independently characterising detector noise is more informative than differentiating first and interpreting every feature as fluid memory.

A forced shear wave gives a second viscosity measurement

An equilibrium stress correlation estimates viscosity without imposing any flow. A separate driven measurement can test that coefficient against a directly observed penetration depth. Take a Newtonian fluid filling $y > 0$, driven by a laterally oscillating plate at $y = 0$ with speed $U_0 \cos \omega t$. Neglect boundaries at large y, compressibility and nonlinear convective effects in this spatially uniform tangential shear. The momentum equation reduces to

$$\partial_t u = \nu \partial_y^2 u,$$

with kinematic viscosity $\nu = \eta/\rho$.

Seek $u = \mathrm{Re}[U_0 e^{-\lambda y - i\omega t}]$ and choose the square root with positive real part. Then $\lambda = (1 - i)/\delta$,

$$\delta = \sqrt{2\nu/\omega}.$$

Both amplitude decay and phase lag carry information: at distance one penetration depth the amplitude is e^{-1} of the wall speed and the lag magnitude is one radian. A single complex velocity map can determine viscosity from either slope and use their agreement to test the simple model.

For a fluid of density 1000 kg per cubic metre and measured penetration depth 0.126 mm at 20 Hz, $\nu = \omega \delta^2/2 \simeq 1.00 \times 10^{-6}$ square metres per second and $\eta \simeq 1.00$ mPa s. A stress-fluctuation integral yielding a value twice as large would predict a penetration depth larger by $\sqrt{2}$, approximately 0.178 mm. Such a comparison can discriminate a missing component or volume factor in the equilibrium calculation if density, temperature and fluid composition match.

At high enough frequency the fluid may have a complex frequency-dependent viscosity rather than the zero-frequency coefficient from a Green–Kubo area. In linear response that generalisation is the one-sided Fourier transform of the equilibrium stress correlation with appropriate prefactors. A single-exponential stress memory gives a frequency-dependent response with a characteristic roll-off and phase lag. Comparing a high-frequency penetration measurement directly to a zero-frequency area without checking the timescale can therefore produce a genuine difference, not a statistical failure.

Remove a centre-of-mass mode only when the observable requires it

In a finite periodic particle simulation total momentum may be exactly conserved. If the initial centre of mass has a random velocity and no thermostat removes it, every particle's measured laboratory coordinate includes the common drift $V_{\mathrm{CM}} t$. Its mean-squared displacement gains a ballistic term $\langle V_{\mathrm{CM}}^2 \rangle t^2$. In a canonical ensemble the centre-of-mass velocity variance in one direction is $k_B T / M_{\mathrm{tot}}$. Though it falls with system size, at sufficiently long times its quadratic contribution exceeds the linear diffusive part.

For 1000 equal particles of mass m, an equilibrium centre-of-mass mode contributes one thousandth of the single-particle instantaneous velocity variance. In a long-time autocorrelation this small constant never integrates to a finite diffusion coefficient. Subtract centre-of-mass displacement or work in a zero-total-momentum ensemble before taking the long-time integral of an internal tracer's motion. In a real sample with externally imposed drift, decide whether the physical quantity sought is diffusion about that drift or absolute laboratory displacement; the subtraction has a physical meaning rather than merely a cosmetic improvement.

Conversely, indiscriminately subtracting the mean velocity of a tiny observation subregion separately at every time can remove genuine hydrodynamic long-wavelength motion. A strongly applied single-particle thermostat can alter current correlations and suppress transport even while maintaining the correct static temperature histogram. Check the thermostat strength and system size; compare an equilibrium-prepared weakly thermostatted trajectory with an energy-conserving production segment where appropriate. Maintaining the correct one-time distribution is not sufficient to preserve all dynamical correlation functions.

The distinction also matters when removing an estimated mean stress before integrating its fluctuations. The true shear mean may vanish by symmetry, but a short noisy record has a nonzero sample mean. Subtracting a slowly varying trend on the same timescale as a genuine stress tail can erase physical viscosity. Use independent samples, separately measured drifts and explicit sensitivity to the subtraction window. A persistent nonzero off-diagonal stress under nominal isotropic conditions may indicate incomplete equilibration, a deformed box or an unrelaxed solid rather than a fluid transport coefficient.

Heat and particle currents can share fluctuations

Consider two thermodynamic fluxes J_1, J_2 driven weakly by conjugate forces X_1, X_2, chosen so entropy production density is $J_1X_1 + J_2X_2$. Linear response writes $J_i = \sum_j L_{ij}X_j$. Near time-reversal-symmetric equilibrium, with flux parities chosen appropriately and no magnetic-field reversal complications, reciprocal relations give

$$L_{12} = L_{21}.$$

The symmetric part must be positive semidefinite so that entropy production is nonnegative for every small force. Thus $L_{11}, L_{22} \geq 0$ and $L_{12}^2 \leq L_{11}L_{22}$.

Suppose a thermal measurement enforces $J_1 = 0$ by allowing a compensating force X_1 to develop, as when a concentration or electrochemical gradient cancels a coupled particle flux. Solving the first row gives

$$X_1 = -L_{12}X_2/L_{11}.$$

Substitution into the second yields response coefficient $L_{22} - L_{21}L_{12}/L_{11}$. The measured constrained heat response can therefore be smaller than the coefficient at fixed $X_1 = 0$. The difference has a simple operational origin in the boundary condition enforced during the experiment.

Take consistent normalised coefficients $L_{11} = 4, L_{22} = 3, L_{12} = L_{21} = 2$. At zero first force the heat flux slope is three. At zero first flux it is $3 - 4/4 = 2$. Both are positive and consistent with the same equilibrium correlation matrix. Comparing a simulated heat-current autocorrelation corresponding to the unconstrained coefficient against a zero-particle-flux experiment requires the mixed-current correlation too. Apparent disagreement by a factor of one and a half can arise entirely from this distinction.

Green–Kubo cross-correlations must use properly defined conjugate fluxes and units. Redefining heat current by adding a constant multiple of particle current changes the matrix entries but must leave a physically defined zero-particle-flux heat measurement invariant when the forces are transformed consistently. This reference invariance is analogous to subtracting a per-particle energy reference in grand-canonical covariance. It provides a powerful test of a multicomponent transport implementation without relying on any empirical transport value.

A lattice trajectory provides an exactly known correlation integral

For a one-dimensional nearest-neighbour walk with steps of length a every interval Δt, suppose successive step directions have correlation $\langle s_{j+k} s_j \rangle = r^k$ with $|r| < 1$, and zero mean. Squaring the sum of n steps gives

$$a^2[n + 2\sum_{k=1}^{n-1}(n-k)r^k].$$

In the large-n limit divide by $2n\Delta t$ to obtain

$$D = a^2(1+r)/[2\Delta t(1-r)].$$

The uncorrelated random-walk value is recovered at $r = 0$.

With $a = 1$ micrometre and $\Delta t = 0.10$ s, uncorrelated diffusion is five square micrometres per second. Persistence $r = 0.5$ raises it to fifteen, while anticorrelation $r = -0.5$ lowers it to $5/3$. All three walks have the same step-size variance, so an estimate based only on one-frame displacement misses the entire correlation effect. The finite-lag autocorrelation sum and a long-time displacement fit should recover the same value on a long simulated record.

The discrete formula also clarifies the special zero-lag weight in numerical integration. An integral approximated from sampled symmetric correlations gives half weight to the zero lag on a one-sided trapezoidal rule. A careless double count of the instantaneous self term predicts twice the uncorrelated diffusion. Benchmark a numerical estimator with this exact walk before applying it to molecular data with unknown late tails and noisy velocity derivatives.

Measure an off-diagonal diffusion coefficient after rotating the sample

In an anisotropic sheet a tracer's two-dimensional displacement covariance need not align with the camera axes. Let its long-time principal diffusion coefficients be D_1, D_2 and the first principal axis make angle θ with laboratory x. In matrix form the laboratory tensor is $R\,\mathrm{diag}(D_1, D_2)R^{\mathsf{T}}$, with R the planar rotation. At sufficiently large lag in an effectively unbounded drift-free region, the displacement covariance matrix tends to $2Dt$. An off-diagonal covariance is a measurable correlation between displacement components even if the noise is independent in the material principal directions.

Take $D_1 = 2.0$ and $D_2 = 0.50$ square micrometres per second with $\theta = 30^\circ$. Laboratory components are $D_{xx} = 1.625$, $D_{yy} = 0.875$ and $D_{xy} = 0.650$, in the same units. At 0.10 s the cross covariance is 0.130

square micrometres. Rotating the physical sheet by ninety degrees interchanges the two diagonal values and reverses the sign of the off-diagonal one. An instrument-fixed camera readout artefact would not generally follow this material rotation.

To infer the principal axes from data, fit each covariance component versus lag using a common drift subtraction and independent position calibration. Diagonalise the symmetric fitted diffusion matrix. Positive semidefiniteness requires

$$D_{xy}^2 \leq D_{xx} D_{yy}.$$

A strong violation indicates incompatible estimates, unit conversion errors or statistical uncertainty, not a physically permissible strongly anisotropic diffusion tensor. Near-degenerate principal values make the orientation poorly determined even when the mean diffusion coefficient is precise. Report uncertainties on eigenvalues and orientation separately.

A finite field of view introduces a further bias. Fast particles along the high-diffusion direction leave the image sooner, and a tracking pipeline retaining only long trajectories can favour slow or fortuitously confined paths. Compare different fields of view and minimum track durations, or fit a measurement model incorporating censoring. A physical material anisotropy should persist after correcting these effects; apparent anisotropy aligned with the short direction of a rectangular observation window demands particular scrutiny.

A weak calibrated force along different axes supplies an independent check through the tensor Einstein relation under equilibrium conditions. If the diffusion tensor predicts off-axis mean drift, the mobility experiment should observe its sign and magnitude after reversing the force. A discrepancy could reveal active forcing, nonstationarity, an incorrect temperature or a mismatch between the mobility and displacement observables. It cannot be diagnosed from a scalar mean square alone.

Another rotational invariant is the trace of the diffusion tensor. In the example its value stays 2.50 square micrometres per second for every sample orientation, while the determinant stays one square micrometre to the fourth per second squared. A comparison of two camera orientations therefore provides numerical checks independent of the fitted angle. A drift mistakenly left in the displacement covariance can violate the expected linear dependence on lag, making the inferred effective eigenvalues grow with the selected fitting window. Subtract the ensemble mean displacement at each lag and record that mean separately as a flow measurement. Conversely do not subtract an individual best-fit line from every short trajectory without modelling the resulting loss of long-time fluctuations; that procedure can bias both diffusion coefficients downward.

Choose an integration window for a transport coefficient

Integrate an oscillatory velocity correlation over a finite record. Compare a short apparent plateau with the independently computed infinite-time result.

Open this experiment online

Try the calculation

Correlation decay time τ (ps) 1

Oscillation frequency Ω (rad/ps) 1

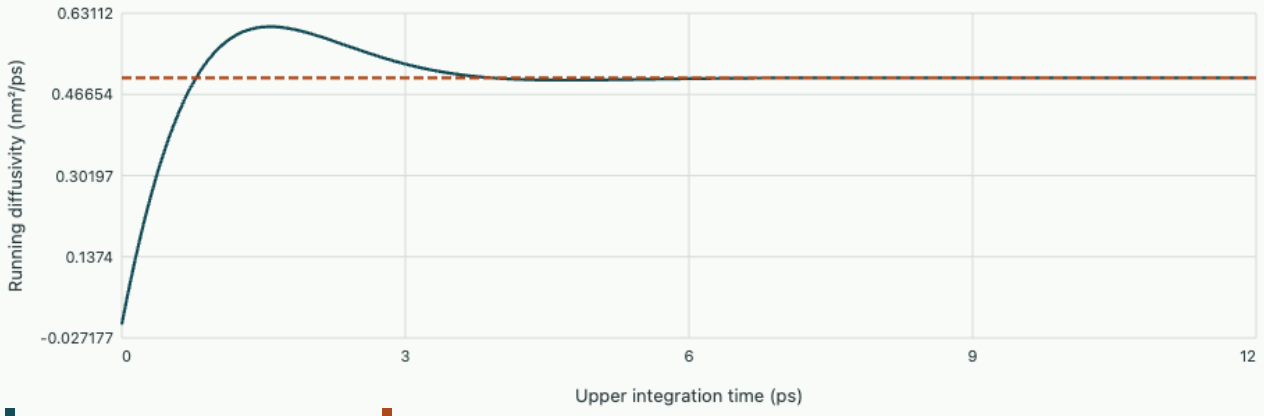

With C(0) = 1 nm²/ps² the exact 1D diffusion coefficient is 0.5 nm²/ps. At 2τ the integral gives 0.58969 nm²/ps.

Tutorial · Invert diffusion measurements at two lags

In a one-dimensional diffusive window a measured displacement variance is 0.024 square micrometres at 10 ms and 0.064 at 50 ms. Assume long-time Brownian displacements plus independent frame-localisation variance per position measurement. Estimate D and localisation precision.

▼ **Fit slope and intercept separately**

The variance slope is $(0.064 - 0.024)/(0.050 - 0.010) = 1.00$ square micrometre per second, so $D = 0.500$ square micrometres per second. The intercept is 0.014 square micrometres, equal to twice the single-frame localisation variance. The implied rms localisation error is $\sqrt{0.007} = 0.0837$ micrometres. Forcing the fit through the origin at the shortest lag would give $D = 1.20$, more than twice the inferred value.

Acquire more lags to test the linear model. A static drift contributes a term growing as t^2 if its mean displacement is not removed, while finite exposure can lower the physically measured short-lag variance. Two points always define a line and cannot independently validate either noise model.

Tutorial · Estimate a hidden slow tail in viscosity

A stress correlation is modelled as $A[e^{-t/\tau} + 0.02e^{-t/(40\tau)}]$. What fraction of its true area comes from the low-amplitude tail? How much does a cutoff at 5τ capture?

▼ **Weight a correlation by its duration**

The fast and slow areas are $A\tau$ and $0.80A\tau$. The slow contribution is 44.4 percent of total. Integration to five fast times captures $A\tau[(1-e^{-5})+0.80(1-e^{-5/40})]\simeq 1.087A\tau$, or about 60.4 percent. Visual disappearance of the dominant peak therefore does not establish convergence of the transport coefficient.

The slow contribution must be supported by independent long blocks or a physically motivated model. A fitted tiny constant baseline can generate arbitrarily large apparent area if extrapolated indefinitely. Subtracting a sample mean estimated from the same short record can conversely suppress a genuine slow mode. Compare records much longer than the proposed tail time before trusting the extrapolation.

Tutorial · Is an algebraic tail compatible with ordinary diffusion?

Suppose a measured long-time velocity correlation remains positive and scales as $t^{-\alpha}$. Determine when its Green–Kubo area exists and infer the displacement growth for $0<\alpha<1$ and for $\alpha=1$. Why would a finite periodic simulation require care?

▼ **Integrate the time-dependent slope before naming a diffusion coefficient**

For $\alpha>1$ the tail integral converges. For $0<\alpha<1$, the running area grows as $t^{1-\alpha}$, so displacement variance grows as $t^{2-\alpha}$ under the stated persistent positive tail. At $\alpha=1$, the running area grows logarithmically and the displacement variance is of order $t\log t$. Sign oscillations, finite-size termination or another late-time crossover can alter these conclusions.

A periodic box admits only discrete long-wavelength modes and may cut off a hydrodynamic tail. Repeating at larger sizes and comparing the range before signals traverse the box helps identify a bulk regime. Unwrap particle coordinates crossing periodic boundaries before calculating a displacement. A wrapped coordinate can falsely saturate and imitate a confined particle even in a freely diffusing periodic fluid.

Cambridge Kinetic Theory · correlations, response and transport

Approximate an eigenstate and check what it predicts

Use projection, orthogonalisation and energy stationarity to construct approximations. Compare them with phase-integral and explicitly propagating pictures of the same mathematical operators.

1. Eigenvalues, modes and spectral gaps
2. Householder reflections and practical QR
3. Arnoldi iteration and GMRES
4. How a moderately accurate mode can give a precise frequency
5. Estimating a cantilever frequency from energy
6. Recover a trajectory from a family of actions
7. Read a period from the area enclosed by an orbit
8. Why two finite-element wave solvers propagate the same mode at different speeds
9. Noether's theorem, boundary terms and measurable constants
10. Kepler orbits, the eccentricity vector and the loss of apsidal closure
11. Symplectic transformations, Poisson brackets and the geometry of measurements
12. Action, angle and the slowing clock at a separatrix
13. Canonical perturbations, small denominators and resonant capture

Eigenvalues, modes and spectral gaps

Numerical linear algebra

Estimate a dominant vibration or diffusion mode, explain the role of eigenvalue separation and use a residual to check a computed eigenpair.

FROM FIRST PRINCIPLES

Watch a power iteration select a mode

Start with a real symmetric matrix A and express a trial vector v_0 in its orthonormal eigenvectors u_j. Order the eigenvalues so $|\lambda_1| > |\lambda_2| \geq \cdots$ and suppose the starting component c_1 is nonzero. Applying A repeatedly multiplies each component by its own eigenvalue,

$$v_0 = \sum_j c_j u_j,$$

$$A^k v_0 = \lambda_1^k \left[c_1 u_1 + \sum_{j>1} c_j (\lambda_j/\lambda_1)^k u_j \right].$$

Normalizing after each application keeps vector lengths manageable without changing the relative size of the components. The largest competing component therefore shrinks relative to the first at a geometric rate. A scalar estimate of the eigenvalue comes from projecting the action back onto the trial vector,

$$\rho(v) = \frac{v^T A v}{v^T v}.$$

For an exact eigenvector the quotient equals its eigenvalue. Equal leading magnitudes can prevent selection of a single direction. A negative leading eigenvalue can make the displayed vector alternate sign even while its one-dimensional direction converges.

DERIVATION AND APPLICATIONS

Derive the eigenvalue shift produced by weak coupling

Two individually measured modes can move apart in frequency or stiffness when they are coupled. Model a real symmetric operator in an orthonormal basis by diagonal values α, β and off-diagonal coupling g. All three numbers must have the same units. A nonzero eigenvector exists only when the determinant of the shifted matrix vanishes,

$$(\alpha - \lambda)(\beta - \lambda) - g^2 = 0.$$

Set $m = (\alpha + \beta)/2$ and $d = (\alpha - \beta)/2$. The quadratic becomes $(\lambda - m)^2 = d^2 + g^2$, so the two eigenvalues and their separation are

$$\lambda_{\pm} = m \pm \sqrt{d^2 + g^2},$$
$$\lambda_+ - \lambda_- = 2\sqrt{d^2 + g^2}.$$

Even when the uncoupled values coincide, a nonzero coupling leaves a gap $2|g|$. For well-separated diagonal entries, assume $\alpha > \beta$ and expand the square root for $|g| \ll \alpha - \beta$. The largest eigenvalue has the shift

$$\lambda_+ - \alpha = \frac{g^2}{\sqrt{d^2 + g^2} + d} \simeq \frac{g^2}{\alpha - \beta}.$$

The rationalized form gives an accurate small shift without subtracting almost equal eigenvalues. The eigenvalue sum remains exactly $\alpha + \beta$ because the trace is unchanged. A mass-normalized vibration model typically returns squared angular frequencies, so any measured ordinary frequencies must first be converted to the variables of the operator.

WORKED EXAMPLE

Predict a shift away from a crossing

With $\alpha = 10$, $\beta = 6$ and $g = 0.5$, the upper shift is $0.25/(\sqrt{4.25} + 2) = 0.0616$; the leading approximation is 0.0625. The lower eigenvalue is 5.9384 and the upper is 10.0616. The opposite signs of the shifts preserve the trace of 16.

SECOND APPLICATION

Recover coupling and detuning from a measured gap

A second two-mode system has eigenvalues 19 and 25, giving a mean of 22 and half-gap of 3. An independent uncoupled calibration gives $|\alpha - \beta| = 4$, so $d^2 = 4$. The inferred coupling magnitude is

$$\sqrt{3^2 - 2^2} = \sqrt{5} = 2.24.$$

The eigenvalues do not determine the sign of g, which also changes when one basis vector is reversed.

What an eigenproblem measures

In an undamped mechanical model $M\ddot{u} + Ku = 0$, insert $u = v\cos(\omega t)$. Nonzero amplitudes require

$$Kv = \omega^2 Mv.$$

The normal frequencies are square roots of generalised eigenvalues and the mode shapes are determined only up to scale and sign. For symmetric positive definite M, substituting $w = M^{1/2}v$ gives the ordinary symmetric operator $M^{-1/2}KM^{-1/2}$. This step explains why mass weighting matters when comparing shapes from two different meshes.

For an ordinary symmetric matrix A with orthonormal eigenvectors, a normalised trial vector $v = \sum c_i u_i$ has Rayleigh quotient $\rho = v^T A v = \sum_i \lambda_i c_i^2$. Hence it lies between the smallest and largest eigenvalues. Near an

isolated mode its value can be much more accurate than the vector direction: shifting a small amount of amplitude into another orthogonal mode changes the quotient by a term proportional to the square of that amplitude.

Select a mode by repeated application

The power iteration repeatedly forms Av_k and normalises it. Before normalisation,

$$A^k v_0 = \sum c_i \lambda_i^k u_i.$$

If one eigenvalue has strictly largest magnitude and the initial vector overlaps its eigenspace, the ratio of any competing component to the leading one gains a factor $(\lambda_i/\lambda_1)^k$. Normalisation avoids overflow while preserving those ratios. A negative leading eigenvalue can cause sign changes between successive normalised vectors. Comparing their raw difference then misrepresents convergence.

For $A = \operatorname{diag}(4, 2)$ and a starting vector proportional to $(1, 1)^T$, the ratio of the second component to the first is $1/2$ after one multiplication and $1/4$ after two. A vector proportional to $(4, 1)^T$ has Rayleigh quotient $(4 \cdot 16 + 2)/17 = 66/17 = 3.88235$. After four steps the ratio is $1/16$, and the quotient is

$$1026/257 = 3.99222.$$

Check that the quotient error decreases roughly fourfold per additional step, although the directional error decreases only twofold.

Use shifts to find another part of the spectrum

Power iteration selects magnitude. To find a mode near a target σ, apply the inverse of $A - \sigma I$ implicitly by solving a linear system each iteration. The eigenvectors are the same and the eigenvalues are $1/(\lambda_i - \sigma)$. For eigenvalues 2 and 4 with $\sigma = 2.1$, the transformed magnitudes are 10 and about 0.526. The eigenvalue near 2 dominates for a suitable starter. A fixed shift permits factor reuse, but choosing it exactly equal to an eigenvalue makes the shifted solve singular.

Test a candidate pair by forming $r = Av - \rho v$ with $\|v\|_2 = 1$. In the orthonormal basis,

$$\|r\|^2 = \sum_i (\lambda_i - \rho)^2 c_i^2.$$

At least one eigenvalue must lie within $\|r\|$ of ρ. This is a numerical certificate for a symmetric problem. To certify a particular eigenvector, the gap to neighbouring eigenvalues also matters. At a repeated eigenvalue every normalised linear combination of its modes is valid. Comparing subspaces or projecting one span onto the other is then more meaningful than matching individual component signs.

Follow the whole spectrum by QR iteration

To find several eigenvalues of a small dense operator, factor $A_k - \mu I = QR$ and form $A_{k+1} = RQ + \mu I$. Orthogonality gives $A_{k+1} = Q^T A_k Q$, an exact similarity. Hence every step preserves eigenvalues, trace and

determinant while changing the coordinates used to represent the operator. This use of QR is different from solving a triangular least-squares system. Repeated unshifted steps amount to an orthogonalised power iteration on multiple directions. A shift near a separated eigenvalue can make the last coordinate nearly invariant and accelerate its separation.

For a symmetric matrix a preliminary sequence of orthogonal similarities reduces it to tridiagonal form. When the coupling between the last two coordinates is negligible relative to the desired accuracy, setting that coupling and its transpose to zero separates a scalar eigenproblem. This removal is called deflation. If the off-diagonal magnitude is e, setting the pair to zero is a symmetric perturbation of spectral norm $|e|$, giving a direct measure for that decision. For a general real nonsymmetric matrix, the converged result can also contain real two-by-two blocks representing complex eigenvalue pairs.

A shifted step with a check on the trailing mode

Take $A = \left(\begin{smallmatrix} 4 & 1 \\ 1 & 2 \end{smallmatrix}\right)$ and $\mu = 2$. Normalising the first shifted column gives $(2, 1)^T/\sqrt{5}$. Complete the orthogonal basis with $(-1, 2)^T/\sqrt{5}$. Multiplying $Q^T A Q$ gives diagonal $(22/5, 8/5)$ and off-diagonal $-1/5$. The trace remains six and the determinant remains seven. The last-coordinate Rayleigh estimate, initially two, is now 1.6 compared with the exact lower eigenvalue $3 - \sqrt{2} \approx 1.586$. The residual of this coordinate has fallen from one to 0.2. A subsequent shift can use the improved estimate. Preserving eigenvalues alone would never certify that a coordinate is ready for deflation.

Which boundary modes are physical?

An elastic assembly suspended in free space admits rigid translations, and in three dimensions also rigid rotations when its geometry allows them. In a correctly assembled model these motions cost zero strain energy, so the stiffness matrix has a nullspace and some generalised eigenvalues are zero. Attempting unshifted inverse iteration on that stiffness matrix cannot locate a lowest positive vibration by simply inverting the operator. One can instead constrain rigid motions consistently or project into the mass-orthogonal complement of their span.

For a two-mass system with unit masses and a unit spring connecting them,

$$K = \left(\begin{smallmatrix} 1 & -1 \\ -1 & 1 \end{smallmatrix}\right).$$

Translation $(1, 1)^T$ has zero squared frequency; opposite displacement $(1, -1)^T$ has squared frequency two. Grounding the first mass eliminates translation but changes the second mode to an oscillator with frequency one. Projection leaves the original free-system vibration at $\sqrt{2}$. These are physically different comparisons; the treatment of a nullspace determines which frequency an eigenvalue calculation is actually reporting.

Convert an eigenpair residual into an angle bound

Given a real symmetric matrix A, a unit-vector approximation v and a scalar θ, write its residual $r = Av - \theta v$. Expand $v = \sum_j c_j q_j$ in an orthonormal eigenbasis. The squared residual is $\sum_j (\lambda_j - \theta)^2 |c_j|^2$. This identity always implies that at least one eigenvalue lies within $\|r\|$ of θ. It does not by itself show whether the reported vector represents one particular physical mode. Two almost coincident eigenvalues can produce a tiny residual for many different mixtures of their eigenvectors.

To bound leakage out of a target eigenspace, let γ be a lower bound on the distance from θ to every eigenvalue outside that space. The squared residual is at least $\gamma^2 \sum_{\text{outside}} |c_j|^2$. Consequently the sine of the angle between v and the target space is at most $\|r\|/\gamma$. This result assumes a symmetric or Hermitian operator under the chosen inner product. For a mass-weighted mechanical problem, transform to mass-normal coordinates or perform the argument in the proper mass inner product. An ordinary Euclidean angle between unscaled coordinates may have little relevance to vibrational energy.

Resolve a cluster more reliably than either member

Take $A = \operatorname{diag}(10, 10.02, 30)$, and $v = (1, 1, 0)^T/\sqrt{2}$. Its Rayleigh quotient is 10.01. Multiplication gives $r = (-0.01, 0.01, 0)^T/\sqrt{2}$, of length 0.01. Relative to the cluster near ten this residual is small, yet the angle from the first coordinate eigenvector is 45 degrees. Using only the other low eigenvalue as a separation gives $\gamma = 0.01$, so the bound on that individual angle has sine at most one and provides no useful orientation guarantee.

Now take the target as the entire first two-dimensional subspace. Its gap from the reported quotient to the outside eigenvalue is

$$30 - 10.01 = 19.99.$$

The same residual implies leakage no greater than $0.01/19.99 \approx 0.000500$. In fact the true leakage of this particular vector is exactly zero. If the measured eigenvalues are squared angular frequencies, the two low modes have nearly the same frequency but can mix strongly under a small asymmetric manufacturing perturbation. Their combined subspace may remain very stable and may therefore be the more sensible object to compare with measurements.

A laboratory comparison can project independent excitation patterns onto the computed cluster and compare the total energy in that cluster. Such a result is invariant under rotations of its basis, which a simple table of pairwise eigenvector differences is not. If the gap to the rest of the spectrum closes as well, the subspace estimate becomes weak. Measuring the residual alone cannot repair that missing separation information.

For example the sum of squared projections of the excitation $(1, 0, 0)^T$ onto the first two computed orthonormal modes is one for any orthonormal basis of the exact low cluster. Individual projections can range between zero and one as the basis rotates. A test based on cluster projection assesses the reproducible quantity without requiring a uniquely labelled first mode.

Watch the QR iteration resolve two close eigenvalues

Rotate each diagonal matrix with eigenvalues $1, \rho$ by 0.6 radians, then factor $A_k = Q_k R_k$ and form $A_{k+1} = R_k Q_k$. Each step is an orthogonal similarity. Dividing the off-diagonal by the gap lets all three matrices start with the same normalized mixing.

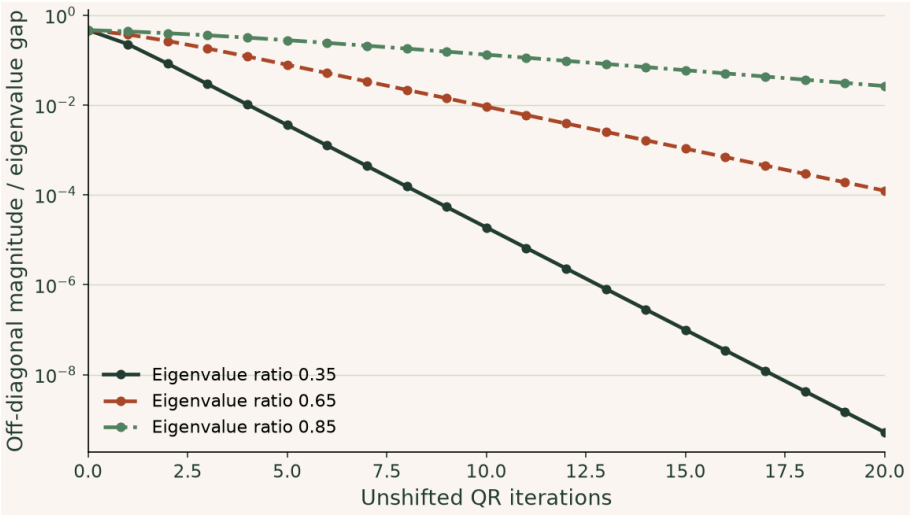

The tails are straight on a logarithmic vertical axis. Once the dominant direction is nearly aligned, one step multiplies the remaining mixing by approximately the eigenvalue ratio. At a ratio near one, an unchanged spectrum can coexist with slow convergence to a diagonal matrix.

▶ **Read values from the plot**

Change parameters · Watch the QR iteration resolve two close eigenvalues

Change one input at a time and compare the recalculated curves with the argument above. The browser evaluates this small model locally; the accompanying numerical values state its assumptions and reference checks. Reset restores the reference inputs, and the data can be downloaded as CSV.

Open this calculation with adjustable inputs

Try the calculation

Smaller eigenvalue ratio ρ **0.8**

Initial rotation (degrees) **35**

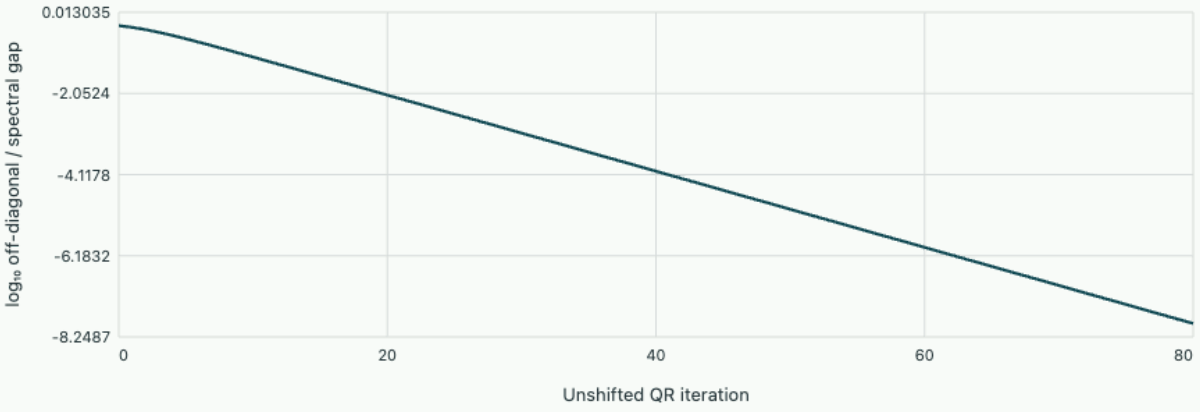

Orthogonal-similarity mixing

Asymptotic contraction factor 0.8. The eigenvalues remain 1 and 0.8 throughout the iteration.

Householder reflections and practical QR

Numerical linear algebra

Build one reflection explicitly, use it to triangularise a rectangular system and locate the residual in transformed coordinates.

FROM FIRST PRINCIPLES

Construct a reflection that zeros a column tail

A reflection reverses the component of a vector perpendicular to a chosen plane and leaves the component in the plane fixed. If a nonzero vector v is normal to the plane, the perpendicular projection of x is $v(v^T x)/(v^T v)$. Subtracting twice that projection defines

$$Hx = x - 2v\frac{v^T x}{v^T v},$$

$$H = I - 2\frac{vv^T}{v^T v}.$$

The reflection preserves lengths because the parallel and perpendicular components remain orthogonal. To reduce a column x, target an axis vector $y = \alpha e_1$ of equal length, so $|\alpha| = \|x\|_2$. Set $v = x - y$. Expanding the dot products gives $2v^T x = v^T v$ because x and y have equal norm. Consequently

$$Hx = x - v = y.$$

Choosing α opposite in sign to x_1 avoids cancellation in v_1. Reflecting only the unprocessed rows of each successive column yields an upper triangular factor while preserving previously generated zeros. An already zero column needs no transformation.

DERIVATION AND APPLICATIONS

Recover the least-squares residual from Householder coordinates

A QR solver can report how much an orthogonal transformation sends into rows below a triangular factor. That part has a precise measurement interpretation. Assume a real matrix A with $m > n$ and independent columns. Apply the same sequence of Householder transformations to the augmented matrix $[A\ b]$. Their product Q^T is orthogonal. Partition its result into an upper block with n rows and a lower block,

$$Q^T[A\ b] = \begin{pmatrix} R & c \\ 0 & d \end{pmatrix}.$$

Because orthogonal transformations preserve the Euclidean norm, the squared residual of any trial vector decomposes into the sum of two nonnegative terms. Only the first can be changed by the fit,

$$\|b - Ax\|_2^2 = \|c - Rx\|_2^2 + \|d\|_2^2.$$

The optimal x solves $Rx = c$, leaving residual norm $\|d\|_2$. In exact arithmetic the original-coordinate residual is obtained by reversing the transforms on a vector whose upper block is zero,

$$r = b - Ax = Q\begin{pmatrix}0\\ d\end{pmatrix},$$

$$A^T r = R^T[\,I\;\;0\,]\begin{pmatrix}0\\ d\end{pmatrix} = 0.$$

The lower coordinates are combinations of observations. Their individual values depend on the particular reflections, whereas their squared norm is invariant. Retaining that norm directly avoids subtracting two large nearly equal squared norms when the fit is excellent. If some triangular pivots are numerically negligible, ranking the columns and checking the model's effective rank become necessary before interpreting all n fitted coefficients.

WORKED EXAMPLE

Inspect residuals after fitting a constant

Fit $A = (1, 1)^T$ to readings $b = (8, 10)^T$. Rotate to the orthonormal directions $(1, 1)/\sqrt{2}$ and $(-1, 1)/\sqrt{2}$. The transformed data are $18/\sqrt{2}$ and $2/\sqrt{2}$, while the transformed model is $(\sqrt{2}, 0)^T x$. The fit gives $x = 9$, residual norm $\sqrt{2}$, and original residual $(-1, 1)^T$. Its sum is zero as the constant column requires.

SECOND APPLICATION

Compare the numerical consequences of subtracting squared norms

Suppose transformed coordinates of another fit have $\|c\| = 10^8$ and $\|d\| = 0.1$. Computing $\|b\|^2 - \|c\|^2$ would subtract two values near 10^{16} to recover 0.01. Binary64 spacing near 10^{16} is 2, so the small addition can disappear entirely. Keeping the transformed tail returns residual norm 0.1 directly, subject to normal transformation errors.

Eliminate a column without stretching it

An orthogonal transformation applied simultaneously to the data matrix and right-hand side leaves every Euclidean residual length unchanged. A Householder reflector is $H = I - 2vv^T/(v^T v)$ for a nonzero vector v. Decompose any vector into components parallel and perpendicular to v. The parallel component changes sign and the perpendicular component stays fixed, proving both length preservation and $H^2 = I$. Implementation requires just a dot product and a scaled vector subtraction rather than a dense matrix.

To make the entries below the first element of a column x vanish, choose $v = x - \alpha e_1$ where $|\alpha| = \|x\|_2$. A short calculation gives $v^T v = 2(\|x\|^2 - \alpha x_1)$ and $v^T x = \|x\|^2 - \alpha x_1$, hence

$$Hx = x - v = \alpha e_1.$$

Choose α with sign opposite to x_1 to avoid subtracting nearly equal positive numbers when computing v_1. A zero column requires a separate decision about rank.

Follow the data into the triangular system

For $x = (3, 4)^T$, use $\alpha = -5$ and $v = (8, 4)^T$. The reflection matrix is $\begin{pmatrix} -.6 & -.8 \\ -.8 & .6 \end{pmatrix}$. Take a one-column least-squares problem with measurements $b = (4, 3)^T$. Multiplication sends x to $(-5, 0)^T$ and b to $(-4.8, -1.4)^T$. The optimum coefficient matches the first transformed measurement, giving $c = 0.96$. The second transformed measurement remains a residual of magnitude 1.4. Back in physical coordinates the predicted vector is $(2.88, 3.84)$, with residual $(1.12, -0.84)$ and the same length.

For several columns, apply the first reflector to the entire remaining block. Construct the next reflector only on rows below the first pivot, protecting the already triangular part. Saving the reflector vectors permits application of Q or Q^T later without explicitly forming either full matrix. Additional right-hand sides can then use the same orthogonal transformations and triangular solve.

Detect when a coefficient is poorly determined

If a new column is nearly in the span of earlier columns, little of it survives below the preceding pivots. A small trailing diagonal of R signals potential coefficient sensitivity, but its scale depends on column units and order. An instrument model with one parameter measured in micro-units can naturally have a small column. Scale the columns according to the intended range of the corresponding parameters before interpreting a near-zero diagonal.

Column-pivoted QR selects a remaining column with a large unrepresented component at each stage. It can help reveal effective rank and identify an independent set of measurements or basis functions. A numerical rank decision still needs a tolerance tied to data uncertainty or floating-point arithmetic. Reporting a fitted coefficient to ten digits after discarding an almost invisible direction would conceal the limitation that the factorisation has revealed.

Check a nearly aligned column

Consider $x = (1, 10^{-8})^T$. Its norm differs from one by only about 5×10^{-17}, below the spacing between one and the next larger double precision number. Computing the exact difference $x_1 - \|x\|$ as that subtraction may therefore give zero. Choosing the target with opposite sign makes the first entry of the reflector vector about two and the remaining small entry is retained. This calculation explains the sign convention geometrically: two opposite points on a sphere provide valid reflection targets, and only one directs the subtraction away from severe cancellation.

Update a triangular fit when observations arrive sequentially

A Householder QR factorisation is efficient when the entire design matrix is available. Suppose instead that a new row arrives after the original observations have been reduced to an upper triangular factor R, a transformed data vector z, and an accumulated residual sum of squares S. The original least-squares objective equals $\|Rx - z\|^2 + S$. Appending measurement row a^T and value b adds $(a^T x - b)^2$. Stacking a^T below R produces a matrix that is triangular except for one row. Orthogonal plane rotations can restore triangular form by eliminating each incoming entry in turn.

To remove a pair $(r, s)^T$, let $\rho = \sqrt{r^2 + s^2}$, $c = r/\rho,\ d = s/\rho$. Multiplication by the rotation with rows (c, d) and $(-d, c)$ gives $(\rho, 0)^T$. Apply precisely the same two-row operation to all remaining columns and to the transformed data. Its final bottom data entry has no corresponding design column, so its square adds to S. If $r = s = 0$ there is nothing to eliminate. A numerical implementation obtains the length with a scaled hypotenuse evaluation to avoid overflow from squaring a large observation.

Insert a point into an existing two-parameter fit

Suppose the existing transformed problem has $R = I,\ z = (1, 2)^T,\ S = 0$. A new sensor reports $x_1 + x_2 = 5$. The best prior estimate was $(1, 2)^T$, which predicts three and therefore disagrees with the incoming reading by two. To reconcile all observations, stack the new row $(1, 1 \mid 5)$ below $(1, 0 \mid 1)$ and $(0, 1 \mid 2)$. Rotate the first and last rows with $c = d = 1/\sqrt{2}$. The first becomes $(\sqrt{2}, 1/\sqrt{2} \mid 6/\sqrt{2})$; the last becomes $(0, 1/\sqrt{2} \mid 4/\sqrt{2})$.

The second rotation combines leading entries 1 and $1/\sqrt{2}$, yielding $\rho = \sqrt{3/2}$, $c = \sqrt{2/3}$, $d = 1/\sqrt{3}$. Its surviving right-hand side is

$$2\sqrt{2/3} + 4/\sqrt{6} = 8/\sqrt{6}.$$

Thus the second-coordinate estimate is

$$(8/\sqrt{6})/\sqrt{3/2} = 8/3.$$

Back substitution in the first row gives

$$x_1 = (6 - 8/3)/2 = 5/3.$$

The final bottom data entry is $-2/\sqrt{3} + 4/\sqrt{3} = 2/\sqrt{3}$, contributing $4/3$ to the residual sum of squares.

Verify directly: the three residual magnitudes are all $2/3$, so their squared sum is $4/3$. The new reading moved both coordinates equally because their original transformed information was identical and its row is symmetric. No historical measurement needed to be reread. Conversely deleting an earlier reading can expose rank deficiency and is a different numerical operation; simply running the same update with a negated data value would add another squared term rather than remove information.

When a stream contains different measurement variances, divide each arriving row and its right-hand side by its independently known standard deviation before applying the rotations. The accumulated S then measures squared standardised discrepancies. If the new reading shares an error with earlier readings, dividing its row alone is insufficient. One must first account for the conditional error and correlated design contribution or refactor a suitable joint whitening transformation. This is a measurement-model limitation, even when the rotations themselves remain perfectly orthogonal.

Repeated readings of a common reference are one practical source of precisely this kind of dependence.

Change parameters · Explore the worked calculation

Change one input at a time and compare the recalculated curves with the argument above. The browser evaluates this small model locally; the accompanying numerical values state its assumptions and reference checks. Reset restores the reference inputs, and the data can be downloaded as CSV.

Open this calculation with adjustable inputs

Try the calculation

$\log_{10}$ fitted component magnitude **8**

Residual-tail magnitude **0.1**

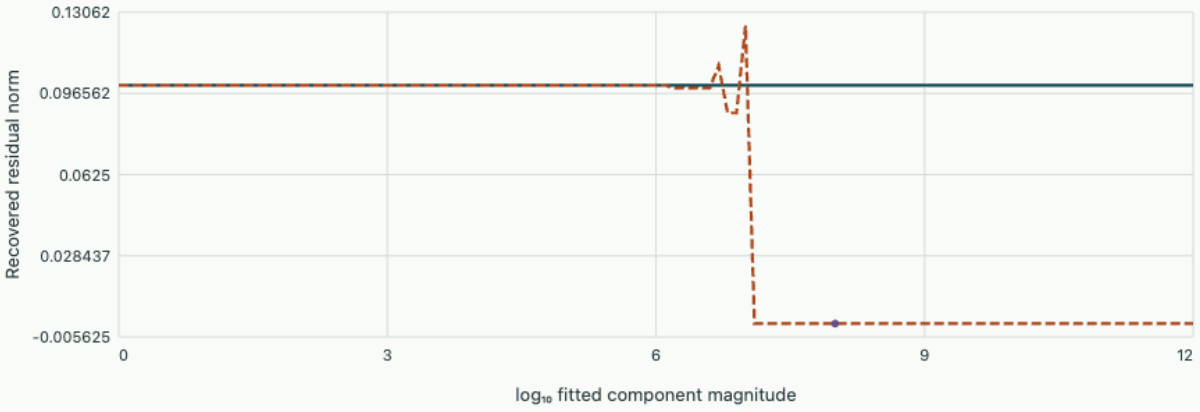

Stored orthogonal tail | Difference of squared norms | Selected subtractive residual

At fitted magnitude 10^8, the subtractive residual is 0 while the preserved tail is 0.1. Binary64 rounding can erase the smaller squared term.

Arnoldi iteration and GMRES

Numerical linear algebra

Construct an orthonormal basis for a nonsymmetric operator and use a small least-squares problem to choose an approximation.

FROM FIRST PRINCIPLES

Project a large operator onto successive Krylov directions

For $Ax = b$, start with residual $r_0 = b - Ax_0 \neq 0$. Choose the first basis vector as $v_1 = r_0/\beta$, where $\beta = \|r_0\|_2$. Repeated matrix products generate a search space, but their columns can become nearly parallel. At step j, remove previously resolved components of Av_j before normalizing the remainder,

$$h_{ij} = v_i^T A v_j,$$

$$w = Av_j - \sum_{i=1}^{j} h_{ij} v_i,$$

$$h_{j+1,j} = \|w\|_2.$$

If $h_{j+1,j} > 0$, take $v_{j+1} = w/h_{j+1,j}$. The resulting Arnoldi identity $AV_k = V_{k+1}\bar{H}_k$ expresses each large product using a small upper-Hessenberg matrix. A correction $V_k y$ gives the residual norm

$$\|r_0 - AV_k y\|_2 = \|\beta e_1 - \bar{H}_k y\|_2.$$

Orthonormality justifies the equality. GMRES solves the smaller least-squares problem to minimize the actual residual over this search space. When the new norm is zero the subspace is already invariant, and no further independent direction is available through another ordinary Arnoldi step.

Background · Krylov subspace

$$\mathcal{K}_k(A, b) = \operatorname{span}(b, Ab, \ldots, A^{k-1}b)$$

The starting vector determines which parts of the spectrum are visible. Orthogonalisation avoids storing nearly dependent raw powers.

DERIVATION AND APPLICATIONS

Calculate exactly when the first GMRES step stalls

For one Krylov step the only available correction is a multiple of the initial residual $r = b - Ax$. The resulting change in the residual is parallel to Ar. Assume a real nonsingular A and $r \neq 0$, which guarantees $Ar \neq 0$. To choose the coefficient α, expand the squared residual norm,

$$\|r - \alpha Ar\|_2^2 = r^T r - 2\alpha (Ar)^T r + \alpha^2 \|Ar\|_2^2.$$

The derivative vanishes at the coefficient of the orthogonal projection of r onto Ar. Hence the new residual is perpendicular to the direction in which it was allowed to change,

$$\alpha_* = \frac{(Ar)^T r}{\|Ar\|_2^2}, \qquad (Ar)^T (r - \alpha_* Ar) = 0.$$

Let θ denote the angle between r and Ar. At the minimum the squared length removed is exactly the squared length of that projection. The remaining fraction of residual is

$$\frac{\|r_{\text{new}}\|}{\|r\|} = \sqrt{1 - \frac{[(Ar)^T r]^2}{\|Ar\|^2 \|r\|^2}} = |\sin\theta|.$$

Thus equal or opposite directions allow an exact one-step solve, while perpendicular directions deliver no change. This is purely geometric and applies without symmetry. It also clarifies why restarting after one step can be costly: later basis vectors could expose a direction inaccessible on the first line.

WORKED EXAMPLE

Compute the residual reduction for a shear

Use $A = \left(\begin{smallmatrix} 1 & 2 \\ 0 & 1 \end{smallmatrix}\right)$ and $r = (0, 1)^T$. Then $Ar = (2, 1)^T$, $\alpha_* = 1/5$, and $r_{\text{new}} = (-0.4, 0.8)^T$. Its norm is

$$\sqrt{0.8} = 0.894.$$

The dot product with Ar is $-0.8 + 0.8 = 0$, an independent check that the selected coefficient minimizes the residual along this line.

SECOND APPLICATION

Detect stagnation despite a nonsingular operator

Take $A = \operatorname{diag}(2, -1)$ and $r = (1, \sqrt{2})^T$. The dot product $r^T Ar = 2 - 2 = 0$ forces the first coefficient to zero. A second direction completes a two-dimensional space because $Ar = (2, -\sqrt{2})^T$ is independent of r. In fact $A^{-1} r = (0.5, -\sqrt{2})^T = -0.5 r + 0.5 Ar$. The two-direction correction solves the problem exactly although a one-step restart would repeat the same state.

A residual as a polynomial

Consider a square nonsingular matrix and an initial residual $r_0 = b - Ax_0$. A correction in the span of $r_0, Ar_0, \ldots, A^{k-1} r_0$ is a polynomial in A applied to r_0. After subtracting the action of A, the new residual is $p_k(A) r_0$ with $p_k(0) = 1$. This constraint matters because the zero polynomial would certify convergence before the calculation began. Generalised minimal residual iteration, or GMRES, minimises the Euclidean norm

over all such residuals. The spaces are nested, so the exact-arithmetic residual norms cannot increase during a cycle.

Replacing large products by small coordinates

Start with $v_1 = r_0/\beta$, where $\beta = \|r_0\|_2$. For each new direction form $w = Av_j$, subtract its projections onto the existing orthonormal columns, and normalise what remains. Each projection coefficient occupies one row of column j of a small matrix. The final normalising length occupies its first subdiagonal. This Arnoldi construction gives

$$AV_k = V_{k+1}\bar{H}_k.$$

If the length vanishes, the current subspace is invariant and no further normalisation is possible.

$$\|b - A(x_0 + V_k y)\|_2 = \|V_{k+1}(\beta e_1 - \bar{H}_k y)\|_2 = \|\beta e_1 - \bar{H}_k y\|_2.$$

The second equality uses orthonormality. Solve this small rectangular least-squares problem, for example by updating triangular factors using plane rotations. The resulting residual is orthogonal to the range of AV_k. Repeated projection may be necessary when rounding spoils orthogonality. A freshly calculated physical residual can test the small model's estimate before reporting convergence.

One eigenvalue and two directions

Take $A = \left(\begin{smallmatrix} 1 & 4 \\ 0 & 1 \end{smallmatrix}\right)$, $x_0 = 0$, $b = e_2$. The first correction is αe_2, so the squared residual is $16\alpha^2 + (1-\alpha)^2$. Differentiation gives $\alpha = 1/17$ and length $4/\sqrt{17}$, about 0.9701. Almost the entire initial residual survives despite every eigenvalue being one. The second direction comes from the component of Ae_2 perpendicular to e_2, namely $4e_1$. With both directions available the exact solution $(-4, 1)^T$ is accessible, and its residual vanishes.

For a diagonal matrix equal to the identity, a single step would suffice for every right-hand side. Here a nonzero off-diagonal entry couples two coordinates and the minimal polynomial is $(z-1)^2$. Counting distinct eigenvalues alone misses that coupling.

Paying for a long memory

Each new direction needs storage and projections against earlier directions. Restarting after a fixed number releases memory, but discards those directions. A new cycle still starts from the best current iterate. It lacks the previous cycle's simultaneous freedom to alter its old coefficients. Record matrix-vector products and the residual after each restart when comparing costs. On the example, two directions retained together already achieve the exact answer, providing a direct check on a proposed implementation.

Read a preconditioned residual

Left preconditioning with M^{-1} makes GMRES minimise $\|M^{-1}(b-Ax)\|$ instead. If $M=\operatorname{diag}(1,1000)$, a physical residual $(0,1)^T$ has preconditioned norm 0.001. A small number in the stopping display could therefore leave the second physical balance comparatively inaccurate. Right preconditioning solves $AM^{-1}y=b$, then forms $x=M^{-1}y$; for exact preconditioner actions its minimisation still uses the original residual. Compare work and physically meaningful residuals when assessing either choice, especially if the equations use different units.

How restarting can erase a necessary direction

GMRES minimises the full residual norm over its current correction space. A one-step restart retains only the first correction direction generated from the new residual. If that direction makes no angle with the gradient of the residual objective in the appropriate sense, progress can cease even for an invertible, perfectly conditioned matrix. To see the geometry, suppose the current residual is r. The one-dimensional trial residual is $r-\alpha Ar$. Minimising its squared norm for real arithmetic yields

$$\alpha=(Ar)^Tr/\|Ar\|^2.$$

Thus no reduction occurs when Ar is perpendicular to r. Restarting immediately recreates precisely the same subproblem.

Unrestarted GMRES can still add Ar and further powers to its Krylov space. In exact arithmetic it can therefore escape one-step stagnation. This distinction matters for storage allocation: reducing the basis length can change the convergence mechanism itself. Merely multiplying the allowed number of restart cycles cannot recover directions if every cycle begins and ends with the same residual. Conversely increasing the restart length helps only if the additional space contains directions that produce a lower residual.

Compare restart lengths for a quarter-turn operator

Set $A=\left(\begin{smallmatrix}0&-1\\1&0\end{smallmatrix}\right)$ and $b=(1,0)^T$, beginning at zero. Its transpose times itself is the identity, so its two-norm condition number is one. Yet $Ab=(0,1)^T$ is perpendicular to the initial residual b. The residual after trial step αb is $(1,-\alpha)^T$, whose squared length is $1+\alpha^2$. The best step is exactly zero. GMRES restarted every step thus leaves the residual norm at one forever.

With a two-vector basis use corrections $x=\alpha b+\beta Ab=(\alpha,\beta)^T$. The residual is $(1+\beta,-\alpha)^T$, which vanishes for $\alpha=0,\beta=-1$. The exact solution is reached. In polynomial language no degree-one real polynomial satisfying $p(0)=1$ can vanish at both eigenvalues $+i,-i$. The degree-two choice $p(z)=1+z^2$ does so because $A^2=-I$. Both formulations explain the same measurement without attributing the failure to singularity or loss of precision.

For a nonsingular real skew-symmetric matrix the one-step numerator vanishes for every real r, since a scalar r^TAr equals its own negative after transposition. This family therefore offers a repeatable stress test for short-restart implementations. A preconditioner or a shift that alters the search geometry can change the result, but any claimed gain should be checked using the residual of the original equations and a comparable count of matrix-vector products.

There is also a simple residual recomputation check at convergence: for the two-step state $(0, -1)^T$, multiplying by the original matrix gives $(1, 0)^T$ exactly. If a code reports the same state after one iteration, it has either used a different search direction or mislabelled its iteration count. Exposing the current basis length alongside the total matrix-vector count makes a restart experiment easier to reproduce.

For complex arithmetic, use the real or complex least-squares minimisation appropriate to the actual correction space; the perpendicularity argument above assumes real data.

Change parameters · Explore the worked calculation

Change one input at a time and compare the recalculated curves with the argument above. The browser evaluates this small model locally; the accompanying numerical values state its assumptions and reference checks. Reset restores the reference inputs, and the data can be downloaded as CSV.

Open this calculation with adjustable inputs

Try the calculation

Shear coupling **2**

Initial residual angle (degrees) **90**

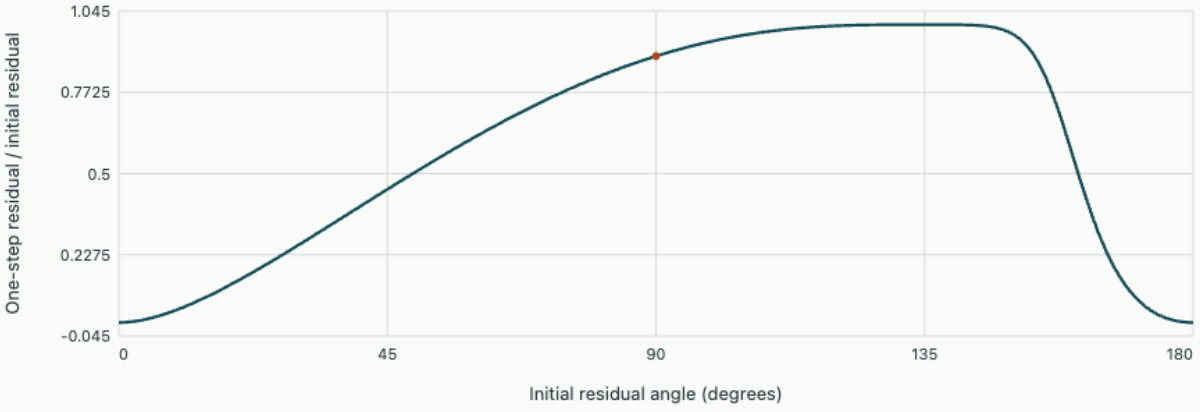

One-dimensional GMRES correction | Selected residual

The shear matrix has both eigenvalues equal to one. Selected first-step residual factor 0.89443. At most two exact-arithmetic directions solve this nonsingular two-variable problem.

How a moderately accurate mode can give a precise frequency

Numerical linear algebra

Use the Rayleigh quotient to estimate an eigenvalue, understand its stationary error and check a proposed mode with a residual.

A frequency estimate from a trial shape

For $Kq = \lambda Mq$ with symmetric K and positive definite mass M, a trial vector v fixes a direction but leaves an amplitude free. Substituting $q(t) = a(t)v$ in kinetic and elastic energy gives effective mass $v^T M v$ and stiffness $v^T K v$. Their ratio $R(v)$ estimates the squared frequency. Rescaling v has no effect. For $M = I$, differentiating $v^T K v / (v^T v)$ gives a gradient proportional to $Kv - Rv$. Its zeros are precisely the eigenvectors.

Angular error and the residual

Take exact orthonormal modes u_1, u_2 and a trial $v = u_1 \cos\theta + u_2 \sin\theta$. Orthogonality removes all cross terms. Then

$$R(v) - \lambda_1 = (\lambda_2 - \lambda_1)\sin^2\theta,$$
$$\|Kv - R(v)v\|_2 = |\lambda_2 - \lambda_1|\,|\sin\theta\cos\theta|.$$

The eigenvalue error is second order in a small angular error. The residual is first order. For general dimension write $v = \sum c_i u_i$ with unit norm. The squared residual is $\sum |c_i|^2(\lambda_i - R)^2$, so at least one eigenvalue lies within one residual norm of R. Identifying which mode requires more information about the spectrum.

Two measurements

For squared frequencies 4 and 9 in s⁻² and a 10-degree error in the first trial mode, $R = 4 + 5(0.03015) = 4.1508$ s⁻². The estimated angular frequency is 2.0373 rad/s. The relative error in frequency is about half the relative error in its square at small errors.

For eigenvalues 4 and 4.01, a 45-degree mixture gives a residual of just 0.005. This very small number is consistent with a large uncertainty in shape. If the two eigenvalues were identical, every combination in the eigenspace would be an exact mode.

Tilt the trial vector

The plot shows the quotient error and residual across trial angles for a chosen spectral gap. Move the angle marker, halve a small angle and compare the two responses.

Open this experiment online

Try the calculation

Separation of squared frequencies (s⁻²) 5

Trial angle (degrees) 10

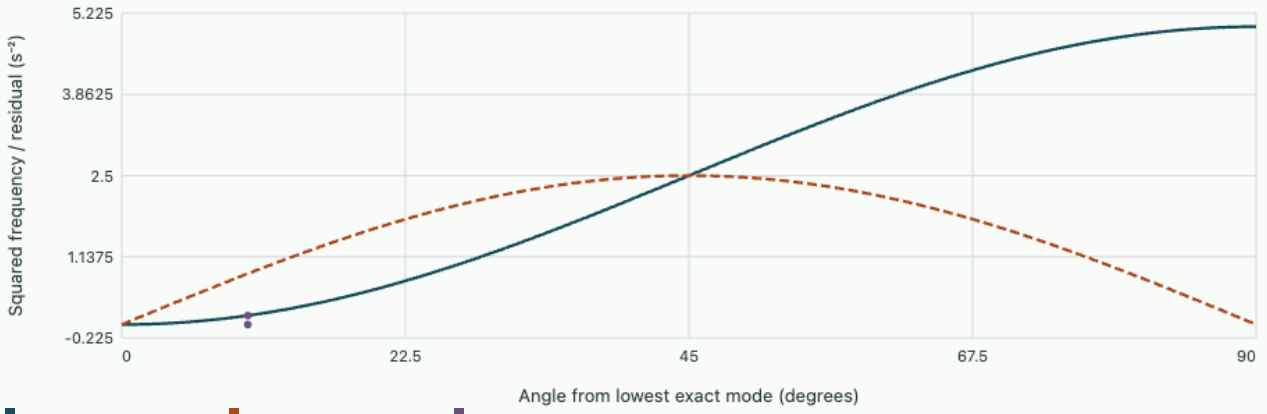

Quotient error | Residual norm | Selected angle

Selected angle gives quotient error 0.15077 s⁻² and residual 0.85505 s⁻².

Tutorial on frequency estimates

For eigenvalues 2 and 10, a trial produces $R = 2.08$. Assuming the two-mode model, estimate its angular error.

▼ **Worked solution**

The mixture fraction is

$$(2.08 - 2)/(10 - 2) = 0.01.$$

Hence $|\sin\theta| = 0.1$, or about 5.74 degrees near the lower mode.

A trial gives $R = 5$ and residual zero. Does it prove that 5 is the lowest eigenvalue?

▼ **Worked solution**

It proves that the trial vector is an exact eigenvector for 5. To certify the lowest eigenvalue one also needs a lower bound or a description of the remaining spectral subspace.

Read the university treatment

Compare how quickly an eigenvector direction converges with the rate for its Rayleigh-quotient eigenvalue estimate. Driscoll and Braun · Power iteration

Estimating a cantilever frequency from energy

Machine dynamics

A clamped beam carries a small sensor at its tip. Its precise mode shape is unknown. Use the static deflection shape to approximate the fundamental frequency and see why concentrating a small mass at the tip has a large effect.

Choose a displacement field that respects the support

Write $w(x,t) = q(t)\phi(s)$ with $s = x/L$, taking q as tip displacement. Choose the shape from a uniform cantilever under a static transverse tip force. The bending moment varies linearly from root to tip. Integrating curvature twice with $w(0) = w'(0) = 0$, then normalizing at $s = 1$, gives

$$\phi = s^2(3 - s)/2.$$

The shape satisfies both essential clamp conditions. Its free-tip shear need not match the exact unloaded vibration mode. The energy formulation can use such admissible trial shapes.

Background · Rayleigh quotient

$$R(u) = u^{\mathsf{T}} K u / (u^{\mathsf{T}} M u)$$

For symmetric stiffness and positive mass this estimates squared angular frequency. A kinematically admissible displacement is required for a useful structural estimate.

Compute effective stiffness and mass from the same field

The strain energy density of a slender beam in bending is $\frac{1}{2} EI w_{xx}^2$. Its translational kinetic energy is $\frac{1}{2}\rho A w_t^2$ per unit length. Changing variables from x to s introduces different powers of L into the two integrals. Include a point mass M at the tip, neglecting its rotary inertia.

$$U = \frac{EIq^2}{2L^3}\int_0^1 (\phi'')^2\, ds,$$

$$K = \frac{\dot{q}^2}{2}\left[\rho A L \int_0^1 \phi^2\, ds + M\right].$$

Since $\phi'' = 3(1 - s)$, the stiffness integral is 3. Squaring $s^2(3 - s)/2$ and integrating $(9s^4 - 6s^5 + s^6)/4$ gives $33/140$. The Euler–Lagrange equation for q is consequently a single oscillator equation with

$$\omega_R^2 = \frac{3EI/L^3}{(33/140)\rho A L + M}.$$

For M=0 this gives the dimensionless coefficient $\sqrt{140/11} = 3.5675$. Solving the exact unloaded cantilever condition $\cosh\beta\cos\beta = -1$ gives

$$\beta_1^2 = 3.5160.$$

The estimate is high by about 1.5%. The Rayleigh minimum principle explains that direction of error for conservative systems with admissible trial shapes. Restricting motion to a single shape restricts the set over which the energy quotient can attain its minimum.

Choose a sensor after estimating its dynamic effect

Take $EI = 10$ N m², $L = 0.50$ m and total beam mass 0.40 kg. The tip stiffness is 240 N/m and the distributed effective mass is only 0.0943 kg. The estimate gives 8.03 Hz without a sensor and 6.49 Hz with a 0.05 kg tip sensor. The sensor weighs one eighth of the beam, yet it has over half the effective mass associated with the chosen tip coordinate.

Place the same sensor at half-span instead. There $\phi(0.5) = 0.3125$. Its kinetic energy contributes $M\phi(0.5)^2 = 0.00488$ kg to the effective mass, so the estimate rises to 7.83 Hz. The sensor measures a smaller local motion as it perturbs the frequency less. Compare that trade-off at different attachment positions with the plot.

Move along the position axis with a fixed sensor mass. Identify the interval in which frequency is most sensitive to changing attachment position. Double the sensor mass and check whether the change in frequency doubles at the tip.

Open this experiment online

Try the calculation

Attached mass (kg) **0.05**

Bending rigidity (N m²) **10**

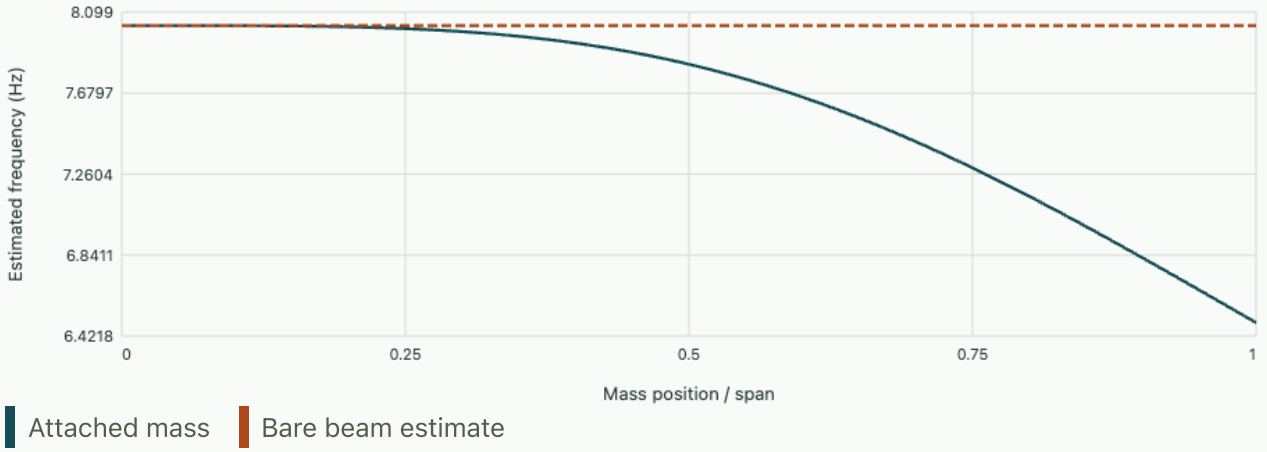

The estimate is 6.49 Hz with the sensor at the tip and 7.83 Hz at half-span. The bare-beam estimate is 8.03 Hz.

A related OneSimulation calculation

The course frame model has several modes and an independently variable payload rotary inertia. Compare a pure translation-mass change with a rotary-inertia change; the scalar point-mass estimate above addresses

the first mechanism.

Open the OneSimulation frame eigenvalue experiment

Tutorial · estimate an unknown attached mass

The same bare beam is fitted with a small tip part and resonates at 6.0 Hz. Estimate the attached mass. Give one reason the inferred mass could differ from a scale reading even when the frequency measurement is accurate.

▼ **Show the mass estimate**

The required effective mass is

$$240/(2\pi \cdot 6)^2 = 0.1689\,\mathrm{kg}.$$

Subtract 0.0943 kg for the beam to obtain about 0.0746 kg. Clamp compliance would reduce the effective stiffness and frequency, which this fit would interpret as extra tip mass. The trial shape and omission of rotary inertia also limit the estimate for a part with appreciable extent.

Further reading · Bower, Applied Mechanics of Solids · Rayleigh–Ritz methods

Recover a trajectory from a family of actions

Classical mechanics

In Hamilton–Jacobi theory a partial differential equation for one function replaces several trajectory equations. A falling bead makes the different roles of position, energy and arrival time explicit.

Use the end of the path as a variable

Let a bead move between a fixed initial event and an adjustable final event. Denote by $S(q,t)$ the action evaluated on the actual trajectory ending at (q,t). For a small change in the endpoint, the interior variation cancels by the Euler–Lagrange equation. If time stays fixed the boundary term is $p\,\delta q$. If time changes, part of the endpoint displacement comes from continuing along the existing path by $\dot q\,\delta t$. The remaining independent displacement is $\delta q - \dot q\,\delta t$. Thus

$$\delta S = p(\delta q - \dot q\,\delta t) + L\,\delta t = p\,\delta q - H\,\delta t.$$

It follows that $p = \partial_q S$ and $-H = \partial_t S$. Eliminate p to obtain

$$\partial_t S + H(q, \partial_q S, t) = 0.$$

The equation involves partial derivatives because it describes an entire family of possible endpoints. By contrast, $dS/dt = L$ along a single trajectory. A useful dimensional check is that a derivative of action with respect to length is momentum, whereas a derivative with respect to time is energy.

Separate time and build a clock from energy

For a time-independent Hamiltonian try $S = W(q;E) - Et$. With $H = p^2/(2m) + U(q)$, this gives

$$(\partial_q W)^2 = 2m(E - U).$$

A trajectory on the increasing-position branch has positive momentum, so an integral from a fixed reference position q_0 supplies W. Differentiate at fixed final position with respect to the family parameter E:

$$W = \int_{q_0}^{q} \sqrt{2m(E-U(s))}\,ds,$$

$$\partial_E W = \int_{q_0}^{q} \frac{m\,ds}{\sqrt{2m(E-U(s))}}.$$

The integrand of the second integral is $1/v$. Hence $\partial_E W$ is the travel time from the reference point. Setting $\partial_E S = 0$ chooses a departure time of zero with this lower integration limit. A different constant amounts to shifting the clock. At a turning point $p = 0$, the one-way branch must be changed to continue the motion.

Work two arrival-time problems

Let $q = x$ point downwards, $U = -mgx$, and release the bead through $x = 0$ with speed u. Its energy is

$$E = mu^2/2.$$

The integral yields

$$t = (\sqrt{u^2 + 2gx} - u)/g.$$

Squaring the relation after isolating the square root gives $x = ut + gt^2/2$, the same trajectory as Newton's equation. For $u = 2.0$ m/s, $g = 9.81$ m/s² and $x = 3.0$ m the arrival time is about 0.604 s. The endpoint momentum uses the final speed, about 7.93 m/s, rather than the average over the fall.

Reverse the experiment by pointing x upward in the same field. If $u = 4.0$ m/s, the maximum height is 0.815 m. A gate 0.50 m above launch is crossed twice. From $x = ut - gt^2/2$ the two times are 0.154 s and 0.661 s. The need for two branches of W is experimentally visible in those two timestamps. A calculation that takes the positive square root everywhere loses the return journey.

Compare two launch speeds

The plot shows upward trajectories and their common observation height. Lower the launch speed towards the value that barely reaches the gate. Watch what happens to the separation between the two arrival times.

Open this experiment online

Try the calculation

Upward launch speed (m/s) **4**

Gate height (m) **0.5**

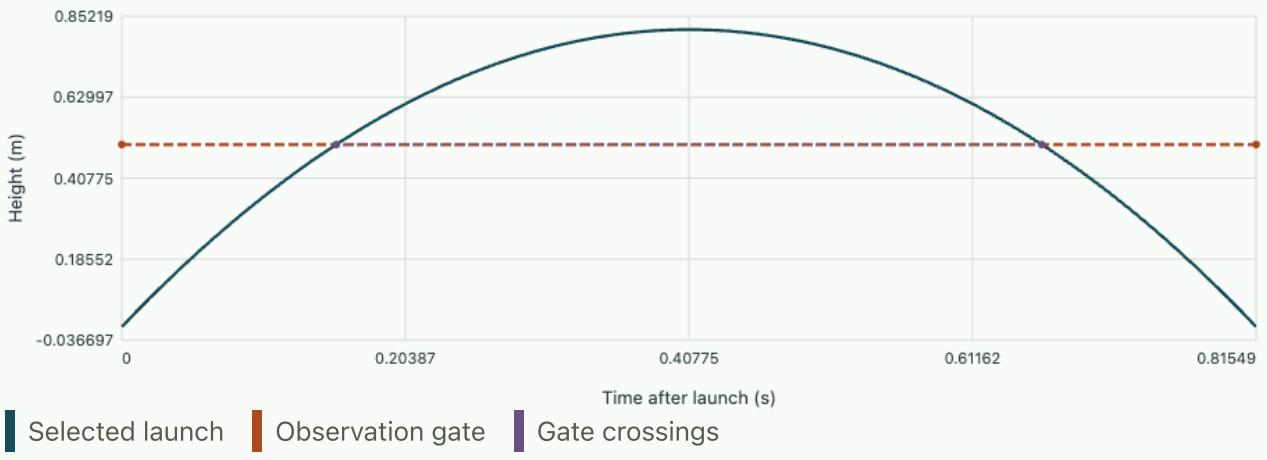

Peak height 0.8155 m. Ascending crossing at 0.1541 s and descending crossing at 0.6614 s.

Tutorial on the double root

A gate records two crossings 0.40 s apart in a uniform field of 9.81 m/s². How high does the trajectory rise above that gate?

▼ **Work through the solution**

The two crossings are symmetric about the maximum. The fall from the maximum takes 0.20 s, giving $g(0.20)^2/2 = 0.1962$ m. Equivalently the speed at the gate is $g(0.40)/2 = 1.962$ m/s and its kinetic energy per unit mass gives the same height.

Estimate the arrival-time sensitivity of a gate on the rising branch. At $u = 4.0\,\mathrm{m/s}$ compare $h = 0.50$ m with $h = 0.80$ m.

▼ **Work through the solution**

Implicit differentiation of $h = ut - gt^2/2$ at fixed launch speed gives

$$dt/dh = 1/(u - gt).$$

On ascent this is $1/\sqrt{u^2 - 2gh}$. Its value grows from $0.402\,\mathrm{s/m}$ at 0.50 m to $1.81\,\mathrm{s/m}$ at 0.80 m. Close to the top a small height error produces much larger timing uncertainty because the trajectory crosses slowly.

Does adding a constant C to a potential alter the predicted arrival times from Hamilton–Jacobi theory when the initial position and speed are unchanged?

▼ **Work through the solution**

Both the total energy label and potential shift by C. Hence $p = \sqrt{2m(E - U)}$ and the abbreviated action as a path integral of p are unchanged for that orbit. The full action acquires $-Ct$. Its spatial derivative and the time crossing predictions remain the same. Comparing the same numerical energy before and after the shift would specify different trajectories.

Further reading · Fließbach, Hamiltonformalismus · Tong, Hamilton–Jacobi theory.

Read a period from the area enclosed by an orbit

Classical mechanics

A track with a V-shaped potential has straight force segments but a strongly amplitude-dependent period. Its area in phase space gives a useful way to predict both the period and a slow change of the track.

Construct the closed area

For one bound degree of freedom at energy E, position runs between two turning points. In a mechanical Hamiltonian $p^2/(2m) + U(q)$, the upper branch of the phase portrait has positive momentum and the lower branch has its negative. The directed integral around one full orbit counts the enclosed area. Define the action J using the area divided by 2π. This choice makes its conjugate phase an angle measured in radians:

$$J(E) = \frac{1}{2\pi} \oint p\, dq = \frac{1}{\pi} \int_{q_-}^{q_+} \sqrt{2m(E - U(q))}\, dq.$$

Now change the energy slightly. The integrand vanishes at each regular endpoint, so the moving-endpoint contributions to the derivative vanish. In the interior $\partial p/\partial E = m/p = 1/v$. The derivative of the upper-branch area therefore measures the time to reach the opposite turning point.

$$\frac{dJ}{dE} = \frac{T(E)}{2\pi},$$

$$\omega(E) = \frac{dE}{dJ}.$$

This argument also explains why the construction becomes delicate near a separatrix. A trajectory can spend an arbitrarily long time near an unstable equilibrium, and T can diverge even though the enclosed area stays finite.

Background · Action variable

$$J = (2\pi)^{-1} \oint p\, dq$$

This area measure is defined for a closed cycle. It can change strongly near a separatrix even under a slow parameter variation.

Carry out the integral for a constant restoring force

Use $U(q) = F|q|$ with $F > 0$. The turning amplitude is $A = E/F$. Integrate the positive branch from 0 to A and use reflection symmetry. Substitution $s = E - Fq$ gives

$$J = \frac{4\sqrt{2m}}{3\pi F} E^{3/2},$$

$$T = \frac{4\sqrt{2mE}}{F} = 4\sqrt{\frac{2mA}{F}}.$$

The answer has an independent mechanical check. From a turning point the bead accelerates uniformly with magnitude F/m until the centre. Covering a distance A from rest takes $\sqrt{2mA/F}$, a quarter period. Doubling the amplitude increases the full period by $\sqrt{2}$. The abrupt change of force direction at the centre leaves the velocity continuous.

Slowly steepen the track

A slowly adjusted parameter can preserve the orbit action to leading order when the motion remains periodic and many cycles fit inside the adjustment time. For this trap m stays constant, so constant J requires $E^{3/2}/F$ to stay fixed. Therefore $E \propto F^{2/3}$,

$$A = E/F \propto F^{-1/3},$$

$$T \propto F^{-2/3}.$$

Work done by adjusting the slope changes the energy despite the approximate conservation of action.

Start with $m = 0.25$ kg, $F = 2.0$ N and $A = 0.080$ m. The initial energy is 0.160 J and the period 0.566 s. A slow eightfold increase of F makes the energy four times larger, 0.640 J, and the amplitude half as large, 0.040 m. The period becomes 0.141 s. In contrast, switching slopes instantaneously at the centre leaves the kinetic energy and position unchanged at that instant. The same final slope would then give amplitude 0.010 m. The route by which a parameter changes matters.

Compare closed orbits at constant action

The axes are position and momentum. Change the slope ratio. Predict how the turning positions and the central momentum move in opposite directions. Their changes balance in the enclosed area.

Open this experiment online

Try the calculation

Final to initial force ratio

2

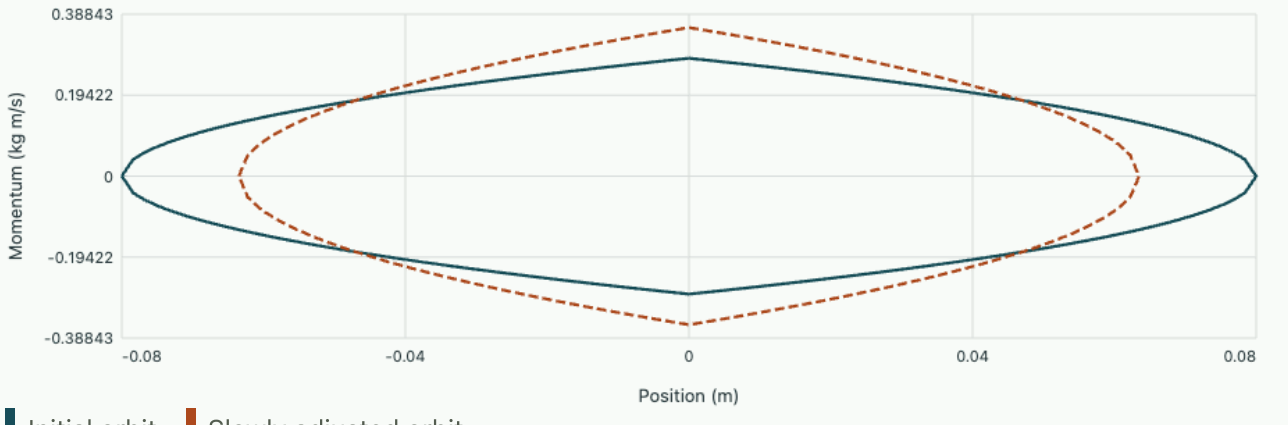

Final amplitude 0.0635 m, energy 0.254 J and period 0.3564 s.

Tutorial on inferring a change

After a slow adjustment the period is 64% of its original value. Find the force and amplitude ratios.

▼ Work through the solution

From $T_f/T_i = (F_f/F_i)^{-2/3}$, the force ratio is

$$0.64^{-3/2} = 1.953125.$$

The amplitude ratio is its inverse cube root, 0.800. Squaring the period formula with these values provides an independent check.

Derive a relation between the area action and the product of energy and period in the V trap. What fraction of the total energy is kinetic on a time average over one orbit?

▼ Work through the solution

The explicit formulas give $J = ET/(3\pi)$, a quantity with units of action. Along the orbit $p\,dq = p\dot{q}\,dt = 2K\,dt$. Hence $2\pi J = 2T\langle K\rangle$ and $\langle K\rangle = E/3$, leaving $2E/3$ in potential energy. This is an independent check using a measurable time average.

A trap is weakened gradually until the amplitude doubles. Find the predicted changes in energy and period, and state when the slow-ramp argument is most vulnerable.

▼ Work through the solution

The amplitude varies as $F^{-1/3}$, so doubling it requires reducing F to one eighth. The energy then falls to one quarter and the period increases fourfold. A ramp duration adequate initially can become too short compared to the final period. Checking the endpoint period alone is insufficient if the control has sharp changes during the ramp.

Further reading · Tong, action–angle coordinates and adiabatic invariants.

Why two finite-element wave solvers propagate the same mode at different speeds

Numerical partial differential equations

Build the element mass matrix by integrating shape functions and track the effect of diagonal lumping on the wave spectrum.

From a continuum wave to nodal masses

For $u_{tt} = c^2 u_{xx}$, multiply by a test function and integrate by parts. On a periodic line the boundary contributions cancel. A piecewise-linear element of width h uses shape functions $1-\xi, \xi$, with $\xi \in [0,1]$. Integrating products of shape functions gives

$$M_e = (h/6)\begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}.$$

Integrating derivative products gives

$$K_e = (c^2/h)\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}.$$

Row-sum lumping replaces each mass row by its sum on the corresponding diagonal, $h/2$. This preserves the mass associated with a uniform velocity but changes the response to a velocity that varies within an element.

Background · Numerical dispersion

$$c_{\mathrm{ph,h}} = \omega_h(k)/k$$

A discretised wave can retain its amplitude and still propagate at the wrong speed. Compare phase and group errors over the actual resolved wavelengths.

Diagonalise the assembled periodic problem

A periodic nodal mode has $u_j = e^{ij\theta}$. Each neighbour contributes $e^{\pm i\theta}$ relative to node j. Thus K multiplies by $2c^2(1-\cos\theta)/h$, consistent M by $h(2+\cos\theta)/3$, and lumped M by h. Dividing stiffness by mass gives the discrete squared frequency. Relative to the exact frequency $c\theta/h$,

$$\frac{\omega_L}{c\theta/h} = \frac{2\sin(\theta/2)}{\theta},$$

$$\frac{\omega_C}{c\theta/h} = \frac{2\sin(\theta/2)}{\theta}\sqrt{\frac{3}{2+\cos\theta}}.$$

Both tend to one for long waves; the leading quadratic errors have opposite signs. This is a spatial effect, present even with exact time integration.

Two wavelengths in one mesh

At four cells per wavelength $\theta = \pi/2$. The lumped ratio is $2\sqrt{2}/\pi \approx 0.9003$; the consistent ratio is about 1.1027. One wave solver lags by roughly ten per cent and the other leads by roughly ten per cent. At eight cells per wavelength the ratios move much closer to one, but repeated propagation over many wavelengths can still accumulate visible phase error.

At the shortest grid mode $\theta = \pi$, the consistent frequency is $\sqrt{12}c/h$ and the lumped frequency $2c/h$. A central-difference time integrator stable for $h_t\omega \leq 2$ therefore requires a smaller step for the consistent mass. Accuracy at one wavelength and stability at the grid scale are different design criteria.

Open this experiment online

Try the calculation

Cells per selected wavelength 8

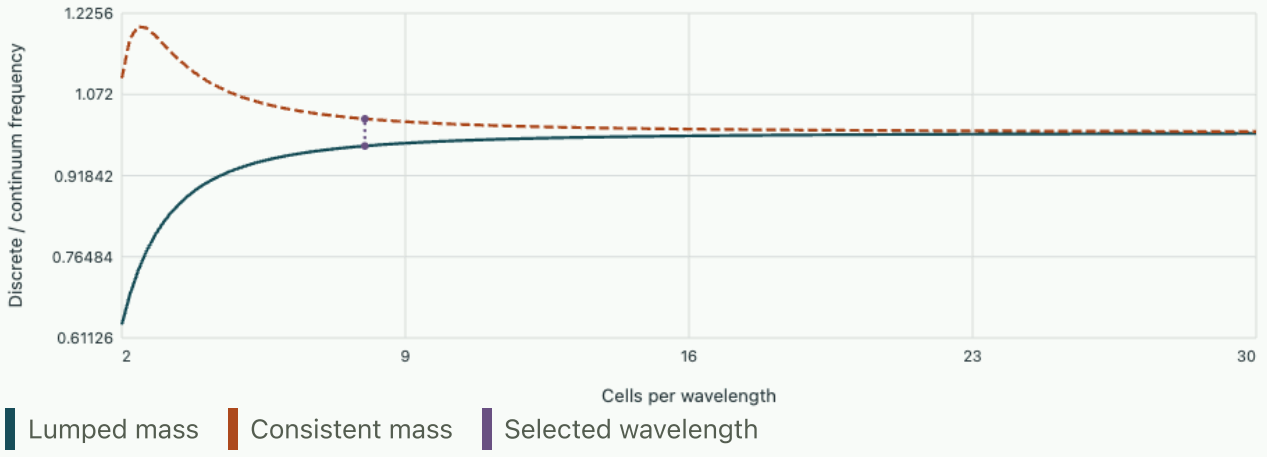

At 8 cells per wavelength the lumped ratio is 0.9745 and the consistent ratio 1.0259.

Tutorial on interpreting a convergence test

Which mass gives the exact kinetic energy for a linear velocity varying from v to $-v$ on one element?

▼ **Worked solution**

The continuum energy with unit density is $hv^2/6$. The consistent matrix reproduces it by exact integration. The lumped kinetic energy is $hv^2/2$, because it assigns both end masses the full endpoint speed.

How would you isolate spatial dispersion from time stepping in a measurement of a cavity eigenfrequency?

▼ **Worked solution**

Solve the generalised eigenproblem $Kq = \omega^2 Mq$ directly for increasingly fine meshes or converge time steps on each fixed mesh. Then compare with the continuum frequency and report both h and the time tolerance.

Noether's theorem, boundary terms and measurable constants

Analytical mechanics from constraints to invariants

A constant of motion can survive when neither laboratory energy nor angular momentum is separately conserved. Construct the conserved quantity from an infinitesimal transformation of the action. Then use boosts, a screw symmetry, and a rotating potential to determine what measured combinations should stay fixed.

Vary a path and its time labels

Let $q_i(t)$ be a path for $L(q, \dot{q}, t)$. Consider a continuous family of transformations $t' = t + \varepsilon\tau(q, t)$, $q_i'(t') = q_i(t) + \varepsilon\xi_i(q, t)$, to first order in a constant parameter. The transformation at the corresponding new time has displacement ξ_i; at the original time the path perturbation is $\eta_i = \xi_i - \dot{q}_i\tau$. Distinguishing these variations prevents a time translation from being mistaken for a vertical displacement in a configuration-versus-time plot. Differentiating the transformed coordinates gives

$$\delta\dot{q}_i/\varepsilon = \dot{\xi}_i - \dot{q}_i\dot{\tau},$$

with dots on the generator functions meaning derivatives along the path.

The transformed integration element is $dt' = (1 + \varepsilon\dot{\tau})dt$. A variational symmetry need only preserve the action up to an endpoint term. If its first-order change in the Lagrangian times the time element is $\varepsilon\, dF/dt\, dt$, the equations of motion with fixed endpoints are unchanged. Substitute the variations to write

$$\sum_i \left(L_{q_i}\xi_i + p_i\dot{\xi}_i - p_i\dot{q}_i\dot{\tau}\right) + L_t\tau + L\dot{\tau} = \frac{dF}{dt},$$

$$p_i = L_{\dot{q}_i}.$$

On a physical path $L_{q_i} = \dot{p}_i$. Combine the first two terms into $d(\sum p_i\xi_i)/dt$. Define $H = \sum p_i\dot{q}_i - L$; its derivative on a solution is

$$\dot{H} = -L_t.$$

The terms with τ then combine into $-d(H\tau)/dt$. Thus $Q = \sum_i p_i\xi_i - H\tau - F$ is conserved. Reversing the sign of the generator reverses the sign of the charge, which has no bearing on conservation. Scaling the group parameter also rescales the charge, so specifying the transformation matters when reporting its units.

Recover the familiar examples and test broken symmetry

For a common infinitesimal translation along fixed unit vector $\hat{n}$, every particle has $\xi_a = \hat{n}$, $\tau = 0$. A potential depending solely on pair separations is invariant, so $F = 0$. The conserved charge is total momentum along $\hat{n}$. For an infinitesimal rotation about the origin,

$$\xi_a = \hat{n} \times \mathbf{r}_a;$$

the scalar triple product gives $\hat{n} \cdot \sum_a \mathbf{r}_a \times \mathbf{p}_a$, the chosen angular momentum component. For time translation $\tau = 1, \xi = 0$, an explicitly time-independent Lagrangian gives $Q = -H$. In each case the

symmetry refers to the system including the prescribed external fields. Internal pair forces may conserve momentum while a fixed trap exerting an external force breaks the translation symmetry of the particles.

The algebra also gives an exact diagnostic for an approximate symmetry. If the variation of $L dt$ differs from the boundary term by $\varepsilon \mathcal{D}\, dt$, the same steps give

$$\dot{Q} = \mathcal{D}.$$

In a weak potential $V(x) = V_0 + \epsilon U(x)$, translation of x gives

$$\dot{p}_x = -\epsilon U'(x).$$

A drifting computed momentum is therefore not automatically a failure of the time integrator. Compare its change with the integrated physical symmetry-breaking force. Reversing the small external bias helps separate the first-order physical drift from a numerical offset that does not reverse.

Derive a boost charge that needs an endpoint term

Consider a free Newtonian particle in one dimension and a small Galilean boost $x' = x + \varepsilon t$. Here $\xi = t$, $\tau = 0$, and $L = m\dot{x}^2/2$ changes to first order by $m\dot{x}\varepsilon = d(mx\varepsilon)/dt$. The Lagrangian is not pointwise invariant, but its change is a total derivative. The charge is $Q = pt - mx$; its constancy is equivalent to $x - (p/m)t$ remaining equal to the initial position. Dropping the endpoint term would give pt, whose derivative is generally the nonzero momentum itself.

For mutually interacting particles with translationally invariant internal forces, summing the argument gives

$$Q = Pt - M R_{\mathrm{cm}}.$$

Constancy of total momentum and boost charge together imply that the centre of mass moves uniformly along a straight line. Take two pucks of masses 0.20 and $0.30\,\mathrm{kg}$ with $x_1(0) = 0$, $x_2(0) = 1.0\,\mathrm{m}$, and velocities 2.0 and $-0.5\,\mathrm{m/s}$. Their total momentum is $0.25\,\mathrm{kg\,m/s}$ and centre of mass starts at $0.60\,\mathrm{m}$. At $t = 0.40\,\mathrm{s}$, including any intervening isolated elastic collision, it must be at $0.80\,\mathrm{m}$. The boost charge $0.25t - 0.50R_{\mathrm{cm}}$ remains $-0.30\,\mathrm{kg\,m}$. Internal impact changes the individual momenta but preserves both total quantities.

Rotate and translate together in a helical landscape

Let a particle move in cylindrical coordinates in a potential $V(r, z - c\varphi)$, with the combination physically interpreted on a helical branch or within a periodic potential compatible with 2π rotation. Rotation through a small angle ε combined with translation $z \to z + c\varepsilon$ leaves the potential unchanged. The kinetic energy $m(\dot{r}^2 + r^2\dot{\varphi}^2 + \dot{z}^2)/2$ also stays unchanged. The generator has $\xi_\varphi = 1, \xi_z = c$, hence $Q = p_\varphi + c p_z$ is conserved. The constant c has units of length per radian, so both terms have angular-momentum units.

Differentiate directly as a check. Writing $u = z - c\varphi$, Hamilton's equations give

$$\dot{p}_\varphi = +c\, V_u,$$

$$\dot{p}_z = -V_u.$$

Each can change, yet the combination has zero derivative. A guided particle receiving a burst of axial momentum can exchange a specific amount of angular momentum with the helical potential. For $c = 0.030\,\mathrm{m}$ and $\Delta p_z = 0.080\,\mathrm{kg\,m/s}$, the relation predicts $\Delta p_\varphi = -0.0024\,\mathrm{kg\,m^2/s}$. A circularly symmetric potential would conserve angular momentum on its own and represents a different field configuration.

Exchange energy and angular momentum in a rotating pattern

A potential $V(r, \varphi - \Omega t)$ with constant pattern speed is unchanged when time increases by ε and the particle coordinates rotate by $\Omega\varepsilon$. Take $\tau = 1, \xi_\varphi = \Omega$. The conserved combination is $H - \Omega p_\varphi$, up to the sign chosen for Q. A torque from the rotating apparatus can change ordinary energy and ordinary angular momentum. Their changes obey $dH = \Omega\, dp_\varphi$ along the trajectory when the Hamiltonian contains only the prescribed rotating potential as explicit time dependence. The external power equals angular speed times the torque acting through that pattern.

Suppose $\Omega = 12\,\mathrm{rad/s}$ and a particle gains $0.015\,\mathrm{kg\,m^2/s}$ in axial angular momentum during a passage past the rotating potential. Its mechanical energy must gain $0.180\,\mathrm{J}$. A numerical run reporting $0.195\,\mathrm{J}$ calls for accounting for other time dependence, sampling uncertainty or integration error. If Ω itself slowly varies, introduce a pattern angle $\alpha(t)$, $\Omega = \dot\alpha$. One still finds $\dot H = \Omega \dot p_\varphi$, but $d(H - \Omega p_\varphi)/dt = -\dot\Omega p_\varphi$. Testing an exact rotating-frame invariant with a ramped motor would therefore diagnose the wrong experiment.

Find the conserved momentum in a magnetic field directly from the action

A charge q of mass m moves in a uniform static field $B\hat z$. Choose vector potential $A_y = Bx$, scalar potential zero. Its Lagrangian in the perpendicular plane is $L = m(\dot x^2 + \dot y^2)/2 + qBx\dot y$. The position y is cyclic, so $p_y = m\dot y + qBx$ is constant. The physical force is also invariant under translations in x, although x appears in this chosen Lagrangian. Under $x \to x + \varepsilon$, its change is $qB\dot y\,\varepsilon = d(qBy\varepsilon)/dt$. The corresponding Noether charge is $p_x - qBy = m\dot x - qBy$. Thus both spatial translation symmetries survive once the endpoint term is included.

Write $\omega_c = qB/m$, which retains the sign of the charge. Rearranging the two constants gives the coordinates of the guiding centre $X_c = x + \dot y/\omega_c,\ Y_c = y - \dot x/\omega_c$. Differentiate using $m\ddot x = qB\dot y,\ m\ddot y = -qB\dot x$ to see both are independent of time. A uniform field changes velocity direction but does no work. In a force experiment the centre of the measured circular arc, its radius and its angular orientation can be estimated without measuring any vector potential. The endpoint terms explain why two gauge choices can assign different canonical momenta while predicting the same observed guiding centre.

There is a useful algebraic surprise. Denote the translation charges by $K_x = p_x - qBy$ and $K_y = p_y$ in the chosen gauge. Their canonical bracket is $\{K_x, K_y\} = -qB$. Each is conserved, yet the bracket does not vanish. The result is a constant, so Jacobi's identity is satisfied. The mechanical effect of their infinitesimal flows on positions commutes; a central constant in the charge bracket has zero Hamiltonian vector field. Conservation, commutation of the generated coordinate motions, and zero bracket of the generating functions are related but distinct statements. This example supplies a direct check when applying symmetry counting in phase space.

For an illustrative particle with $q/m = 2.0\,\mathrm{C/kg}$, $B = 0.50\,\mathrm{T}$, initial $(x, y) = (0.20, 0.10)\,\mathrm{m}$ and velocity $(0.30, -0.40)\,\mathrm{m/s}$, the signed cyclotron frequency is $1.0\,\mathrm{s^{-1}}$. The guiding centre is $(-0.20, -0.20)\,\mathrm{m}$, and

the radius from that centre is $0.50\,\mathrm{m}$. The speed is $0.50\,\mathrm{m/s}$, consistent with radius $v/|\omega_c|$. Replacing the field by its negative at the same initial state moves the guiding centre to the opposite side of the initial velocity and reverses the sense of rotation. A constant numerical offset in the vector potential should leave both reconstructed centres unchanged.

Use scaling to decide when a formal transformation creates a charge

Many equations of motion have scaling relations. A map taking one solution into another need not preserve the action up to an endpoint term, and therefore need not supply a Noether charge. To test an example, let $L = T - V$ in three dimensions with V homogeneous of degree n, so $V(\lambda \mathbf{r}) = \lambda^n V(\mathbf{r})$. Take $\mathbf{r} \to e^{\varepsilon}\mathbf{r}$, $t \to e^{a\varepsilon}t$. Kinetic energy times dt scales as $e^{(2-a)\varepsilon}$, while potential energy times dt scales as $e^{(n+a)\varepsilon}$. Equating the exponents gives $a = (2-n)/2$, a map that rescales the entire action by a common factor and preserves stationary paths. Exact invariance of the action additionally requires the common exponent to vanish, giving $a = 2, n = -2$.

For a potential proportional to $1/r^2$, the continuous symmetry has $\xi = \mathbf{r}, \tau = 2t$, $F = 0$, and charge $D = \mathbf{r}\cdot\mathbf{p} - 2tH$. Direct differentiation gives $d(\mathbf{r}\cdot\mathbf{p})/dt = 2T - \mathbf{r}\cdot\nabla V = 2T + 2V = 2H$, verifying conservation. At fixed positive energy this particular scalar product grows linearly in time. This does not require every coordinate to grow linearly; it follows from the scaling structure of the combined kinetic and potential energies. Collision singularities may limit the interval on which a classical solution exists.

For a Newtonian potential $n = -1$, the path scaling gives $a = 3/2$, precisely the scaling of orbital periods with orbital size. The common action exponent is $1/2$, so this scaling does not produce the preceding conserved D. Instead the virial identity $d(\mathbf{r}\cdot\mathbf{p})/dt = 2T + V$ implies $2\langle T\rangle + \langle V\rangle = 0$ on a bounded periodic trajectory after time averaging. This distinction can be checked numerically on one elliptical orbit: the virial combination generally oscillates during a cycle, whereas its integrated full-cycle average vanishes. Mistaking a scaling of solutions for exact invariance of the action would turn that averaged result into a false instantaneous conservation law.

Derive a centre-of-energy statement from a boost charge

For a free relativistic particle in one spatial dimension use $L = -mc^2\sqrt{1 - \dot{x}^2/c^2}$. An infinitesimal Lorentz boost, with a parameter carrying velocity units and the sign convention $x' = x + \varepsilon t$, has generators $\xi = t, \tau = x/c^2$. Proper time is invariant, so $L dt$ is invariant and $F = 0$. The Noether charge is $K = pt - Hx/c^2$, where $H = \gamma mc^2$, $p = \gamma m\dot{x}$. Its derivative vanishes because $p = H\dot{x}/c^2$. At low speed H/c^2 approaches m, and the Newtonian boost charge reappears up to relativistic corrections.

For an isolated set of freely propagating particles away from local collisions, sum the charges to obtain

$$K = Pt - (E_{\mathrm{tot}}/c^2)X_E,$$

where $X_E = \sum E_i x_i / E_{\mathrm{tot}}$ is the centre of energy. If collisions conserve total energy and momentum and occur at the same spacetime point for their participants, the summed charge passes continuously across them. Its constancy implies

$$dX_E/dt = c^2 P/E_{\mathrm{tot}}.$$

In an interacting field problem energy stored between particles also belongs in the weighted centre; assigning every bit of energy to particle positions can violate the balance. A field-carrying interaction illustrates why the Newtonian centre of mass alone is insufficient in a relativistic energy accounting.

Imagine equal positive rest masses at $x = 0$ and $x = 2.0\,\mathrm{m}$. At the observation instant their energies are $2mc^2$ and mc^2, with the more energetic particle moving to the right and the second at rest. Their ordinary mass centre is $1.0\,\mathrm{m}$, while their energy centre is $2/3\,\mathrm{m}$. The total momentum is $\sqrt{3}\,mc$, the total energy $3mc^2$, so the centre of energy moves at $c/\sqrt{3}$. The numerical difference is not a correction to the measured coordinates but a change in the physical quantity whose motion follows a boost symmetry. A centre-of-momentum frame for the pair moves at the same speed.

Diagnose symmetry from noisy finite-time balances

In an experiment one rarely has symbolic access to every component of the Lagrangian. A practical symmetry test measures the candidate charge at two times and independently integrates the physical source expected to change it. For the rotating potential, use $R = \Delta H - \int_{t_1}^{t_2} \Omega(t)\, dL_z$. If Ω is measured on the same clock as the trajectory, the integral can be estimated by summing the mean Ω across each angular-momentum increment. This avoids taking a second time derivative of position merely to reconstruct the torque. An equivalent form obtained by parts is $\Delta(H - \Omega L_z) + \int L_z\, d\Omega$. Both versions should agree if endpoint times and signs are consistent.

Take a short programmed ramp in which Ω increases linearly from 10 to $12\,\mathrm{s^{-1}}$ and L_z from 0.030 to $0.050\,\mathrm{kg\,m^2/s}$ over the same time. If the rotating pattern is the only external work source, the energy change is mean speed times angular-momentum change,

$$11(0.020) = 0.220\,\mathrm{J}.$$

The product ΩL_z changes by $0.600 - 0.300 = 0.300\,\mathrm{J}$, so $H - \Omega L_z$ falls by $0.080\,\mathrm{J}$. The missing amount equals

$$-\int L_z\, d\Omega = -0.040(2) = -0.080\,\mathrm{J}.$$

These two balances provide a nontrivial sign test when the drive speed is varied deliberately.

When measurements are noisy, an observed discrepancy must be compared with the uncertainties of both charge and physical source. An erroneous angular zero in a central problem can affect the torque reconstruction while leaving a scalar speed unchanged; a simple time shift between the motor encoder and position camera can mimic an extra source of work. Change the drive direction, repeat with constant speed, and estimate the effect of clock offsets using resampled records. These experiments test the specific symmetry under investigation. A plot that appears almost flat on an arbitrary vertical scale supplies neither a physical tolerance nor a discrimination among competing error sources.

Derive a conservation law for an elastic field

Noether's argument also applies when the coordinates form a continuous field $u(x, t)$. Consider a uniform string or rod with density per length ρ_l, displacement field u and stiffness coefficient C, so

$\mathcal{L} = \rho_l u_t^2/2 - C u_x^2/2$. A constant shift $u \to u + \varepsilon$ changes no strain and leaves the density invariant. The Euler–Lagrange equation becomes $\partial_t(\rho_l u_t) + \partial_x(-C u_x) = 0$, a local conservation law. For an axial rod the density $\rho_l u_t$ is ordinary longitudinal momentum per unit length, and the boundary term is axial traction with its outward sign. A uniform displacement symmetry survives a free translation of the whole rod but is broken if an end displacement is imposed relative to a clamp.

Integrate from $x = 0$ to L. The total physical momentum changes as $dP/dt = C u_x(L) - C u_x(0)$. A fixed end can exert a nonzero reaction and spoil conservation of the rod's momentum even when it does no power because its actual velocity vanishes. The distinction between a momentum source and an energy source is experimental: a rigid wall reflecting a pulse can accept net impulse while taking no mechanical energy. Including the support in the system restores the corresponding balance. Checking only the field density's invariance in the interior neglects boundary conditions, which are part of a physical symmetry.

The uniformity of material parameters also gives an invariance under translations of the material coordinate x in an infinite or periodic homogeneous domain. The associated density is proportional to $-\rho_l u_t u_x$, often called wave pseudomomentum under this convention. Derive its local balance by multiplying $\rho_l u_{tt} = C u_{xx}$ by $-u_x$ and applying the product rule. One obtains

$$\partial_t(-\rho_l u_t u_x) + \partial_x[C u_x^2/2 + \rho_l u_t^2/2] = 0.$$

This conserved quantity differs from the ordinary momentum found from shifting u itself. A small-amplitude sinusoidal travelling wave can have zero mean material velocity over one period yet positive mean transport of this wave quantity.

For $u = A\cos(kx - \omega t)$,

$$u_t = \omega A \sin(kx - \omega t),$$
$$u_x = -kA \sin(kx - \omega t).$$

The cycle-averaged wave density is $\rho_l \omega k A^2/2$ while ordinary momentum density averages to zero. With $\rho_l = 0.20\,\mathrm{kg/m}$, $A = 1.0\,\mathrm{mm}$, $k = 5.0\,\mathrm{m^{-1}}$, and $\omega = 100\,\mathrm{s^{-1}}$, the average wave density is $5.0 \times 10^{-5}\,\mathrm{kg/s}$. Its units are momentum per unit length, but equal dimensions do not make the two conserved charges interchangeable. To identify the charge, specify whether the symmetry shifts every displacement equally or shifts the material-coordinate pattern.

Test the evolution of a phase-space patch

Follow a family of initial conditions in a Hamiltonian flow. Compare changes in its shape with the conservation of its phase-space area.

Open this experiment online

Try the calculation

Time after release (s) **4**

Drag rate γ (1/s) **0.2**

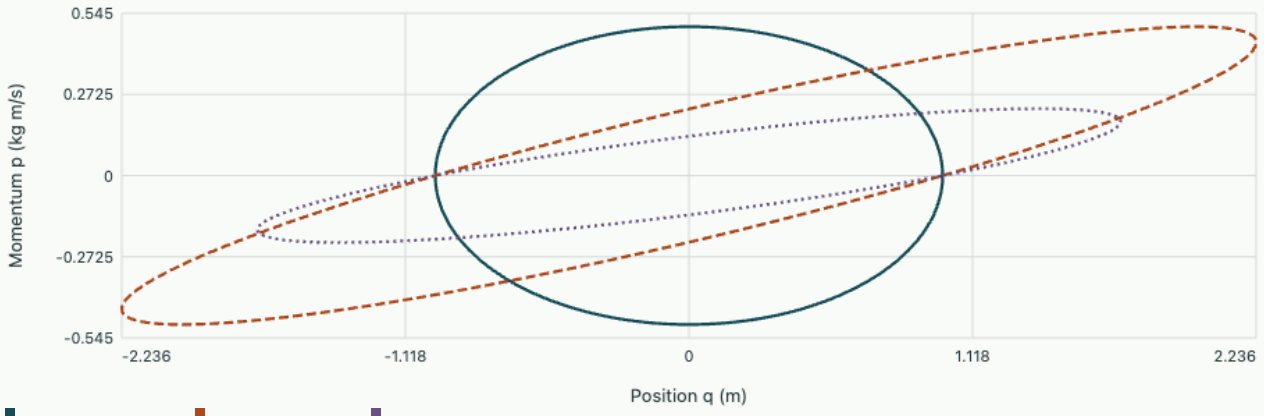

Initial cloud | Free flight | With drag

Free position spread 2.236 m. The damped ellipse retains 44.93% of its initial area and has position spread 1.702 m.

Derive an identity without assuming the equations of motion

The usual on-shell proof substitutes the Euler–Lagrange equations just before a total derivative appears. Retaining the difference gives an equally useful off-shell identity. Define the residual $R_i = dp_i/dt - \partial L/\partial q_i$. For a true variational symmetry and any differentiable trial path, repeat the calculation with $\eta_i = \xi_i - \dot{q}_i \tau$. The result is

$$dQ/dt = \sum_i R_i \eta_i.$$

Thus the drift of a Noether charge evaluated on an approximate path is a particular projection of its dynamical residual. A path can have zero drift of one charge while violating equations in directions to which that charge is insensitive.

For a free particle in two dimensions, perfect conservation of p_x enforces zero residual acceleration in the x direction but says nothing about a spurious force in y. Conversely tracking two independent translational charges checks both components. Energy conservation tests a different projection, $dH/dt = \dot{\mathbf{q}} \cdot \mathbf{R}$ for a time-independent Lagrangian. An error force always perpendicular to velocity can leave the speed and energy unchanged while bending an originally straight trajectory. That is one reason an invariant energy graph cannot replace a force residual or a comparison of trajectories.

A numerical planar free particle follows a circle of constant speed even though the programmed potential vanishes. Its energy is conserved exactly. Which Noether quantities expose the problem most directly?

▼ **Work through the solution**

The two laboratory translation momenta should each be constant. On the reported circle their vector changes direction continuously, so their residuals immediately reveal the missing transverse equation. A check of only the magnitude of momentum would also pass because that magnitude is fixed by energy. The angular momentum about the circle centre can remain constant too. Testing a set of charges selected for the actual symmetry and measured geometry identifies information lost by each scalar test on its own.

If the Lagrangian acquires a small explicitly time-dependent conservative potential $V_1(t)q^2$, what quantity should replace a flat-energy test?

▼ **Work through the solution**

The energy change on a physical path is $\Delta H = \int \dot{V}_1(t) q(t)^2 \, dt$. Positive or negative changes are allowed according to how the externally imposed stiffness is varied during the motion. Compare the measured change with this integral while keeping the time dependence inside the Hamiltonian itself. A flat-energy requirement would incorrectly penalise physical work by the control, whereas a residual after the input integral tests the dynamics.

Separate a discrete symmetry from a continuous one

A potential $V(q) = a(q^2 - b^2)^2$ is invariant under the discrete reflection $q \mapsto -q$. Reflecting both position and velocity maps any solution into another solution with equal energy. There is no continuously tunable parameter that takes the identity smoothly to this isolated reflection while preserving the same one-dimensional potential at every intermediate step. No infinitesimal translation generator follows from this symmetry, and the momentum generally changes along either path. A pair of measured trajectories launched from mirrored initial data should match after reflection, but neither trajectory is required to have zero or constant momentum.

For $a = 1$ and $b = 1$ in consistent scaled units, release particles from $q = 0.50$ and -0.50 at rest. The force is $-4q(q^2 - 1)$, giving initial accelerations $+1.50/m$ and $-1.50/m$. The two accelerations are mirror images, not individually zero. A numerical method that accidentally introduces a constant bias force breaks the reflection comparison and shifts the two turning times in opposite senses. Energy can remain nearly conserved with respect to a slightly tilted effective potential, so trajectory pairing is a complementary diagnostic. Reflection symmetry imposes a relationship between solutions even when Noether's continuous-symmetry construction supplies no new conserved scalar.

Tutorial on identifying symmetries

A Lagrangian is replaced by $L' = L + dG(q,t)/dt$. Does the numerical value of an appropriately derived Noether charge have to change?

▼ **Work through the solution**

The new momentum is $p_i' = p_i + G_{q_i}$, and $H' = H - G_t$. Under the given symmetry the new boundary term is $F' = F + \sum_i G_{q_i}\xi_i + G_t\tau$. Substitution into $p'\xi - H'\tau - F'$ cancels the extra terms exactly. Charges can still be shifted by a physically irrelevant additive constant through the choice of boundary term.

Comparing raw canonical momentum in two such conventions without also transforming F, H confuses intermediate variables with observations.

A central potential is changed into an axisymmetric potential $V(r, \theta)$. Which angular-momentum tests remain exact?

▼ Work through the solution

Rotation about the symmetry axis remains a continuous symmetry, so its angular-momentum component is conserved. Tilt rotations generally change the potential, so the transverse components need not stay constant. The torque may cause the angular momentum direction or magnitude to evolve while maintaining its axial component. A code test that requires all components to remain fixed would reject many legitimate axisymmetric trajectories.

A table exerts an unmodelled approximately constant horizontal drag $0.010\,\mathrm{N}$ opposite the motion of the two pucks for $0.40\,\mathrm{s}$. Estimate how the translation and boost checks change.

▼ Work through the solution

Total momentum changes by $-0.0040\,\mathrm{kg\,m/s}$. The derivative of $Pt - MR_{\mathrm{cm}}$ is $t\dot{P}$, since $P = M\dot{R}_{\mathrm{cm}}$. Its change is therefore $-0.010(0.40)^2/2 = -0.00080\,\mathrm{kg\,m}$ under constant drag. The differing time dependence gives an additional way to distinguish a sustained external force from an initial position calibration error.

Follow the action of an infinitesimal generator on phase space.

David Tong treats continuous symmetries and Noether's theorem in Lagrangian dynamics.

Kepler orbits, the eccentricity vector and the loss of apsidal closure

Analytical mechanics from constraints to invariants

An orbit can conserve energy and angular momentum yet gradually rotate its periapsis. The inverse-square problem possesses one more conserved vector that fixes the direction of closest approach. Derive it, turn a single position and velocity measurement into a conic, and measure what a weak perturbation does to the inferred apsidal direction.

Reduce the two-body system without freezing either mass

Two bodies of masses m_1, m_2 have positions $\mathbf{r}_1, \mathbf{r}_2$. Define total mass $M = m_1 + m_2$, centre $\mathbf{R} = (m_1\mathbf{r}_1 + m_2\mathbf{r}_2)/M$, relative coordinate $\mathbf{r} = \mathbf{r}_1 - \mathbf{r}_2$, and reduced mass $\mu = m_1 m_2/M$. Substitution into the kinetic energy cancels the cross terms, giving $M\dot{\mathbf{R}}^2/2 + \mu\dot{\mathbf{r}}^2/2$. Newtonian gravitational potential is $-\kappa/r$ with $\kappa = Gm_1m_2$. The centre moves uniformly when isolated, while the relative motion obeys $\mu\ddot{\mathbf{r}} = -\kappa\hat{r}/r^2$. Dividing by the reduced mass gives the acceleration parameter $k = \kappa/\mu = GM$; replacing it by Gm_1 is the test-particle approximation.

Use specific energy $\epsilon = v^2/2 - k/r$ and specific angular momentum $\mathbf{h} = \mathbf{r} \times \mathbf{v}$, both per unit reduced mass. A central force makes $\dot{\mathbf{h}} = 0$, fixing the plane of nonradial motion. The area of the triangle swept in time dt is $r^2 d\varphi/2 = h\,dt/2$, yielding the constant areal rate. The radial energy is $\dot{r}^2/2 + h^2/(2r^2) - k/r$. Its turning points determine the accessible radii, but energy and h alone cannot specify where the major axis lies in the chosen inertial plane.

Construct a vector whose derivative vanishes

The derivative of $\mathbf{v} \times \mathbf{h}$ is $(-k\hat{r}/r^2) \times \mathbf{h}$. Use $\mathbf{h} = r^2\dot{\varphi}\,\hat{z}$ with an oriented plane in which motion is positive. Since $\hat{r} \times \hat{z} = -\hat{\varphi}$, the derivative becomes

$$k\dot{\varphi}\hat{\varphi} = k\,d\hat{r}/dt.$$

Consequently the eccentricity vector

$$\mathbf{e} = \frac{\mathbf{v} \times \mathbf{h}}{k} - \hat{r}$$

is constant. The dimensional Runge–Lenz vector is proportional to this choice: for relative momentum $\mathbf{p} = \mu\mathbf{v}$,

$$\mathbf{A} = \mathbf{p} \times \mathbf{L} - \mu\kappa\hat{r} = \mu\kappa\mathbf{e}.$$

Every term in the specific expression is dimensionless. The step equating the derivative to $k\,d\hat{r}/dt$ uses exactly the inverse-square radial force. General central forces retain areal velocity but lose this cancellation.

Dot $\mathbf{e}$ with $\hat{r}$. The triple product $(\mathbf{v} \times \mathbf{h}) \cdot \hat{r} = h^2/r$ gives

$$e\cos(\varphi - \varpi) = h^2/(kr) - 1,$$

where ϖ is the fixed direction of $\mathbf{e}$. Thus $r = p/[1 + e\cos(\varphi - \varpi)]$, $p = h^2/k$, the conic in focus-centred polar coordinates. Squaring $\mathbf{e}$ and eliminating v^2 with the energy gives

$$e^2 = 1 + 2\epsilon h^2/k^2.$$

Bound ellipses have $\epsilon < 0, e < 1$, parabolic escape has $\epsilon = 0, e = 1$, and positive energy gives hyperbolas for nonzero angular momentum. Radial trajectories require separate care because their orbital plane and apsidal angle are degenerate.

Recover the orbit from a single state vector

In units chosen so $k = 1$, a measured relative position is $(2, 0, 0)$ and velocity is $(0.2, 0.6, 0)$. Compute $h_z = 1.2$, $\epsilon = (0.04 + 0.36)/2 - 1/2 = -0.30$, and $p = 1.44$. The cross product $\mathbf{v} \times \mathbf{h} = (0.72, -0.24, 0)$ gives $\mathbf{e} = (-0.28, -0.24, 0)$, with $e = \sqrt{0.136} \approx 0.369$. The direction of periapsis is about 220.6° from the positive x axis. At the measured point $\dot r > 0$, so the body is travelling from periapsis toward apoapsis. Dropping the radial component of velocity would rotate the eccentricity vector and change the inferred shape.

The semimajor axis is

$$a = -k/(2\epsilon) = 1.667.$$

Periapsis and apoapsis radii $a(1 - e)$ and $a(1 + e)$ are approximately 1.05 and 2.28, consistent with the observed radius of two. The orbital period is $2\pi\sqrt{a^3/k} = 13.52$ time units. Derive this last result by dividing ellipse area πab by the areal speed $h/2$. Using $b = a\sqrt{1 - e^2}$ and $h^2 = ka(1 - e^2)$ cancels eccentricity exactly. This shows why changing the initial radial velocity can change shape and orientation while the period depends only on binding energy.

Determine time along an ellipse and check limiting cases

Introduce the eccentric anomaly E by $x = a(\cos E - e)$, $y = a\sqrt{1 - e^2}\sin E$ in axes centred on the focus, with x pointing toward periapsis. Then $r = a(1 - e\cos E)$. Compute $x\,dy - y\,dx = ab(1 - e\cos E)\,dE$ and equate it to $h\,dt$. Integrating from periapsis gives

$$n(t - t_p) = E - e\sin E,$$

where $n = \sqrt{k/a^3}$. This relation turns orbital geometry into an observation clock. The true anomaly is generally different from the linearly advancing mean anomaly because the body sweeps angle faster at small radius.

For $e = 0.6$ and the time $t - t_p = T/4$, the mean anomaly is $\pi/2$. Newton iteration for $F(E) = E - 0.6\sin E - \pi/2$ uses denominator $1 - 0.6\cos E$, which stays positive. Starting near $E = 2.0$ yields $E \simeq 2.091$ radians and $r \simeq 1.30a$. The orbit has already moved beyond its semimajor radius at a quarter period. At $e = 0$, E equals the mean and true anomalies up to an arbitrary origin, and $\mathbf{e} = 0$ has no definable periapsis direction. Estimating a periapsis angle for a nearly circular noisy orbit is therefore poorly conditioned even when the radius and period are well measured.

Perturb the force and measure apsidal precession

For any central acceleration $f(r)\hat{r}$, put $u = 1/r$. Since $\dot{\varphi} = hu^2$, radial velocity is $-h\,du/d\varphi$. Substituting into the radial acceleration gives Binet's equation $u'' + u = -f(1/u)/(h^2u^2)$, where primes denote angular derivatives. Choose $f(r) = -k/r^2 - \alpha/r^3$. Then $u'' + (1 - \alpha/h^2)u = k/h^2$. If $0 < \alpha < h^2$, a radial oscillation requires angular advance $2\pi/\sqrt{1 - \alpha/h^2}$, greater than one full turn. Successive periapses advance by that value minus 2π, approximately $\pi\alpha/h^2$ for a weak perturbation.

With $\alpha/h^2 = 0.010$, the exact precession per radial cycle is approximately $0.03165\,\mathrm{rad} = 1.81^\circ$. Ordinary energy and angular momentum remain constant in this conservative central potential. The Kepler eccentricity vector evaluated from instantaneous position and velocity now changes. Track periapses using zeros of the radial velocity with positive radial acceleration, then compare their inertial angles. A numerical Kepler integrator can precess even for $\alpha = 0$; refine the time step and repeat the unperturbed control before assigning a small precession to the physical perturbation.

Infer an orbit using quantities that remain well conditioned

Suppose positions and velocities are measured with finite precision. A circular orbit has exact eccentricity vector zero, so even tiny perpendicular perturbations can swing the direction of that vector through large angles. Keep its Cartesian components when computing a fit near circularity. For an initially transverse circular state at radius r and speed $v_c = \sqrt{k/r}$, perturb the velocity by radial δv_r and tangential δv_t. In the local $(\hat{r}, \hat{\varphi})$ basis the eccentricity components are $rv_t^2/k - 1$ and $-rv_rv_t/k$. At first order they become $2\delta v_t/v_c$ and $-\delta v_r/v_c$. This linear response is well behaved even though taking an inverse tangent to quote periapsis longitude can become meaningless.

For a circular speed of $10.0\,\mathrm{m/s}$, radial and tangential measurement errors each of $0.010\,\mathrm{m/s}$ correspond to component uncertainties about 0.0010 and 0.0020. A reported eccentricity of 0.0012 is not accompanied by a precise preferred direction under those uncertainties. Repeated orbital samples reduce independent noise but do not remove a persistent tangential velocity bias. Measuring the radius and period independently supplies an additional constraint on k and the orbital energy. A model that regards the primary as fixed also introduces a systematic mass error when the secondary has non-negligible mass.

At the opposite extreme, a near-parabolic orbit can have a well-defined eccentricity direction but a badly conditioned energy estimate. Its energy is the small difference of $v^2/2$ and k/r. Since $a = -k/(2\epsilon)$, a small energy uncertainty produces a very large uncertainty in semimajor axis and period. Report the measured state vector, eccentricity vector and energy uncertainty before interpreting an enormous best-fit period. If the energy uncertainty straddles zero, bounded and unbounded conics are both consistent with the measured data. Integrating a single best-fit ellipse for many revolutions does not resolve that ambiguity.

Calculate the velocity impulse to transfer between circular radii

The same two-body invariants reconstruct a transfer between circular trajectories without requiring a time-stepped simulation. Set a lower radius r_1 and an upper radius $r_2 > r_1$. An ellipse tangent to both has periapsis r_1, apoapsis r_2, semimajor axis $a_t = (r_1 + r_2)/2$, and eccentricity $(r_2 - r_1)/(r_2 + r_1)$. At either apsis its speed follows from energy:

$$v_t(r) = \sqrt{k(2/r - 1/a_t)}.$$

A first instantaneous tangential velocity change moves the particle from its lower circular trajectory to the transfer ellipse. A second tangential change at the opposite apsis matches the upper circular speed. These formulas describe impulses idealised as short compared with the orbital period.

Choose normalised $k = r_1 = 1$, $r_2 = 2$. Initial speed is one, transfer semimajor axis 1.5, and periapsis speed $\sqrt{4/3} = 1.1547$. The first change is 0.1547. At the upper radius the transfer speed is $\sqrt{1/3} = 0.57735$, while the circular speed there is

$$1/\sqrt{2} = 0.70711.$$

The second change must also be forward, by 0.12976. The flight time between the two impulses is $\pi\sqrt{1.5^3} = 5.77$ time units, half the transfer period. Simply increasing the initial speed until the path reaches radius two does not make it circular upon arrival. Matching position and angular momentum at the final boundary requires the second change.

Verify the energy increment from an impulse independently. A change $\Delta\mathbf{v}$ at fixed position gives

$$\Delta\epsilon = \mathbf{v}\cdot\Delta\mathbf{v} + |\Delta\mathbf{v}|^2/2.$$

At the lower radius our impulse raises ϵ from -0.5 to $-1/3$, an increase of $1/6$. Substitution of the rounded 0.1547 gives the same value. The same velocity increment generally changes energy more when applied where the initial speed is larger and aligned with the impulse. This is a statement about mechanical work at different initial states, not an exception to conservation between impulses.

A laboratory analogue uses a central potential engineered over a limited radius range. A small finite-duration tangential push changes radius as well as speed during the push. To check the impulse approximation compute the change in angular momentum as $\int \mathbf{r}\times\mathbf{a}_{\text{applied}}\,dt$ and energy as $\int \mathbf{v}\cdot\mathbf{a}_{\text{applied}}\,dt$. Compare these with the instantaneous predictions and shorten the pulse while holding its integrated velocity change fixed. If the return radius still drifts, map the force against radius to distinguish a finite-pulse error from an actual deviation from inverse-square attraction.

Continue the time law across the escape boundary

A hyperbolic Kepler orbit can be described without calling a divergent ellipse period a clock. Define a positive hyperbolic length $a_h = k/(2\epsilon)$ for $\epsilon > 0$, and let $e > 1$. With periapsis on the positive x axis use $x = a_h(e - \cosh H)$,

$$y = a_h\sqrt{e^2 - 1}\sinh H.$$

Then $r = a_h(e\cosh H - 1)$, and the swept-area relation gives

$$t - t_p = \sqrt{a_h^3/k}(e\sinh H - H).$$

The hyperbolic anomaly H increases monotonically from negative to positive along the passage. At large positive H the ratio dr/dt approaches $\sqrt{k/a_h} = \sqrt{2\epsilon}$, the asymptotic speed predicted by the conserved energy.

At the parabolic boundary $e = 1, \epsilon = 0$, put $D = \tan(\varphi/2)$ with φ measured from periapsis. If q_p is the periapsis distance, the semilatus rectum is $p = 2q_p$ and $r = q_p(1 + D^2)$. Since $dt = r^2 d\varphi/h$, $d\varphi = 2dD/(1 + D^2)$ and $h = \sqrt{2kq_p}$, integration gives Barker's relation $t - t_p = \sqrt{2q_p^3/k}(D + D^3/3)$. This expression is finite on either side of the periapsis for finite radius, even as the time to reach infinity diverges. A parabolic passage therefore has an entirely useful local time-of-flight law despite having no period.

For $k = 1$, $q_p = 1$ in scaled units, a parabolic trajectory reaches $\varphi = 90^\circ$ at $D = 1,\ r = 2$. Its elapsed time is

$$4\sqrt{2}/3 = 1.886.$$

Using the instantaneous angular speed at periapsis for the entire quarter turn would predict $(\pi/2)/\sqrt{2} = 1.111$, much too short. The areal speed stays constant, but angular speed decreases by a factor of four by the time radius has doubled. Integrating $r^2 d\varphi/h$ along the conic captures that slowdown directly.

Near zero energy, formulas involving a semimajor axis of enormous magnitude can lose numerical accuracy through cancellation. For example $r = a(1 - e\cos E)$ involves subtracting two near-equal terms near periapsis in a high-eccentricity ellipse. Expressions in $q_p = a(1 - e)$ and half-angle quantities, such as $r = q_p + 2ae\sin^2(E/2)$, avoid one damaging subtraction. A state-based numerical propagation should preserve angular momentum, compare the periapsis passage time with a quadrature, and test both small positive and small negative energies. Agreement only in the well-bound ellipse regime is a weak test of an algorithm used for near-escape trajectories.

Relate the Runge–Lenz direction to a velocity-space circle

The spatial Kepler conic has a simpler partner in velocity space. In the orbital plane,

$$d\mathbf{v}/d\varphi = \mathbf{a}/\dot{\varphi} = (-k\hat{r}/r^2)/(h/r^2) = -(k/h)\hat{r}.$$

Since $d\hat{\varphi}/d\varphi = -\hat{r}$, integration gives

$$\mathbf{v} = (k/h)\hat{\varphi} + \mathbf{v}_c,$$

where $\mathbf{v}_c$ is a constant vector. Thus the velocity vector traces a circle of radius k/h, or part of that circle on an unbound orbit. At periapsis a prograde orbit has speed $(k/h)(1 + e)$, so the circle centre is shifted in the prograde tangential direction by ke/h. The offset direction is perpendicular to the eccentricity vector.

One can recover k/h and e from a set of calibrated velocity samples by fitting a circle in the orbital plane and comparing its radius with the offset of its centre from zero velocity. This method still needs the correct inertial frame: adding a constant observer velocity translates every velocity-space point, moving the fitted centre while preserving the circle radius. A time sequence of relative velocities between the two masses removes a

common boost. If the force departs from inverse square, the derivative $d\mathbf{v}/d\varphi$ acquires radius dependence and the velocity circle generally distorts. Simultaneously checking the spatial conic and the velocity circle makes a more discriminating test than plotting either shape alone.

Derive which central orbits remain nearly closed

A circular orbit at radius r_c in a smooth potential per mass $U(r)$ has $h^2 = r_c^3 U'(r_c)$. Its angular frequency obeys $\Omega^2 = U'(r_c)/r_c$. Small radial deviations at fixed h see the effective potential $U + h^2/(2r^2)$, so their squared epicyclic frequency is

$$\kappa_r^2 = U''(r_c) + 3U'(r_c)/r_c.$$

During one radial oscillation the orbital angle advances roughly $2\pi\Omega/\kappa_r$. If the two frequencies differ, the apsidal orientation generally rotates, even though energy and angular momentum remain constants. This local result distinguishes near-circular closure from the much stronger claim that all bounded orbits close for a given force law.

For an attractive radial acceleration $-kr^{-s}$, $U' = kr^{-s}$ and $\kappa_r^2 = (3-s)\Omega^2$. A stable circular orbit requires $s < 3$. For inverse square, $s = 2$, radial and angular frequencies agree, so the apsis returns to its starting orientation each revolution. For an isotropic harmonic potential the force is proportional to r, so $s = -1$ and $\kappa_r = 2\Omega$. There are two radial minima per full turn; its closed ellipse is centred on the force centre rather than focused there. A rational near-circular ratio at one energy or radius does not by itself prove global closure for all eccentricities. The Runge–Lenz derivation supplied much stronger information for the inverse-square case.

Suppose a weak field calibration drift changes the local logarithmic force slope from $s = 2$ to $s = 1.98$, approximately constant over a narrow radial range. The near-circular radial frequency is 1.00995Ω. Consecutive closest approaches are separated by $2\pi/\sqrt{1.02} \approx 6.221\,\mathrm{rad}$, giving a retrograde change about $0.062\,\mathrm{rad}$ per radial period. The sign is opposite the advance induced by the additional attractive inverse-cube contribution considered before. Estimating the radial oscillation frequency and orbital frequency independently can therefore diagnose deviations that an energy test would miss. At larger eccentricity the sampled force covers a broader radial interval and the local-power approximation must be replaced by the full turning-point integral.

Reconstruct a Kepler ellipse from one state

Vary initial radial and tangential velocity separately. Infer the direction of the eccentricity vector and compare both turning radii with the energy equation.

Open this experiment online

Try the calculation

Initial radial velocity **0.1**

Initial tangential velocity **0.9**

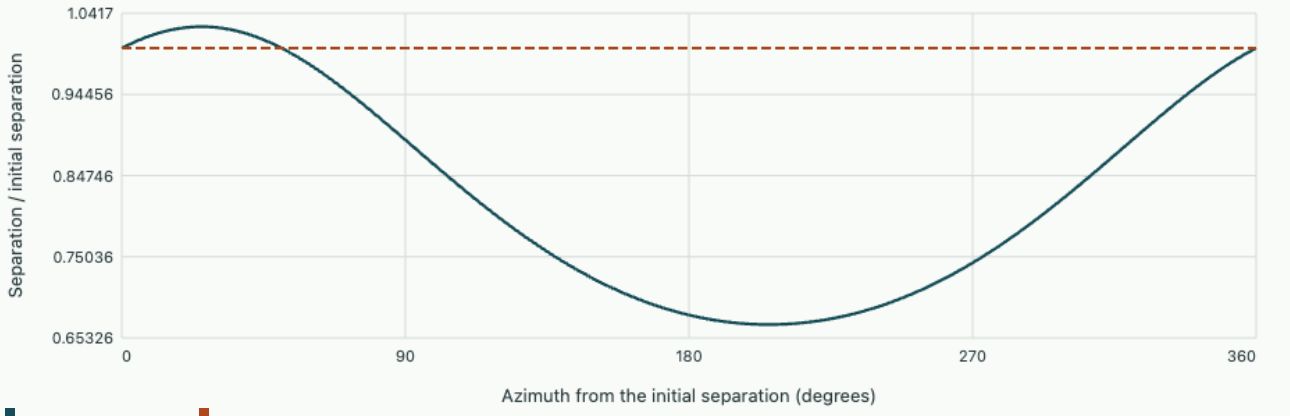

Bound orbit Initial separation

Eccentricity 0.21024. Closest approach 0.66929; largest separation 1.0256.

Use conservation of the eccentricity vector to test an orbital integrator

A code integrating $\ddot{\mathbf{r}} = -k\mathbf{r}/r^3$ can be tested on a moderate ellipse where both apses are well resolved. Choose $k = 1$, semimajor axis $a = 1$, eccentricity $e = 0.6$, and start at periapsis $\mathbf{r} = (0.4, 0)$, $\mathbf{v} = (0, 2.0)$. Then $h = 0.8$, $\epsilon = -0.5$, and the eccentricity vector is exactly $(0.6, 0)$. The apoapsis must lie at $(-1.6, 0)$ with speed 0.5, and the return time is 2π. This one initial condition tests two-dimensional direction, speed, position and clock against independent invariant predictions.

A kick–drift method for a central force can preserve angular momentum exactly at each substep. A radial kick changes velocity along position and gives zero torque; a free drift changes position along velocity and likewise preserves the cross product. Such a method can nevertheless rotate the eccentricity vector slowly or distort the time of periapsis. A symmetric velocity-Verlet step often produces bounded energy errors but still a systematic numerical periapsis shift. Record the eccentricity vector just after each interpolated periapsis crossing over a fixed number of physical periods and compare its direction as the step is halved. Do not locate a closest approach merely by selecting the nearest stored time sample if that timing error is comparable to the shift under study.

A high-eccentricity orbit requires a smaller near-periapsis time step than its full period alone suggests. At the example periapsis, a local bend time is $r_p/v_p = 0.20$ in the scaled units, about one thirtieth of the period. For $e = 0.95$ at the same semimajor axis it becomes $0.05/\sqrt{39} \approx 0.0080$, almost eight hundred times smaller than the period. A time step that produces a smooth-looking wide ellipse at $e = 0.6$ may skip much of the rapid turning near $e = 0.95$. Refine according to the fastest relevant local dynamics as well as the final integration length.

How could a simulation distinguish a real apsidal advance from an error in fitting the coordinate origin?

▼ **Work through the solution**

Translate the calculated orbit to a candidate force centre and evaluate both the radial force law and the eccentricity vector at many positions. A constant wrong origin changes the measured radius and radial direction in a phase-correlated way even within a single orbit. A physical perturbation about the true centre can preserve zero torque and produce a gradual cycle-to-cycle change. Check angular momentum about the candidate centre, the direction of the measured acceleration, and convergence under grid or time-step refinement. Several diagnostics are needed because a short arc alone may admit many fitted centres and apsidal directions.

Determine the total mass from a visual binary

If the two bodies can both be resolved, measure their relative semimajor axis a and period T. Kepler's area relation gives $M = 4\pi^2 a^3/(GT^2)$, the sum of masses. Motion about the common centre divides the same relative semimajor axis in inverse proportion to the masses: $a_1/a_2 = m_2/m_1$. For a visual orbit with $a_1/a_2 = 0.50$ and a dynamically inferred total of three mass units, the bodies have masses two and one in those units. An orbit of only one luminous body around an unobserved centre normally supplies fewer independent constraints.

A projected semimajor axis and a measured angular period require additional geometric information before applying the mass law. The apparent ellipse on the sky can be shortened by inclination and the focus of the spatial ellipse does not in general map to a focus of the apparent ellipse. Distances converting an angular semimajor axis to metres enter cubed: an independent distance bias of two percent produces approximately a six-percent mass bias at fixed angular scale and period. This makes the gravitational invariant a sensitive quantitative measurement rather than simply a qualitative statement that the relative trajectory looks elliptical.

Tutorial on orbital inference

A body passes periapsis of an ellipse at r_p with purely transverse speed v_p. Express its eccentricity in terms of these measurements. Which ambiguity occurs if the apsis has not been identified?

▼ **Work through the solution**

At an apsis the eccentricity vector is $(rv^2/k - 1)\hat{r}$. At periapsis it points outward along $\hat{r}$, giving $e = r_p v_p^2/k - 1$. At apoapsis the coefficient is negative and its magnitude is the eccentricity. A purely transverse velocity identifies a turning point of radius but does not by itself distinguish which apsis without comparing v with the circular value $\sqrt{k/r}$ or measuring radial acceleration.

At fixed orbital radius the speed is increased tangentially from circular speed to exactly $\sqrt{2}$ times circular speed. What has changed in a, e and the apsidal interpretation?

▼ **Work through the solution**

The specific kinetic energy now equals k/r, giving total energy zero. The resulting orbit is parabolic with $e = 1$ and no finite semimajor axis or orbital period. The launch point is periapsis since radial speed is initially zero and the transverse speed exceeds the circular value. Taking the ellipse period formula with a very large fitted a near this energy boundary makes tiny measurement errors produce enormous period uncertainties.

Two solutions report nearly identical energy histories but their eccentricity vectors rotate in opposite directions as time advances. What can a controlled comparison establish?

▼ Work through the solution

Compute angular momentum and eccentricity vector from positions and velocities using one common force parameter and frame. Run both at decreasing time steps over a fixed physical number of orbits, with the same periapsis event interpolation. If the rotation vanishes under refinement in the inverse-square problem, the original precession was a discretisation effect. If a prescribed conservative non-Kepler term remains present, a finite shift can survive with energy conserved. Energy conservation is one scalar condition and cannot determine the apsidal direction.

Test a state vector and compare the predicted turning radii.

Torsten Fließbach's mechanics course includes dynamics in central potentials.

Symplectic transformations, Poisson brackets and the geometry of measurements

Analytical mechanics from constraints to invariants

Two phase-space plots can enclose the same volume while predicting different responses to an impulse. Derive the stronger condition that makes a change of variables canonical, test it on measured linear maps and use Poisson brackets to propagate observables without solving separately for every coordinate.

Derive the matrix condition from Hamilton's equations

Collect n canonical positions and momenta into $z = (q_1, \dots, q_n, p_1, \dots, p_n)^T$. Write Hamilton's equations as $\dot{z} = J\nabla_z H$, where J has zero diagonal blocks, identity upper right and minus identity lower left. For a time-independent change $Z = f(z)$ with Jacobian $S = \partial Z/\partial z$, the chain rule gives

$$\dot{Z} = SJS^T \nabla_Z H.$$

To preserve the canonical form for every differentiable Hamiltonian require

$$SJS^T = J.$$

For an invertible square matrix this is equivalent to $S^T J S = J$. In geometric terms the bilinear antisymmetric form $\omega(u, v) = u^T J v$ has the same value on any transported pair of tangent vectors.

Taking determinants gives $(\det S)^2 = 1$, and symplectic maps have positive determinant equal to one. One way to see the sign is to take the top exterior power of the preserved two-form: ω^n fixes an oriented volume. In one canonical pair the identity $S^T J S = (\det S) J$ makes determinant one sufficient. In two or more pairs a single volume condition cannot enforce every entry of the matrix condition. For example in ordering (q_1, q_2, p_1, p_2), scaling by diagonal entries $(2, 1, 1, 1/2)$ preserves four-dimensional volume. But $\{Q_1, P_1\} = 2$ and $\{Q_2, P_2\} = 1/2$, so neither conjugate pair has the required bracket. A volume-preserving histogram test alone would miss this failure.

Identify brackets of observables and what they measure

Define $\{A, B\} = (\nabla A)^T J \nabla B = \sum_i (A_{q_i} B_{p_i} - A_{p_i} B_{q_i})$. For any explicitly time-dependent observable $A(z, t)$, differentiation along Hamiltonian trajectories gives

$$\dot{A} = \{A, H\} + \partial_t A.$$

The bracket is antisymmetric and obeys the product rule. Computing it with the chain rule in the new variables proves that canonical transformations preserve brackets, even between nonlinear observables. Conversely the fundamental brackets $\{Q_i, P_j\} = \delta_{ij}$, $\{Q_i, Q_j\} = \{P_i, P_j\} = 0$ are precisely the entries of $SJS^T = J$.

For a particle set $L_z = xp_y - yp_x$. Its Hamiltonian flow as a generator sends x to $x - \varepsilon y$, y to $y + \varepsilon x$ and rotates momenta by the same infinitesimal angle. The angular momentum components satisfy $\{L_x, L_y\} = L_z$ with cyclic permutations. Derive one entry directly for $L_x = yp_z - zp_y$ and $L_y = zp_x - xp_z$. The z, p_z pair contributes $xp_y - yp_x$, with the other terms zero. Rotations about two different axes generally

fail to commute. The resulting commutator has the form of a rotation about the third axis even though each underlying infinitesimal motion preserves every canonical bracket.

Jacobi's identity $\{A,\{B,C\}\}+\{B,\{C,A\}\}+\{C,\{A,B\}\}=0$ implies that the bracket of two time-independent conserved quantities is conserved under the same Hamiltonian. If the Hamiltonian is invariant under rotations about x and y, its invariance about z follows from this relation and $\{L_x,L_y\}=L_z$. Conservation alone need not make two observables commute. A central force can conserve all three angular momentum components, but their mutual brackets remain nonzero. For independent commuting integrals in an integrability argument one must check the involution condition explicitly.

Construct a canonical map for any smooth point transformation

Suppose new coordinates depend only on old coordinates, $Q=f(q)$. Preserve the canonical one-form exactly,

$$p^T dq = P^T dQ = P^T Df\, dq.$$

Thus $p=(Df)^T P$ and $P=(Df)^{-T}p$. Applying the same coordinate matrix to momenta works only for a restricted case such as an orthogonal change of coordinates. An extension of position must be accompanied by a reciprocal rescaling of its conjugate momentum. For $Q=q^2$ on $q>0$, for example, $P=p/(2q)$. The free Hamiltonian becomes $2QP^2/m$; even a physically free particle can acquire a coordinate-dependent expression for its kinetic energy.

Polar coordinates demonstrate why a momentum variable need not equal a measured tangential component. In the plane $x=r\cos\varphi, y=r\sin\varphi$. Substituting into $p_x dx+p_y dy$ gives

$$p_r=\mathbf{p}\cdot\hat{r},$$

$$p_\varphi = r\,\mathbf{p}\cdot\hat{\varphi}=L_z.$$

The physical tangential momentum is p_φ/r. At radius $0.40\,\mathrm{m}$ and transverse momentum $0.60\,\mathrm{kg\,m/s}$, the canonical angular momentum is $0.24\,\mathrm{kg\,m^2/s}$. Treating the transverse component itself as conjugate to angle would make $\{\varphi,p_\varphi^{\mathrm{wrong}}\}=1/r$. The coordinate singularity at $r=0$ also prevents these canonical polar variables from covering the entire original phase space with one chart.

Generate shears and test a numerical map

A type-two generating function $F_2(q,P,t)$ defines $p=\partial_q F_2$, $Q=\partial_P F_2$, and $K=H+\partial_t F_2$ where the transformation can be locally inverted. The differential $d(F_2-QP)$ equals $p\,dq-P\,dQ+(K-H)\,dt$. Thus the actions differ only by an endpoint term; taking another exterior derivative at fixed time gives the symplectic form condition. For $F_2=qP+sP^2/(2m)$, the map is $p=P, Q=q+sp/m$, free drift over time s. For $F_2=qP+sV(q)$, the map is $P=p-sV'(q), Q=q$, a potential kick. Their composition is canonical for any step size even though an alternating drift and kick need not give the exact trajectory.

Take a unit-mass, unit-frequency oscillator with $h=0.5$. A kick followed by drift gives $P=p-hq$, $Q=q+hP=(1-h^2)q+hp$. Its amplification matrix has determinant one and trace $2-h^2=1.75$, so its eigenvalues form a unit-modulus conjugate pair. The numerical phase per step is $\arccos(1-h^2/2)=0.50536$, slightly ahead of the exact phase 0.50000. Starting at $q=1,p=0$, the first

step gives $Q = 0.75, P = -0.50$, and physical energy becomes 0.40625 instead of 0.5. Exact bracket preservation does not imply exact conservation of the original energy at every point. If $h > 2$ the same matrix remains symplectic but has a real expanding and contracting eigenvalue; geometric preservation is compatible with numerical instability.

For a numerical implementation with analytic equations unavailable, approximate the Jacobian of a one-step map at selected states by perturbing one canonical component at a time. Compute $S^T J S - J$ after nondimensionalising coordinates and momenta in conjugate pairs. Repeat with smaller perturbations and at different integration tolerances. A determinant close to one in four dimensions is insufficient. Testing the transported brackets of one position–momentum pair and one cross-pair can reveal a swapped column or misplaced mass factor before a long simulation exhibits visible energy drift.

Track canonical area for a transported experimental ensemble

A narrow group of trajectories lets an experimenter estimate a tangent map without applying a separate perturbation to each coordinate at the same time. Prepare a small covariance ellipse in one canonical pair $z = (q, p)$, propagate each member through a known linear Hamiltonian device, and compute $\Sigma_{\mathrm{out}} = S \Sigma_{\mathrm{in}} S^T$. The ellipse area is proportional to $\sqrt{\det \Sigma}$, and a symplectic map in one pair preserves that area. Its projections can grow or contract individually: a focusing device can reduce position spread while increasing momentum spread or correlation. Measuring only the two marginal widths ignores their covariance and can falsely suggest area growth.

For a free unit-mass drift of duration two in scaled units,

$$S = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}.$$

Start with $\Sigma = \mathrm{diag}(1, 1)$. The outgoing covariance has entries $\Sigma_{qq} = 5, \Sigma_{pp} = 1, \Sigma_{qp} = 2$. The apparent rectangle of marginal standard deviations grows by $\sqrt{5}$, yet the determinant remains $5 - 4 = 1$. A plot that suppresses the correlation has discarded exactly the information needed for the area test. Noise in the two channels also contributes to observed covariance, so calibrate sensor variances independently before attributing a small apparent growth to physical dissipation or numerical smoothing.

For two pairs it is possible for a perfectly canonical coupling to exchange projected emittances. Take a uniform rotation of the two positions and the same rotation of their conjugate momenta. It is canonical by construction. If the two incoming pairs have different covariance areas, a ninety-degree rotation swaps them completely; at forty-five degrees each outgoing projected pair contains a mixture. The full four-dimensional symplectic structure remains intact. An experimentally available two-dimensional marginal then tests only part of the dynamics. Counting phase-space points in a volume and measuring cross-pair brackets answer different questions about a strongly coupled system.

The eigenvalues of $J\Sigma$ are invariant under the linear canonical covariance transport. Indeed $JS = S^{-T} J$ for symplectic S, so $J\Sigma_{\mathrm{out}} = S^{-T} (J \Sigma_{\mathrm{in}}) S^T$, a similarity transformation. For a positive covariance these eigenvalues occur as conjugate imaginary pairs, and their positive magnitudes give symplectic covariance scales. This algebraic test can distinguish loss of correlations in a fitted transfer matrix from actual exchange between pairs. It assumes linear propagation of the ensemble. A nonlinear canonical map can stretch a finite ellipse into a curved shape whose covariance determinant changes while its exact enclosed area remains preserved.

Use generating functions to find a moving normal-mode basis

Consider a two-dimensional isotropic oscillator and rotate its coordinate axes by a prescribed time-dependent angle $\alpha(t)$. Let $Q = R(\alpha)^T q$, $P = R(\alpha)^T p$ for a proper planar rotation. At any fixed time the map is orthogonal and symplectic. To derive the Hamiltonian in these moving variables, use $F_2(q, P, t) = q^T R(\alpha) P$. Its time derivative at fixed old q and new P is $-\dot{\alpha}(Q_x P_y - Q_y P_x)$, with the sign corresponding to axes rotating counterclockwise in the laboratory. The transformed Hamiltonian is therefore

$$K = H_0(Q, P) - \dot{\alpha} L_z.$$

For constant rotation rate and an isotropic potential the new Hamiltonian has no explicit time dependence. Its value is the combination energy minus rotation rate times angular momentum encountered in the Noether argument. Hamilton's equations now include $-\dot{\alpha}\,\hat{z} \times Q$ in $\dot{Q}$ relative to the rotated physical velocity. Thus the transformed canonical momentum P equals the rotated laboratory momentum rather than $m\dot{Q}$. Confusing the two yields incorrect initial conditions for a rotating-frame simulation even if the transformation matrix at the starting instant is the identity.

Take $m = 1.0\,\mathrm{kg}$, a laboratory point initially at $q = (0.20, 0)\,\mathrm{m}$ with $p = (0, 0.30)\,\mathrm{kg\,m/s}$, and axes initially aligned but rotating at $2.0\,\mathrm{rad/s}$. The canonical P_y is $0.30\,\mathrm{kg\,m/s}$. The derivative of the rotating-frame coordinate is

$$0.30 - 2.0(0.20) = -0.10\,\mathrm{m/s}.$$

It moves in the negative rotating-frame y direction while its laboratory velocity points positive y. The angular momentum is $0.060\,\mathrm{kg\,m^2/s}$, so K lies $0.120\,\mathrm{J}$ below the laboratory Hamiltonian. Both position–momentum transformations preserve brackets; the extra term derives from the moving basis, not a failure of canonical geometry.

An oscillator with two nearly equal natural frequencies provides another reason to track the generating function. If one diagonalises a slowly changing stiffness matrix at each time and keeps only the instantaneous diagonal frequencies, the omitted rotation of the eigenvectors couples the nominal modes through a term of this type. When the eigendirections turn rapidly near a degeneracy, the coupling can be appreciable even if both eigenvalues change little. Measuring the orientation as well as the eigenfrequencies allows the size of this effect to be estimated before fitting amplitude transfer to a spurious dissipation term.

Compare exact phase transport with a dissipative fit

A damped oscillator with $m\ddot{q} + c\dot{q} + kq = 0$ has equations $\dot{q} = p/m$, $\dot{p} = -kq - cp/m$. In the physical (q, p) plane their divergence is $-c/m$. A small material phase-space area therefore contracts as $A(t) = A(0)e^{-ct/m}$. No time-independent canonical coordinate change on the same two-dimensional isolated phase space can turn that physical contraction into a Hamiltonian flow preserving area. This conclusion concerns the reduced dissipative description. An expanded system including the degrees of freedom that absorb energy can still be Hamiltonian.

For $c/m = 0.40\,\mathrm{s^{-1}}$, a pair of independent initial perturbations spanning area one in calibrated coordinates spans only 0.449 after two seconds. A numerical step of $0.10\,\mathrm{s}$ fitted to that exact map should have determinant $e^{-0.04} = 0.9608$. Forcing its determinant to one during system identification would remove a real dissipative measurement. Conversely a conservative oscillator propagated by forward Euler at matched units

and $h = 0.10$ has amplification determinant $1 + h^2 = 1.01$, corresponding to a spurious area growth. A geometric diagnostic is interpretable only after deciding what physical system and coordinate pair were actually observed.

The pair q, v differs from canonical q, p by a constant mass scaling for an ordinary single particle. In one pair with constant mass that scaling does not change a measured same-coordinate before-and-after determinant. For a position-dependent mass, constrained coordinates with a varying metric, or explicit time-dependent changes of measurement scale, volume in raw position–velocity coordinates need not equal canonical volume. For example a free particle expressed in $Q = q^2$ has canonical momentum $P = m\dot{Q}/(4Q)$. A detector estimating phase area in $(Q, \dot{Q})$ without the Q-dependent factor can infer contraction or expansion even when canonical phase area is exactly conserved.

Compute a bracket of mechanical momenta and recover the Lorentz force

For a charge e in a prescribed vector potential write mechanical momentum $\boldsymbol{\pi} = \mathbf{p} - e\mathbf{A}(\mathbf{q}, t)$. The canonical brackets are $\{q_i, p_j\} = \delta_{ij}$, but the measured mechanical components satisfy

$$\{\pi_i, \pi_j\} = e(\partial_i A_j - \partial_j A_i) = e\epsilon_{ijk} B_k.$$

At nonzero magnetic field they are not simultaneously canonical momenta conjugate to the Cartesian positions. This result is gauge invariant because it involves only the curl of the vector potential. It offers a physical example where replacing canonical by measured momenta in a software interface changes the matrix of fundamental brackets.

With $H = \pi^2/(2m) + e\phi$, the time derivative is

$$\dot{\pi}_i = \{\pi_i, H\} + \partial_t \pi_i.$$

The kinetic-energy bracket contributes $(e/m)(\boldsymbol{\pi} \times \mathbf{B})_i$; the scalar potential contributes $-e\partial_i \phi$; and the explicit derivative supplies $-e\partial_t A_i$. Their sum is $e[\mathbf{E} + (\boldsymbol{\pi}/m) \times \mathbf{B}]_i$. Dropping the last term would make time-dependent electric acceleration depend on gauge. In a static uniform B_z, the transverse mechanical momenta rotate although their energy is constant. Canonical bracket structure therefore contains observable information about the force even before integrating a trajectory.

A useful finite-difference test sets $A_y = Bx$ and samples the map from canonical to mechanical variables around one state. The measured bracket $\{\pi_x, \pi_y\} = eB$ should be constant in space. A code that insists on zero will attribute the entire magnetic response to an erroneous Hamiltonian gradient. A code using mechanical variables legitimately can instead use their noncanonical bracket matrix; its equations must then preserve the correct transformed structure, not the original constant J. Testing the resulting acceleration at zero velocity and at a transverse velocity distinguishes electric and magnetic terms and catches sign errors in the antisymmetric entries.

Derive a reversible drift–kick–drift scheme and its measured phase

For $H = p^2/(2m) + V(q)$, exact kinetic flow for a duration t is $D_t : (q, p) \mapsto (q + tp/m, p)$. Exact potential flow is $K_t : (q, p) \mapsto (q, p - tV'(q))$. Each is canonical. Compose $D_{h/2} K_h D_{h/2}$ in chronological order. Applying the same construction with $-h$ undoes each substep in reverse order, so the full map is time reversible. Symmetry eliminates the even local-error powers in the effective generator, giving a second-order method for smooth V. The individual Hamiltonian pieces need not commute; if they did, the split map would be exact at any step.

For a unit oscillator, carrying out the algebra gives

$$Q = (1 - h^2/2)q + h(1 - h^2/4)p,$$

$$P = -hq + (1 - h^2/2)p.$$

Verify directly that the determinant is one. The trace gives the same step angle $2\arcsin(h/2)$ as the first-order kick–drift matrix in its stable range, but the new off-diagonal coefficients make the amplitude errors and time reversal different. Starting from $q = 1, p = 0$ at $h = 0.50$, one gets $Q = 0.875, P = -0.50$, energy 0.5078125. A negative step returns to the initial values apart from roundoff. Passing this one reversal test says nothing about absolute phase accumulated over hundreds of periods.

The map preserves $q^2 + (1 - h^2/4)p^2$ exactly for the harmonic test. At $h = 0.5$ the coefficient of p^2 is 0.9375, so the preserved ellipses differ slightly from the circles of physical energy. Their reversible oscillation in physical energy is bounded for $|h| < 2$. The per-step phase is about 0.50536 at this step, giving a phase gain 0.536 radians after one hundred steps. As h is halved while physical integration time stays fixed, the overall phase error falls approximately by four. Measuring an energy oscillation, a reversal residual and a return time therefore probes three different properties of one apparently simple map.

Compare a symplectic map with the phase of a physical orbit

Change the step of the midpoint update. Distinguish preservation of phase-space structure from agreement of the measured oscillation frequency.

Open this experiment online

Try the calculation

Time step h (ω = 1) 0.5

Number of forward steps 20

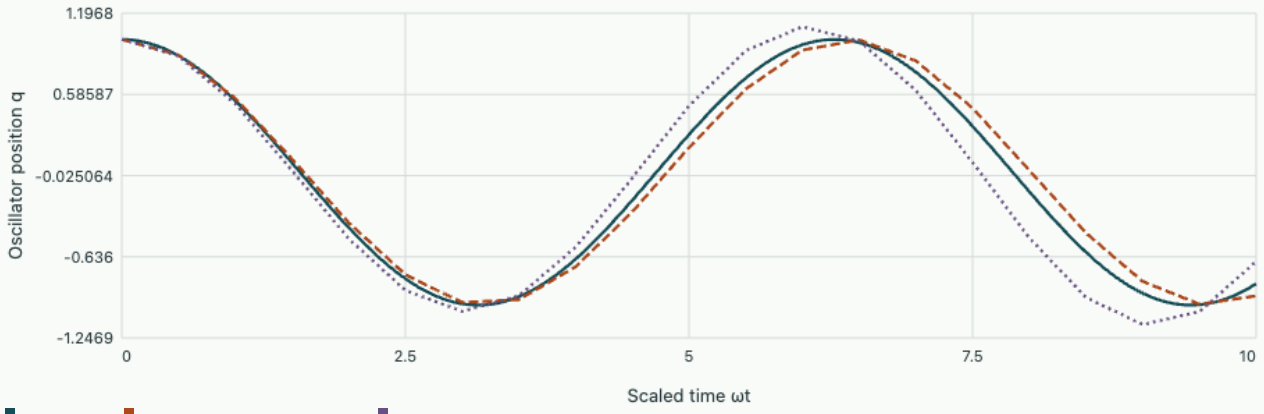

Implicit midpoint phase lag 0.2009 rad. Forward-backward distance from the start is 8.951e-16 for implicit midpoint and 0.3635 for explicit midpoint.

Relate the order of two pulses to their Poisson bracket

A canonical transformation close to identity can be built by applying two brief Hamiltonian operations in different orders. Take one free drift and one potential kick, each of short duration s. A drift then a kick gives

$$q' = q + sp/m,\ p' = p - sV'(q + sp/m).$$

A kick then a drift gives

$$p'' = p - sV'(q),\ q'' = q + s[p - sV'(q)]/m.$$

To second order their coordinate difference is $q' - q'' = s^2V'(q)/m$; their momentum difference is $p' - p'' = -s^2pV''(q)/m + O(s^3)$. The difference vector equals the Hamiltonian vector generated by $G = pV'(q)/m = \{V, T\}$, multiplied by s^2. The bracket sets the leading effect of reversing the operation order.

For the harmonic oscillator at $q = 1, p = 0, m = k = 1$ with $s = 0.10$, a drift then kick gives $(1, -0.10)$; kick then drift gives $(0.99, -0.10)$. The coordinate difference 0.01 equals s^2V'/m, and the leading momentum difference is zero for this initial condition. At $q = 0, p = 1$, the complementary experiment gives a momentum difference of -0.01 and no coordinate difference. Running both initial states tests the two components of the bracket-generated perturbation. A single apparently correct coordinate update would not detect a sign swap in the other component.

An experimental implementation might use a short force pulse as a kick and free propagation between pulses. The pulse must be short enough that motion during the kick is negligible; otherwise its own kinetic drift enters the comparison at the same order as the desired bracket. Extrapolate the measured difference divided

by s^2 to small s, and separately measure acquisition noise, which grows in relative importance under that division. Alternating the pulse order cancels many static calibration shifts but does not magically remove timing offsets correlated with the order. Treat the small bracket signal as a physical difference between two specified protocols.

A transformation preserves all coordinate–momentum brackets but a time integration on the transformed variables produces the wrong laboratory trajectory. Can the bracket check alone identify the error?

▼ **Work through the solution**

Canonical coordinates guarantee the form of Hamilton's equations. Their actual trajectories still require the correct transformed Hamiltonian, including explicit time dependence of the generating function, consistent initial momenta and a sufficiently accurate integrator. Transform an independently known initial state, compare the initial velocity and acceleration of a measurable coordinate, and follow both descriptions back to the same laboratory frame. A bracket test can establish that a change of variables is admissible without establishing that the correct energy function or physical input was attached to it.

Set compatible units before inspecting a symplectic residual

Suppose position is recorded in millimetres and momentum in kilogram metres per second. Scaling each numerical coordinate to unit variance is convenient for fitting a matrix, but independent scaling of each coordinate and momentum can change the bracket matrix from its textbook form. If $z' = Dz$ with diagonal D, then $J' = DJD^T$. For one pair this multiplies its bracket by the product of the position and momentum scale factors. In two pairs these factors can differ. A residual computed against the original J can flag a harmless unit change or, conversely, mask an error if scales happen to compensate.

Choose a characteristic length q_0, momentum p_0, and common action scale $q_0 p_0$ for each pair if the dimensionless chart is to have a common symplectic scale. Otherwise carry the transformed bracket explicitly through the test. With two oscillators of different physical units, a canonical nondimensionalisation may balance the dynamics better than using the same numerical length scale for both. Report the dimensional transformation along with any determinant or defect norm so another calculation can reconstruct which tangent vectors and measurement units were actually compared.

Tutorial on canonical tests

In two degrees of freedom try $Q_1 = q_1,\ Q_2 = q_2,\ P_1 = p_1 + \alpha q_2,\ P_2 = p_2$. Is this map canonical? What is the simplest modification that restores the condition?

▼ **Work through the solution**

The fundamental mixed momentum bracket is $\{P_1, P_2\} = \alpha$. Restore zero by also taking $P_2 = p_2 + \alpha q_1$, since then the extra term $\{p_1, \alpha q_1\} = -\alpha$ cancels. The added momentum vector is now the gradient of the scalar $\alpha q_1 q_2$, whose derivative matrix is symmetric. Both candidate maps preserve volume, reinforcing why volume alone cannot decide.

For the unit oscillator kick–drift map find a quadratic quantity it preserves exactly while its physical energy changes.

▼ **Work through the solution**

Direct substitution gives $Q^2 + P^2 - hQP = q^2 + p^2 - hqp$ for the kick–drift convention above. Multiplying by one half gives a positive modified quadratic invariant when $|h| < 2$. The first step at $h = 0.5$ gives $0.75^2 + (-0.50)^2 - 0.5(0.75)(-0.50) = 1.0$, equal to its initial value. At $h = 2$ the quadratic becomes degenerate and beyond it no longer bounds the motion.

A time-dependent momentum shift is generated by $F_2 = qP - a(t)q$. For a free physical particle find the transformed Hamiltonian and verify constant physical momentum.

▼ **Work through the solution**

The transformation gives

$$p = P - a,\; Q = q,$$
$$K = (P - a)^2/(2m) - \dot{a}\,Q.$$

Hamilton's equation gives $\dot{P} = +\dot{a}$, hence $\dot{p} = 0$. Omitting the explicit $\partial_t F_2$ term would predict constant P and a spurious acceleration of the original free particle. This test uses an observable whose physical behaviour is already fixed independently.

Compare generated canonical changes and numerical determinant tests.

David Tong develops canonical transformations and Poisson brackets in Hamiltonian dynamics.

Action, angle and the slowing clock at a separatrix

Analytical mechanics from constraints to invariants

A phase portrait contains enough information to extract a period by differentiating an area. Use that fact to build a clock for finite-amplitude pendulum motion, determine precisely where slow-parameter invariance fails, and distinguish a librating trajectory from one that makes complete rotations.

Turn a closed phase-space loop into an action variable

For a one-degree-of-freedom Hamiltonian $H(q,p;\lambda)$ at fixed parameter λ, take a connected periodic orbit of energy E. Define its action by $J(E,\lambda)=(2\pi)^{-1}\oint p\,dq$, oriented along the actual motion. The units are energy times time. A smooth family of nested librating loops has $2\pi J$ equal to the enclosed canonical phase area. For circulation on a coordinate cylinder, such as an angle increasing through 2π, the integral still defines an action even though the loop winds around the cylinder. Its area interpretation and sign need a reference branch.

At fixed q,λ, differentiating $H(q,p(E,q))=E$ gives

$$\partial_E p = 1/H_p = 1/\dot{q}.$$

Differentiating the action integral yields $\partial_E J = T/(2\pi)$, because $dq/\dot{q}=dt$ around the orbit. At ordinary turning points $\dot{q}=0$, the derivative integrand has an integrable singularity while the two endpoints of the closed loop agree; one can justify the derivative using a local square-root coordinate. Inverting gives

$$\omega(J)=\partial_J E.$$

If θ is the angle conjugate to J, the transformed Hamiltonian is $E(J)$ and $\dot{\theta}=\omega(J)$ is constant. This canonical angle generally differs from the geometric angle reported by a camera.

For a unit-mass harmonic oscillator the energy curve $p^2+\omega_0^2q^2=2E$ is an ellipse. Its area is $2\pi E/\omega_0$, giving $J=E/\omega_0$ and a frequency independent of amplitude. One may take $q=\sqrt{2J/\omega_0}\sin\theta,\ p=\sqrt{2J\omega_0}\cos\theta$, whose bracket is one. The position turns nonuniformly in coordinate space even though the phase angle advances uniformly. Under a slow change in ω_0, constant J predicts $E\propto\omega_0$ and amplitude proportional to $1/\sqrt{\omega_0}$. Constant action does not mean constant mechanical energy while work is performed by a changing spring.

Integrate the finite-amplitude pendulum

Let $I=ml^2$, $\omega_0^2=g/l$, and measure q from downward vertical. The Hamiltonian $H=p^2/(2I)+I\omega_0^2(1-\cos q)$ has separatrix energy $E_s=2I\omega_0^2$. For a libration of turning angle $q_m<\pi$, define $k=\sin(q_m/2)=\sqrt{E/E_s}$. The positive momentum branch is $p=2I\omega_0\sqrt{k^2-\sin^2(q/2)}$. Set $\sin(q/2)=k\sin u$. Evaluating four quarter-arcs of the loop gives

$$J_{\mathrm{lib}}=\frac{8I\omega_0}{\pi}\left[\mathsf{E}(k)-(1-k^2)\mathsf{K}(k)\right],$$

$$T_{\mathrm{lib}}=\frac{4}{\omega_0}\mathsf{K}(k).$$

Here $\mathsf{K}(k)=\int_0^{\pi/2}(1-k^2\sin^2 u)^{-1/2}du$ and $\mathsf{E}(k)=\int_0^{\pi/2}\sqrt{1-k^2\sin^2 u}\,du$. To verify the first formula after substitution, note that $k^2\cos^2 u=(1-k^2\sin^2 u)-(1-k^2)$. The transformed integral of $k^2\cos^2 u/\sqrt{1-k^2\sin^2 u}$ splits exactly into the two quoted complete elliptic integrals. The period follows either by $dq/\dot q$ or by differentiating the action. That independent check is useful when numerical differentiation of noisy action estimates amplifies sampling errors.

At small k,

$$\mathsf{K}=(\pi/2)(1+k^2/4+\cdots),$$

$$T\approx(2\pi/\omega_0)(1+q_m^2/16+\cdots).$$

At $q_m=60^\circ$, $k=1/2$ and $\mathsf{K}\simeq 1.68575$, giving $T/T_{\mathrm{small}}=1.0732$. A pendulum of length $1.00\,\mathrm{m}$ then has period about $2.153\,\mathrm{s}$, where the small-angle answer is $2.006\,\mathrm{s}$. Timing ten periods rather than one helps resolve the accumulated difference. Measuring the arc amplitude from a projected image requires a geometric correction if the swing plane is tilted relative to the camera.

Approach the saddle from either side

At the upright position $q=\pi$, the potential is maximal and the equilibrium is a hyperbolic saddle in phase space. Put $q=\pi+s$ with small s; the energy near the saddle is $H\simeq E_s+p^2/(2I)-I\omega_0^2 s^2/2$. For an orbit at $E_s-\Delta E$, the time spent in a fixed small neighbourhood includes an integral proportional to $\int ds/\sqrt{s^2-2\Delta E/(I\omega_0^2)}$. As ΔE decreases, this develops a logarithm. Equivalently $\mathsf{K}(k)\sim\ln[4/\sqrt{1-k^2}]$ as $k\to 1^-$. The separatrix encloses finite canonical area but its period is infinite. Hence $\partial_E J$ diverges although J itself remains finite.

On the positive rotation branch $E>E_s$, set $k_r=\sqrt{E_s/E}<1$. A period is the time for q to increase through 2π, so there is only one positive branch of momentum in its action integral. Direct integration gives

$$J_{\mathrm{rot}}=(2/\pi)\sqrt{2IE}\,\mathsf{E}(k_r),$$

$$T_{\mathrm{rot}}=4\sqrt{I/(2E)}\,\mathsf{K}(k_r).$$

At very high energy, the potential is a small modulation and these formulas approach free-rotor values $\sqrt{2IE}$ and $2\pi\sqrt{I/(2E)}$. At the separatrix the positive rotation action approaches $4I\omega_0/\pi$, half the limiting libration action $8I\omega_0/\pi$. The two actions count different closed paths on a cylinder. The factor two is a topological change in the cycle, not a discontinuous jump in the mechanical momentum at a point.

Test slow conservation against the actual orbit period

To treat slowly varying parameters, define instantaneous action–angle coordinates for each fixed λ. The generating function $W(q,J,\lambda)$ acquires explicit time dependence when $\lambda=\lambda(t)$, adding $\dot\lambda\,\partial_\lambda W$ to the transformed Hamiltonian. Its angle dependence produces an oscillating $\dot J=-\dot\lambda\,\partial_\theta\partial_\lambda W$. On a regular periodic orbit, the integral of an angle derivative over a cycle vanishes. If λ changes little during that cycle and there is no resonant complication, this cancellation explains why the action changes little on average. The

requirement uses the actual orbit period at the chosen energy. Near a separatrix it cannot be justified merely by comparing a drive rate to the small-amplitude period.

Consider a parameter ramp lasting $20\,\mathrm{s}$ for the one-metre pendulum. Near the bottom this covers about ten periods. For $1 - E/E_s = 10^{-6}$, the asymptotic $\mathsf{K} \simeq \ln 4000 = 8.29$ gives a period of approximately $10.6\,\mathrm{s}$, so the same ramp spans less than two cycles. Even longer ramps fail the scale separation sufficiently close to the saddle. If a varying parameter changes the separatrix area until it meets the trajectory, continued use of a single analytic action branch is especially unreliable. A trajectory can pass into circulation or return to libration depending on its phase and the ramp history.

Compute an action without measuring momentum continuously

Suppose a camera tracks one-dimensional positions over one full cycle, and velocity estimates are noisy. For a known position-dependent potential and calibrated mass, an independent measurement of the turning energy allows phase-area integration without differentiating the camera trace. On a bound orbit $p_{\pm}(q) = \pm\sqrt{2m[E - V(q)]}$, so $J = \pi^{-1}\int_{q_-}^{q_+}\sqrt{2m[E - V(q)]}\,dq$. One needs a potential calibration, both turning positions and a common energy datum. An additive constant in the potential cancels when E is shifted by the same amount. An unrecognised dissipative force makes the inward and outward energy profiles different and invalidates the assumed single closed energy curve.

Take a particle in the potential $V(q) = a|q|$, $a > 0$, which can approximate opposed constant restoring forces away from a narrow central smoothing zone. Its turning displacements are $\pm E/a$. Integrate on the positive half and use symmetry to find

$$J = 4\sqrt{2m}E^{3/2}/(3\pi a).$$

Differentiating gives

$$T = 2\pi\,dJ/dE = 4\sqrt{2mE}/a.$$

The same period follows by constant acceleration: starting at rest at a positive turning point, the fall to the origin takes $\sqrt{2mE}/a$, one quarter period. Agreement across a nondifferentiable potential cusp is a useful check that the action calculation handles the turning boundaries rather than requiring sinusoidal motion.

For $m = 0.10\,\mathrm{kg}$, $a = 2.0\,\mathrm{N}$, and $E = 0.20\,\mathrm{J}$, the amplitude is $0.10\,\mathrm{m}$, action is $0.00849\,\mathrm{J\,s}$, and period is $0.400\,\mathrm{s}$. If the restoring strength a doubles slowly at fixed action, $E \propto a^{2/3}$ increases by $2^{2/3} = 1.59$, while the turning amplitude shrinks by $2^{-1/3} = 0.794$. A harmonic-oscillator action rule $J = E/\omega$ using a single measured frequency would be incorrect here: the actual $E/\omega = ET/(2\pi)$ exceeds the action by a factor $3/2$. The equality found for a harmonic oscillator follows from its linear energy–action relation and need not generalise.

Separate the angle accumulated during a cycle from the action it encloses

In a slowly modulated oscillator an almost constant action can coexist with a measurable phase error. Consider $\ddot{q} + \omega(t)^2 q = 0$ with smooth positive ω. An approximate oscillation is $q = A(t)\cos\Phi(t)$. At leading

order $\dot{\Phi} = \omega$. Cancelling the first-order sine term gives $2\dot{A}\omega + A\dot{\omega} = 0$, so $A\sqrt{\omega}$ is approximately constant. The amplitude law reproduces conservation of E/ω after averaging. The accumulated phase is $\int \omega(t)\,dt$, which can change by many radians during a slow parameter cycle even if the initial and final frequency and action agree.

For $\omega(t) = 10[1 + 0.10\sin(2\pi t/T_r)]\,\mathrm{s}^{-1}$ and $T_r = 20\,\mathrm{s}$, integration over one full ramp gives $200\,\mathrm{rad}$ of leading phase. The largest value of $|\dot{\omega}|/\omega^2$ is only a few 10^{-3}. To first order the final amplitude returns to its initial value, but the position at the end is set by a phase that is generally not an integer multiple of 2π. A calculation that expects to return to its initial point after the parameter cycle confuses action invariance with angle closure. A small constant frequency calibration error, integrated over a long ramp, can dominate the final phase error while producing little amplitude error.

A sudden frequency switch gives an instructive opposite limit with exact phase dependence. At fixed q, p, changing ω_1 instantly to ω_2 changes the action to $J_2 = p^2/(2m\omega_2) + m\omega_2 q^2/2$. Write the initial state at action J_1 using the harmonic angle. Then

$$J_2/J_1 = (\omega_1/\omega_2)\cos^2\theta + (\omega_2/\omega_1)\sin^2\theta.$$

For a frequency doubling, the ratio ranges from one half at a passage through the centre to two at a turning point. For a uniformly sampled ensemble the average ratio is 1.25, although the pre-switch and post-switch Hamiltonians are both perfectly integrable.

An experiment can therefore test slowness by preparing the same initial action at different phases and comparing the distribution of final actions. A broad phase-dependent distribution for a rapid switch narrows as the ramp becomes slow on a regular family of tori. Near a separatrix a slow external ramp can still yield a broad distribution because the orbital clock itself stalls. The size of the instantaneous energy, measured without phase and the final Hamiltonian, does not distinguish these mechanisms.

Resolve a separatrix in time without integrating arbitrarily close to the saddle

At exactly E_s, energy conservation gives $p = \pm 2I\omega_0\cos(q/2)$ on $-\pi < q < \pi$. Choose the increasing branch passing through $q = 0$ at $t = 0$. Separation of variables in $\dot{q} = 2\omega_0\cos(q/2)$ gives

$$\tan(\pi/4 + q/4) = e^{\omega_0 t}.$$

Therefore $q(t) = 4\arctan(e^{\omega_0 t}) - \pi$,

$$p(t) = 2I\omega_0\,\mathrm{sech}(\omega_0 t).$$

This exact orbit approaches the upright saddle exponentially at both infinite endpoints. The two endpoints represent the same physical angle on the cylinder even though they differ by 2π on an unwrapped plot.

The separatrix momentum falls to one tenth of its maximum after $\omega_0 t = \operatorname{arcosh} 10 \approx 2.993$. At $\omega_0 = 3.13\,\mathrm{s}^{-1}$ that takes only $0.956\,\mathrm{s}$. Reaching one hundredth takes $\operatorname{arcosh} 100/\omega_0 \approx 1.69\,\mathrm{s}$. The logarithmic increase in time per retained significant digit explains why an apparently captured upright state can simply be a trajectory very close to the saddle. Measuring angular velocity as well as angular displacement makes the exponential approach easier to distinguish from a truly stationary upright solution.

The phase area enclosed by both separatrix branches can also be recovered directly. Integrating their difference gives

$$\int_{-\pi}^{\pi} 4I\omega_0 \cos(q/2) dq = 16 I\omega_0.$$

Dividing by 2π yields the libration limit $8I\omega_0/\pi$ obtained using elliptic integrals. By contrast a trajectory winding once on the upper branch integrates only the upper momentum, giving half that value. Testing these two integrals on an unwrapped numerical phase trace can reveal a missing lower turning branch or a routine that silently wraps the angle before accumulating $p\,dq$.

Find a frequency directly from noisy orbital areas

Suppose a calibrated experiment measures the phase areas $A(E) = \oint p\,dq$ at several energies. With $J = A/(2\pi)$, the identity $dA/dE = T$ says that an area-versus-energy slope directly has units of time. Take three close energies $0.09, 0.10, 0.11\,\mathrm{J}$, with measured areas $0.180, 0.202, 0.225\,\mathrm{J\,s}$. A central difference gives $T \simeq (0.225 - 0.180)/0.020 = 2.25\,\mathrm{s}$ at the centre. If each area has an independent uncertainty $0.001\,\mathrm{J\,s}$, the propagated standard uncertainty of that difference is about $0.071\,\mathrm{s}$, before accounting for energy uncertainty and curvature of $A(E)$.

Decreasing the energy spacing would reduce smooth finite-difference bias but amplify the effect of those same area errors. A local fit over several energies can trade variance against the assumption of a smooth slope. Near a separatrix the actual slope varies rapidly, so a wide polynomial fit may smooth away precisely the period increase being sought. Compare with direct return times at selected energies, or fit the logarithmic form justified by the saddle expansion when the orbit is demonstrably in its asymptotic regime. A measured phase area can be highly accurate while an inferred frequency remains poorly constrained because it is a derivative.

A practical area estimate from ordered camera samples uses the polygon formula $\frac{1}{2}\sum_i (p_i + p_{i+1})(q_{i+1} - q_i)$, including the return to the initial point with a consistent sign. Very long time sampling can overrepresent the slow turn neighbourhood but need not corrupt a properly oriented geometric line integral. It can make a uniform-clock finite difference of velocity near a turnaround noisy, however. If a dataset includes several cycles with dissipation, do not simply close the final state back to the beginning of the first cycle. Find individual return events, estimate the action per cycle and test whether its change is consistent with measured losses.

Construct the canonical angle on two sides of a turning point

For a bound one-dimensional potential take a reference crossing at the lower turning point and choose the increasing-q branch. A Hamilton–Jacobi generating integral is

$$W(q, J) = \int_{q_-}^{q} p(q', E(J))\, dq'.$$

Its angle is $\theta = \partial_J W = (dE/dJ)\int_{q_-}^{q} dq'/\dot{q}' = \omega t_{\mathrm{flight}}$ on that branch. At the upper turning point $\theta = \pi$ for a time-reversal-invariant kinetic-plus-potential Hamiltonian. On the return branch the same coordinate q must

be assigned the complementary angle $2\pi - \theta_{\mathrm{outgoing}}(q)$, because its momentum has opposite sign. The coordinate alone is insufficient to reconstruct a unique phase.

For a harmonic oscillator beginning at its leftmost position a quarter-amplitude point is crossed at two different times. In coordinates $q = -A\cos\theta$ take $q = -A/2$. The increasing branch gives

$$\theta = \pi/3;$$

the returning branch gives $5\pi/3$. Their velocities have opposite signs. Using the principal arccosine for every time sample forces both branches onto the same interval and makes an apparently discontinuous phase that never completes a cycle. Measuring the sign of velocity or interpolating successive turning events restores the appropriate branch. Near a saddle, tiny velocities make this branch identification particularly sensitive to camera noise.

Once a reliable angle is available, the phase deviation from the uniform prediction $\theta(t_0) + \omega(t - t_0)$ becomes a useful diagnostic. A static potential with a wrongly calibrated shape can cause a phase error periodic in position even when its return time is fitted correctly. Dissipation usually changes the action as well, while a constant clock error causes a steadily accumulating phase drift. Compare the residual over several cycles at different amplitudes and plot it against reconstructed angle as well as laboratory time. These alternate plots test different physical explanations of the same timing mismatch.

Measure the area of a closed orbit

Inspect how deformation of a potential changes an energy contour. Compute its action before using it to infer an adiabatic frequency.

Open this experiment online

Try the calculation

Final to initial force ratio

2

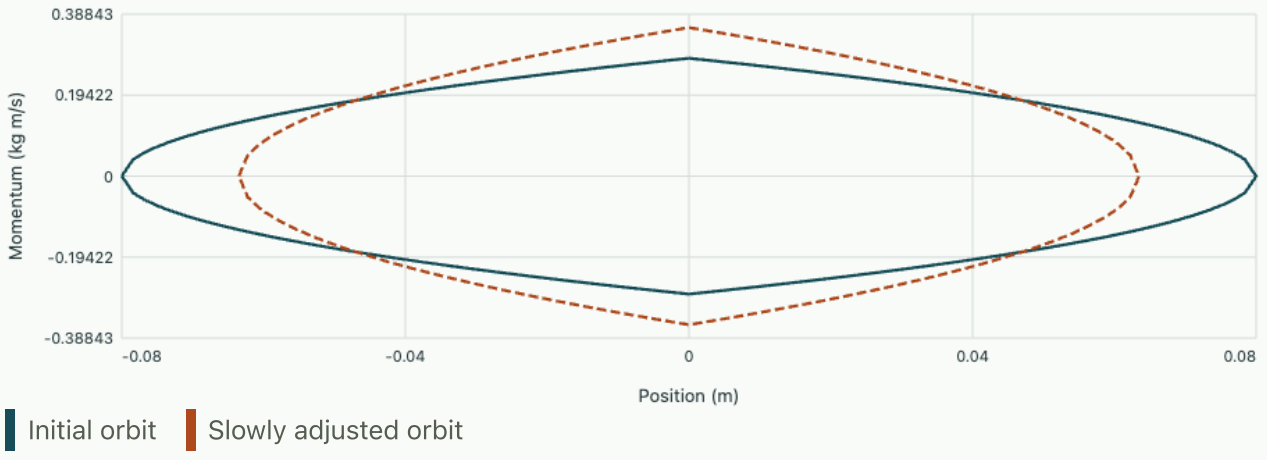

Final amplitude 0.0635 m, energy 0.254 J and period 0.3564 s.

Separate a true separatrix crossing from a change of angle convention

Imagine slowly reducing the pendulum's gravitational strength while a bob initially librates. The separatrix action $J_s = 8I\omega_0/\pi$ decreases with ω_0. Away from the crossing the orbit approximately maintains its own action, so its energy as a fraction of the instantaneous barrier rises. When the shrinking separatrix action reaches that of the orbit, no closed librating torus of that action remains. The system can enter one of the two circulation branches. Its action convention changes because the relevant closed curve on the cylinder now winds rather than traversing two momentum branches. An apparent numerical factor-of-two jump caused by this convention is distinct from the actual dynamics through the slow saddle neighbourhood.

For an initial libration action equal to 0.40 of the original separatrix action at fixed I, a parameter protocol reducing ω_0 to half its original value leaves J/J_s about 0.80, assuming adiabatic tracking remains valid. Reducing it to 0.30 of the initial value would formally give $J/J_s = 1.33$, impossible for a libration. The naive crossing estimate is $\omega_0 = 0.40\omega_{0,\mathrm{initial}}$, but the timing and selected rotation direction require integrating the actual trajectory near the saddle; the adiabatic argument loses uniform accuracy there. A local measure of p and unwrapped q supplies the physical distinction.

A real parameter protocol must specify how gravity strength changes. A pendulum length ramp changes both $I = ml^2$ and $\omega_0 = \sqrt{g/l}$, so its separatrix area scales as $I\omega_0 = m\sqrt{g}\,l^{3/2}$. Reducing length to 0.64 of its starting value reduces the separatrix action by $0.64^{3/2} = 0.512$, whereas replacing only ω_0 in a fixed-I model would predict the opposite trend. A literal change of length also brings radial kinetic energy and actuator work, as derived for the reelled pendulum. Specify the physical control path in parameter space before applying a numerical action law.

An unwrapped angle increases by 2π over a recorded interval but the phase portrait seems to contain a bounded loop. Can both observations be trusted simultaneously?

▼ **Work through the solution**

A true circulation winds around the cylindrical phase space, whereas a libration has zero net unwrapped displacement after a full return. A plot that wraps angles into a fixed interval and draws a line across its boundary can display an artificial closing segment. Conversely, measurement noise near $\pm\pi$ can make a naive unwrapping code insert a spurious 2π jump. Inspect signed momentum and the time-resolved crossing of the angle boundary, then reconstruct the loop in adjacent copies of the coordinate interval. Explicitly identifying the cycle on which the action was integrated resolves the apparent contradiction.

Infer work during an adiabatic change from a phase average

Differentiate the action integral at fixed energy with respect to the external parameter λ. The energy relation gives $\partial_\lambda p|_{E,q} = -H_\lambda/H_p$, so $J_\lambda|_E = -(2\pi)^{-1}\oint H_\lambda\,dt$. Combine this with $J_E = T/(2\pi)$ and differentiate at fixed J. The result is

$$\partial_\lambda E|_J = \langle \partial_\lambda H \rangle_{\mathrm{time}}.$$

A slowly changed parameter therefore does net reversible work at the orbit-averaged rate dictated by its conjugate generalised force. The average must weight time, not distance along the plotted phase curve.

For a harmonic spring with stiffness k, $H_k = q^2/2$ and $\langle q^2 \rangle = A^2/2 = E/k$, giving $\partial_k E|_J = E/(2k)$. Integrating yields $E \propto \sqrt{k}$, consistent with $J = E/\omega$. A 4% slow stiffness increase raises energy by about 1.98% exactly under this law. Changing stiffness at a single turning point would instead add 4% to the energy for that sudden change. The difference quantifies why observing only identical start and end stiffnesses does not identify the work protocol.

Tutorial on interpreting an action experiment

A harmonic oscillator has energy $0.030\,\mathrm{J}$ and frequency $5.0\,\mathrm{Hz}$. Its stiffness is increased slowly until the frequency doubles. Determine action, final energy and relative amplitude.

▼ **Work through the solution**

The action is

$$0.030/(2\pi \cdot 5.0) = 9.55 \times 10^{-4}\,\mathrm{J\,s}.$$

Approximately constant action gives a final energy of $0.060\,\mathrm{J}$. With fixed mass and $E = m\omega^2 A^2/2$, the final amplitude is $1/\sqrt{2}$ times its initial value. A sudden change performed at a particular phase would in general give a different final action.

A simulation estimates pendulum periods by differentiating sampled action values on a coarse uniform energy grid. Why is the last estimated period below the separatrix likely to be inaccurate even if the individual areas are well converged?

▼ **Work through the solution**

A finite difference averages $dJ/dE = T/(2\pi)$ over a finite energy interval. The derivative rises logarithmically near the saddle, so the average can underestimate its endpoint value markedly. Refine adaptively in the distance to E_s, use the elliptic quadrature or measure return times directly. The numerical integrator must also locate the long dwell near the unstable equilibrium without stepping across it artificially.

Two trajectories have energy $E_s + \delta$, one rotating clockwise and one anticlockwise. Can energy and the positive magnitude of rotation action distinguish them?

▼ **Work through the solution**

Both have the same speed at each angle and the same positive action magnitude and period. The sign of momentum or a signed orientation of the fundamental rotational cycle distinguishes the branches. Near the separatrix a numerical classifier looking only at energy and unsigned area can lose track of this information. Count the unwrapped angular change and inspect the sign of p to identify circulation direction.

Measure phase area and compare its change under a parameter ramp.

David Tong develops action–angle variables and adiabatic invariance.

Canonical perturbations, small denominators and resonant capture

Analytical mechanics from constraints to invariants

A weak periodic force can produce a bounded ripple on most trajectories and a large change on a narrow band of initial conditions. Derive the averaging transformation away from resonance, construct the slow resonant dynamics where that transformation fails, and estimate the physically observable width and time scale of a resonance.

Separate the unperturbed clock from a weak correction

Take an integrable Hamiltonian $H(I,\theta,t)=H_0(I)+\epsilon H_1(I,\theta,t)$ in action–angle variables. With $\epsilon=0$, actions are constants and angles advance at $\boldsymbol{\omega}(I)=\partial_I H_0$. A perturbation with Fourier component $h_{\mathbf{n}}(I)\cos(\mathbf{n}\cdot\theta-\Omega t)$ drives action changes proportional to its sine. If the phase sweeps quickly through all angles, positive and negative contributions largely cancel. If $\mathbf{n}\cdot\boldsymbol{\omega}-\Omega$ is small, the perturbation acts with a coherent sign over a long time. We want a transformation that makes both statements quantitative.

Use a near-identity type-two generating function $F_2(\theta,J,t)=\theta\cdot J+\epsilon S_1(\theta,J,t)$. Old actions are $I=J+\epsilon\partial_\theta S_1$, and new angles are $\psi=\theta+\epsilon\partial_J S_1$. Expand $H_0(I)=H_0(J)+\epsilon\boldsymbol{\omega}(J)\cdot\partial_\theta S_1+O(\epsilon^2)$. At first order the new Hamiltonian is $K=H_0(J)+\epsilon[H_1+\partial_t S_1+\boldsymbol{\omega}\cdot\partial_\theta S_1]$. Replacing old by new angles inside that bracket produces only higher-order errors. To leave the averaged part $\langle H_1\rangle$, solve the homological equation

$$(\partial_t+\boldsymbol{\omega}\cdot\partial_\theta)S_1=-\big(H_1-\langle H_1\rangle\big).$$

For an oscillating cosine with detuning $\Delta=\mathbf{n}\cdot\boldsymbol{\omega}-\Omega$, choose $S_1=-h_{\mathbf{n}}\sin(\mathbf{n}\cdot\theta-\Omega t)/\Delta$. The transformation is useful while the resulting action displacement $\epsilon|\partial_\theta S_1|$ is small on the relevant action scale and the derivatives entering the angle transformation remain controlled. Small ϵ alone is insufficient if Δ is even smaller. A singular denominator says that this choice to remove the harmonic is unsuitable, not that the underlying mechanical force becomes infinite.

Obtain a nonlinear frequency shift from the angle average

A free oscillator with mass m, small-amplitude frequency ω_0, and quartic correction $\beta q^4/4$ has unperturbed coordinates $q=\sqrt{2I/(m\omega_0)}\,\sin\theta$. Substituting gives the correction $\beta I^2\sin^4\theta/(m^2\omega_0^2)$. Since $\langle\sin^4\theta\rangle=3/8$, the first-order averaged Hamiltonian is $K=\omega_0 J+3\beta J^2/(8m^2\omega_0^2)$. Its derivative gives

$$\omega=\omega_0+3\beta J/(4m^2\omega_0^2).$$

Expressing $J\simeq m\omega_0 A^2/2$ in terms of leading oscillation amplitude yields the frequency shift $3\beta A^2/(8m\omega_0)$. Oscillatory harmonics at $2\theta,4\theta$ change the waveform but do not contribute directly to this averaged first-order clock.

For $m=0.20\,\mathrm{kg}$, $\omega_0=10.0\,\mathrm{s^{-1}}$, $\beta=80\,\mathrm{N/m^3}$, and $A=0.10\,\mathrm{m}$, the predicted shift is $0.150\,\mathrm{s^{-1}}$, about 1.5%. The nonlinear-to-linear force ratio at maximum displacement is $\beta A^2/(m\omega_0^2)=0.040$, consistent with the expansion. At twice the amplitude the shift quadruples and the expansion parameter grows to 0.16. Measuring period over several amplitudes and plotting the frequency against A^2 tests both the slope and the

onset of curvature from higher orders. A positive β stiffens the spring; a negative quartic term is only a local approximation unless higher powers restore global confinement.

Solve the slow motion in the resonant frame

Consider a rotor $H = I^2/(2\mathcal{I}) - \epsilon V \cos(\theta - \Omega t)$, where $\mathcal{I}$ is moment of inertia and $\epsilon V > 0$ has energy units. Choose $F_2 = (\theta - \Omega t)P$, so $P = I$,

$$\phi = \theta - \Omega t,$$

$$K = P^2/(2\mathcal{I}) - \Omega P - \epsilon V \cos\phi.$$

Completing the square gives $K = (P - P_*)^2/(2\mathcal{I}) - \epsilon V \cos\phi + \text{constant}$ with $P_* = \mathcal{I}\Omega$. This is an exact pendulum in the rotating frame. It retains the resonant cosine instead of attempting to average it away. Phase locking means libration of ϕ around zero; outside its separatrix the relative phase circulates.

The stable fixed point has $P = P_*, \phi = 0$; its small libration frequency is

$$\omega_{\text{lib}} = \sqrt{\epsilon V/\mathcal{I}}.$$

The saddle at $\phi = \pi$ has energy $+\epsilon V$ relative to the same constant. At $\phi = 0$ that energy is reached for $|P - P_*| = 2\sqrt{\mathcal{I}\epsilon V}$, the resonance half-width in action. Its complete width is twice that number. For $\mathcal{I} = 0.050\,\mathrm{kg\,m^2}$, $\Omega = 8.0\,\mathrm{s^{-1}}$ and $\epsilon V = 0.0020\,\mathrm{J}$, the central action is $0.40\,\mathrm{kg\,m^2/s}$; the half-width is $0.020\,\mathrm{kg\,m^2/s}$ and small libration period $31.4\,\mathrm{s}$. The rotor makes many laboratory turns while its slow relative phase completes one libration.

An initial action inside the quoted half-width is necessary but not sufficient for trapping at every phase. At $\phi = \pi$, any nonzero $P - P_*$ is above the saddle energy. Compute the actual rotating-frame energy of the initial state and compare with the saddle. For $\phi = 0, P = 0.415$, the kinetic detuning energy is $0.00225\,\mathrm{J}$; subtracting $0.0020\,\mathrm{J}$ gives $0.00025\,\mathrm{J}$, below the saddle. With the same action at $\phi = \pi$ the value is $0.00425\,\mathrm{J}$, above the saddle. Selecting trajectories by mean frequency alone can therefore mix trapped and circulating populations.

Retain a resonant harmonic in a general integrable system

For one action with $H_0'(I_*) = \Omega$, use the rotating frame and expand $H_0(I) - \Omega I$ about I_*. The linear term vanishes, leaving $H_0''(I_*)(I - I_*)^2/2$. Retain the slowly varying resonant cosine and average the rapidly varying harmonics. When $H_0'' \neq 0$, this gives the same pendulum normal form with an effective inertia and possibly a reversed sign of the stable phase. The action half-width scales as $2\sqrt{\epsilon|V/H_0''|}$. Thus a weak perturbation can influence a band of width proportional to $\sqrt{\epsilon}$, parametrically larger than a typical nonresonant action ripple proportional to ϵ.

For an exactly harmonic oscillator $H_0 = \omega_0 I$ the curvature vanishes. One cannot replace H_0'' with zero in the width formula. An exactly tuned additive drive in a linear oscillator instead produces secular growth until damping, detuning, nonlinearity or a finite drive interval intervenes. In multiple actions, one first transforms to the resonant phase $\mathbf{n} \cdot \theta - \Omega t$ and complementary angles, preserving the canonical one-form. Other near-resonances can prevent isolation of a single pendulum. Apparent chaos inferred from a short irregular time

trace alone needs additional tests, including resolved phase portraits, convergence under time refinement and identification of competing resonances.

Calculate a phase-locking threshold for a drive that sweeps in frequency

The exact rotating rotor also allows a tractable time-dependent drive frequency. Write the pattern angle as $\alpha(t)$, with $\Omega(t) = \dot{\alpha}$. In its frame $K = P^2/(2\mathcal{I}) - \Omega(t)P - \epsilon V \cos\phi$. Hamilton's equations give

$$\dot{\phi} = P/\mathcal{I} - \Omega(t),$$

$$\dot{P} = -\epsilon V \sin\phi.$$

Define the detuning momentum $p = P - \mathcal{I}\Omega(t)$. Even when $\dot{\phi} = 0$ at one instant, phase locking throughout a ramp requires $\dot{p} = 0$ and hence $\sin\phi_0 = -\mathcal{I}\dot{\Omega}/(\epsilon V)$. The restoring torque must be sufficient to accelerate the rotor along with the changing pattern.

For constant $\dot{\Omega}$ and $|\mathcal{I}\dot{\Omega}| < \epsilon V$, a stable locked phase with $\cos\phi_0 > 0$ exists. Perturbing the phase by η gives $\ddot{\eta} + (\epsilon V/\mathcal{I})\cos\phi_0\,\eta = 0$, so the libration softens as the required torque approaches its maximum. At threshold the stable and unstable stationary phases meet. Above it no fixed phase can provide the torque needed by the ramp. The mere existence of a stable locked solution still does not imply that an arbitrary initial phase and momentum will be captured in it. In a conservative system initial phase-space location and separatrix crossing are essential.

With the previous $\mathcal{I} = 0.050\,\mathrm{kg\,m^2}$ and $\epsilon V = 0.0020\,\mathrm{J}$, the largest trackable constant acceleration in this model is $0.040\,\mathrm{rad/s^2}$. At half that value the stable phase lags by $\pi/6$, and the small relative-phase frequency is

$$0.20\sqrt{\cos(\pi/6)} = 0.186\,\mathrm{s^{-1}}.$$

Increasing the drive frequency by $1.0\,\mathrm{rad/s}$ at this rate takes $50\,\mathrm{s}$, only about one and a half libration periods. Adiabatic capture estimates that additionally require many libration periods may therefore be unconvincing even though the elementary torque threshold is satisfied. A measured lag versus sweep rate assesses the static torque condition, while a full phase trajectory tests dynamic capture.

The same equations admit a tilted effective washboard potential in the accelerating frame, $U_{\mathrm{rel}} = -\epsilon V\cos\phi + \mathcal{I}\dot{\Omega}\,\phi$, as long as $\dot{\Omega}$ is constant. Its potential difference over a whole 2π turn is nonzero, so unwrap the relative phase before using its energy. Mapping every phase back to $[-\pi, \pi)$ hides this difference and produces artificial jumps in the inferred energy. In the original variables the physical motor does work. There is no contradiction between a conserved local accelerating-frame energy for a constant ramp and the changing laboratory kinetic energy of a locked rotor.

Derive secular energy exchange in a pair of weakly coupled oscillators

Take two equal masses m with individual natural frequency ω_0 and potential coupling $k_c(q_1 - q_2)^2/2$. The normal coordinates $Q_\pm = (q_1 \pm q_2)/\sqrt{2}$ have frequencies $\omega_+ = \omega_0$,

$$\omega_- = \sqrt{\omega_0^2 + 2k_c/m}.$$

Release mass one at displacement A with both velocities and q_2 initially zero. Superposition gives

$$q_1 = A[\cos\omega_+ t + \cos\omega_- t]/2,$$
$$q_2 = A[\cos\omega_+ t - \cos\omega_- t]/2.$$

The slowly varying difference phase transfers oscillation amplitude between the masses over a time much longer than their rapid periods.

For $\omega_0 = 10\,\mathrm{s}^{-1}$ and $k_c/m = 1.0\,\mathrm{s}^{-2}$, the exact separation is

$$\sqrt{102} - 10 = 0.0995\,\mathrm{s}^{-1}.$$

The first near-complete exchange occurs at $t = \pi/|\omega_- - \omega_+| \approx 31.6\,\mathrm{s}$. Over a single fast period the displacements can look almost like uncoupled oscillations, despite large later energy redistribution. Perturbing two isolated oscillators while trying to average the coupling over both independent fast angles would discard a term depending on their slow difference phase. Retaining the resonant combination predicts the exchange; diagonalising the linear model supplies a reference solution for that approximation.

Detune the uncoupled springs so their natural frequencies have a separation large compared with $k_c/(m\omega_0)$. The coupled normal modes then align more closely with the original oscillators, and the energy exchange cannot reach its former depth. In a nonlinear problem the individual frequencies also depend on action, so the energy transfer can detune itself. To interpret an observed beating envelope, measure both normal frequencies and relative phases. An apparent loss of amplitude in oscillator one can be energy stored in oscillator two, distinguishable from damping through recovery at later times and a balance of the total Hamiltonian.

Estimate the accuracy of averaging over a finite observation time

For a nonresonant perturbation with initially slow action change, one can integrate the first-order drive along the unperturbed trajectory over the actual observation window instead of assuming an infinite-time average. For $H_1 = A\cos(\theta - \Omega t)$, $\dot{I} = A\sin(\Delta t + \phi_0)$ to leading order with $\Delta = \omega_0 - \Omega$. Thus $\Delta I(T) = A[\cos\phi_0 - \cos(\Delta T + \phi_0)]/\Delta$. For $\Delta \neq 0$ its absolute value remains at most $2|A/\Delta|$, and its largest possible value over initial phase is $2|A\sin(\Delta T/2)/\Delta|$. For $\Delta \to 0$ the phase-dependent result tends to $AT\sin\phi_0$. The finite-time expression is regular right at resonance and shows exactly when a secular term emerges.

For $A = 0.001\,\mathrm{J}$, $T = 10\,\mathrm{s}$ and $\Delta = 0.20\,\mathrm{s}^{-1}$, the maximum-amplitude change in action over initial phase is $0.00841\,\mathrm{J\,s}$. At zero detuning the leading largest change is $0.0100\,\mathrm{J\,s}$, only modestly larger because this observation window resolves a frequency band of order $1/T$. For $\Delta = 2\pi/T$ the change after exactly T

vanishes, although the action swung during the interval. A final-minus-initial measurement alone can mistake a completed nonresonant excursion for an entirely unperturbed trajectory. Recording at an additional end time tests that cancellation.

At leading order a force of finite duration produces this sine-over-detuning envelope even in a linear system. It is an observation-time bandwidth, distinct from the dynamical resonance width proportional to $\sqrt{\epsilon}$ in the nonlinear rotor. Repeat the experiment with doubled drive duration at fixed amplitude. Features that narrow inversely with the duration are evidence of finite-time spectral selection. A persistent bounded island in a stroboscopic phase portrait, with the predicted amplitude-dependent width and libration frequency, is evidence of the phase-locking dynamics. If drive amplitude is large enough to change the natural frequency appreciably over the interval, the unperturbed phase-integration formula ceases to describe the trajectory.

Construct a stroboscopic section with the right clock

For a time-periodic Hamiltonian sample (I, θ) every period of the external drive at a fixed drive phase. This creates a two-dimensional area-preserving map for one periodically forced degree of freedom. A phase-locked solution can become a fixed point of this map, and librating trajectories around it can form an island. Sampling at each crossing of a moving coordinate threshold instead can produce a different return map and a different phase weight. A return plot can look irregular simply because the drive phase was not held fixed at the samples.

Consider a rigid rotor commanded at $\Omega = 8\,\mathrm{s}^{-1}$. Its drive period is

$$2\pi/8 = 0.7854\,\mathrm{s}.$$

An acquisition system sampling nominally every $0.7900\,\mathrm{s}$ accumulates approximately $8(0.7900 - 0.7854) = 0.0368\,\mathrm{rad}$ of phase error per sample. After one hundred samples the sampling phase has moved about $3.68\,\mathrm{rad}$, enough to smear a narrow island across a substantial fraction of the plot. Interpolating trajectories to timestamps from a synchronised drive encoder controls that effect. Keep unwrapped relative phase as a separate channel to count actual slips.

A numerical area test on the stroboscopic map should use canonical action and angle or an equivalent canonical coordinate pair. A scatter plot of polar angle against angular velocity for a state-dependent inertia can distort area. Changes of pixel density, trajectory dwell time and resolution near the separatrix further affect visual appearance. To identify a resonance quantitatively, measure the action extent at the stable phase, the local libration period and the saddle crossing energy on independently prepared trajectories. Agreement among those three observations is less easily mimicked by a sampling artefact than the existence of a visually closed loop alone.

Recover a frequency shift by removing a secular term

Canonical averaging yields a frequency correction for the quartic oscillator. An independent derivation makes clear why a small difference in frequency can become a large trajectory error at long times. Write the equation as $\ddot{q} + \omega_0^2 q + (\beta/m) q^3 = 0$. Put $q = Ax$, $\tau = \omega t$, $\lambda = \beta A^2/(m\omega_0^2)$ and seek $\omega = \omega_0(1 + c\lambda + \cdots)$,

$$x = x_0 + \lambda x_1 + \cdots.$$

Leading initial conditions $x(0) = 1, x'(0) = 0$ give

$$x_0 = \cos\tau.$$

At first order one finds

$$x_1'' + x_1 = 2c\cos\tau - \cos^3\tau.$$

Using $\cos^3\tau = (3\cos\tau + \cos 3\tau)/4$, the right-hand side contains a term $(2c - 3/4)\cos\tau$ resonant with the leading motion. Choosing $c = 3/8$ removes the forcing that would otherwise generate a term proportional to $\tau\sin\tau$. The frequency then agrees with the action-angle calculation. The remaining particular solution is $\cos 3\tau/32$. Including the homogeneous term required by $x_1(0) = x_1'(0) = 0$ gives

$$x_1 = (\cos 3\tau - \cos\tau)/32.$$

This separates the leading frequency shift from an actual third-harmonic distortion of the orbit. Fitting only frequency and overlooking the small third harmonic cannot generally reproduce a precise waveform.

In the earlier example $\lambda = 0.040$, so the displacement amplitude at the third harmonic is $A\lambda/32 = 0.125\,\mathrm{mm}$ at $A = 0.10\,\mathrm{m}$. The frequency shift is about $0.15\,\mathrm{s}^{-1}$. After twenty seconds the uncorrected harmonic prediction is nearly three radians out of phase, far larger than the instantaneous third-harmonic displacement. A small-amplitude correction to the right-hand side can therefore cause an order-one error in the long-time phase. Measuring both the third harmonic and frequency gives two independent checks of the same inferred quartic stiffness. A drive electronics harmonic can mimic the former without producing the amplitude-dependent free-oscillation frequency.

Inspect growth on a parametrically driven orbit

Move the drive away from resonance and examine stroboscopic growth. Determine whether changing phase, amplitude or forcing frequency is responsible.

Open this experiment online

Try the calculation

Fractional stiffness modulation **0.08**

Damping γ (1/s) **0.1**

Drive frequency Ω (rad/s) **20**

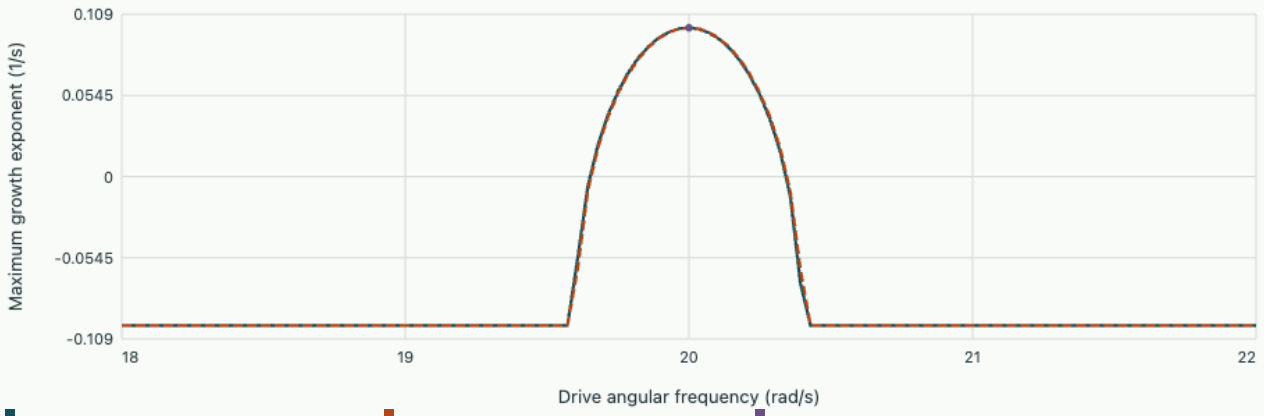

One-period numerical map · Weak-modulation average · Selected exponent

At Ω = 20 rad/s, numerical exponent 0.09997 s⁻¹ and averaged exponent 0.1 s⁻¹. Amplitude doubling time 6.934 s.

Test overlap of nearby resonances without declaring chaos from one curve

Take a rotor driven by two weak periodic patterns with distinct angular speeds Ω_1, Ω_2. For an isolated pattern j, the earlier calculation gives a resonant centre $P_j = \mathcal{I}\Omega_j$ and half-width $\Delta P_j = 2\sqrt{\mathcal{I}\epsilon V_j}$. Comparing the sum of estimated half-widths with the separation $|P_2 - P_1|$ is a practical indication of when the isolated-island approximation may fail. It does not by itself prove every trajectory between the resonances is chaotic. The relative phases, other harmonics and the actual time dependence determine the global structure.

For the same inertia $0.050\,\mathrm{kg\,m^2}$, choose $\Omega_1 = 8.0\,\mathrm{s^{-1}}$, $\Omega_2 = 8.6\,\mathrm{s^{-1}}$, and forcing energies 0.0020 and $0.00050\,\mathrm{J}$. The centres differ by $0.030\,\mathrm{kg\,m^2/s}$; the individual half-widths are 0.020 and $0.010\,\mathrm{kg\,m^2/s}$. Their nominal intervals just touch. Halving both forcing energies shrinks both widths by $1/\sqrt{2}$, leaving a gap. Doubling both makes the estimated intervals overlap. The speed difference corresponds to a beat period $2\pi/0.6 = 10.47\,\mathrm{s}$, another clock that must be resolved when recording nearby motion.

Transforming into the first rotating frame makes the first perturbation stationary and leaves the second with time dependence proportional to $\cos[\phi - (\Omega_2 - \Omega_1)t]$. A stroboscopic map sampled at the beat period is therefore more informative than a map with unsynchronised lab sampling. Prepare nearby initial conditions, compare their separation over multiple beat periods, and repeat with a refined symplectic time step and a modestly varied observation interval. Short-term stretching near an ordinary hyperbolic fixed point is possible even in a fully integrable pendulum. Evidence for widespread chaotic transport requires more than the growth of one separation during one saddle passage.

A student finds an apparent threshold by plotting whether the momentum ever crosses the midpoint between the two resonant centres. What additional evidence distinguishes irregular transport from a single large

regular libration?

▼ **Work through the solution**

A regular libration may cross that midpoint repeatedly inside one island, especially near its separatrix. Plot return pairs at a fixed beat phase, check for closed invariant curves or broader area filling, and track unwrapped relative phase to see whether it slips irregularly between the two patterns. Vary initial conditions and compare the statistical trend with the isolated-width estimate. Confirm that energy balance in the driven system includes explicit work by both patterns; a drift of the autonomous single-pattern energy is expected once the second drive is present.

Use phase reversal to isolate the resonant action transfer

For an initially weak sinusoidal drive, prepare two otherwise identical trajectories with initial relative phases differing by π. Their leading changes of action during a short resonant pulse have opposite signs, since $\dot{I}$ is proportional to $\sin\phi_0$. Half the difference isolates that phase-sensitive term, whereas half the sum can reveal phase-independent heating or higher-order effects. With drive coefficient $A = 0.001\,\mathrm{J}$, duration $2.0\,\mathrm{s}$ and initial phases $\pm\pi/2$, short-pulse first-order predictions differ by $0.0040\,\mathrm{J\,s}$. The test requires that detuning and nonlinear phase drift over the pulse remain small.

Perform the comparison at several durations. A linear growth of the early phase-sensitive difference is compatible with coherent transfer. Subsequent saturation, reversal or a slow oscillation can reflect the resonant pendulum dynamics once the action shift changes the phase. Friction can also saturate a response, but it produces mechanical energy loss and typically modifies the two phase preparations in different ways from conservative trapping. Measuring work from the applied torque and the full phase history adds the independent information needed to separate those explanations. A final amplitude alone gives neither the sign of the action transfer nor its phase dependence.

Tutorial on weak forcing

Far from resonance let a one-mode perturbation be $A\cos(\theta - \Omega t)$ with $A = 0.001\,\mathrm{J}$ and detuning $0.50\,\mathrm{s}^{-1}$. What first-order ripple scale should appear in the old action?

▼ **Work through the solution**

The generating function has amplitude $A/|\Delta| = 0.002\,\mathrm{J\,s}$. For harmonic number one its angle derivative has the same amplitude, so the old action oscillates about the averaged action on that scale. At detuning $0.01\,\mathrm{s}^{-1}$ the nominal ripple would be fifty times larger. Check its size against the action scale and the change of detuning across the ripple before trusting that continuation.

In the rotor example the forcing energy is multiplied by four. What happens to the resonant action, action width and small-amplitude libration period?

▼ **Work through the solution**

The central action is set by $H_0' = \Omega$ and stays $0.40\,\mathrm{kg\,m^2/s}$. The half-width doubles to $0.040\,\mathrm{kg\,m^2/s}$. The libration frequency doubles, so the period halves to $15.7\,\mathrm{s}$. These predictions apply to the exact sinusoidal rotor model; other harmonics of a physical drive may also change as its amplitude increases.

A slowly swept drive seems to carry a phase-locked rotor beyond its original resonance band. Is this inconsistent with the static separatrix estimate?

▼ **Work through the solution**

A varying $\Omega(t)$ moves the resonance centre $\mathcal{I}\Omega(t)$. In coordinates measured from that moving centre the equation for detuning momentum has an additional torque-like term $-\mathcal{I}\dot{\Omega}$. Locking can persist if the available restoring torque can track the ramp and the initial motion stays inside the evolving island. The static width describes a fixed-frequency section. Comparing a captured trajectory to its initial centre after the sweep tests the wrong phase portrait.

Compare a resonance driven by a varying stiffness with a rotating additive potential.

David Tong's Hamiltonian course develops canonical variables and adiabatic invariants.

Applied study

Connect confinement, quantised energies and measurement probabilities through wavefunctions.

What information does a quantum state contain?

Suppose an electron is confined to a narrow region and a sequence of energy measurements shows distinct lines. To interpret them, distinguish a state from an outcome. In one spatial dimension a pure state has a complex amplitude $\psi(x,t)$ normalised so that $\int |\psi|^2 dx = 1$. The Born rule gives the probability in an interval by integrating $|\psi|^2$ there. Relative phases affect interference between amplitudes.

An observable is represented by a self-adjoint operator with suitable boundary conditions. The Hamiltonian for a nonrelativistic particle of mass m in a scalar potential V is $\hat{H} = -(\hbar^2/2m)\partial_x^2 + V$. It both sets the possible energy outcomes and generates the time evolution $i\hbar\partial_t\psi = \hat{H}\psi$.

Derive a spectrum from confinement

Take an infinite square well on $0 < x < L$. The potential vanishes inside and hard walls impose $\psi(0) = \psi(L) = 0$. Separating $\psi(x,t) = \phi(x)e^{-iEt/\hbar}$ reduces the equation to $-\phi'' = k^2\phi$, $E = \hbar^2k^2/(2m)$. The left wall selects a sine. The right wall fixes $k_n = n\pi/L$. The integer starts at one because the zero solution cannot be normalised. Integrating $\sin^2$ over the well fixes the normalisation

$$\phi_n(x) = \sqrt{\frac{2}{L}}\sin\frac{n\pi x}{L},$$

$$E_n = \frac{\hbar^2\pi^2 n^2}{2mL^2}.$$

Take an electron and $L = 1.00$ nm. With $\hbar = 1.05457 \times 10^{-34}$ J·s, $m_e = 9.10938 \times 10^{-31}$ kg and 1 eV $= 1.60218 \times 10^{-19}$ J, the ground energy is 0.376 eV and the second energy is 1.504 eV. The ground state is lowest in energy because it has the gentlest curvature compatible with both walls.

Separate energy measurements and time evolution

Prepare $\psi = (\phi_1 + \phi_2)/\sqrt{2}$. Orthogonality makes $|\langle\phi_n|\psi\rangle|^2 = 1/2$ for each of these two energy outcomes. An ideal projective measurement giving E_1 leaves the state in ϕ_1; an immediate repetition gives E_1 with probability one. Measurements on an ensemble of fresh preparations give mean energy $(E_1 + E_2)/2 = 0.940$ eV.

Before any measurement the two components evolve with different phases. For these real eigenfunctions the density contains the cross term $\phi_1\phi_2\cos[(E_2 - E_1)t/\hbar]$. A position detector can therefore observe a period $h/(E_2 - E_1) = 3.67$ fs for the 1 nm well. Each energy probability remains constant. This is the distinction between changing local probability and conserved energy in a time-independent Hamiltonian.

Transfer the measurement rule to spin

A spin-½ atom prepared in the upper output of a Stern–Gerlach analyser along z enters another analyser oriented along x. In the two-dimensional spin space,

$$|+z\rangle = (|+x\rangle + |-x\rangle)/\sqrt{2}.$$

The second analyser therefore sends half of a fresh ensemble into each output. Retain only the upper x beam and measure along z again. Since $|+x\rangle = (|+z\rangle + |-z\rangle)/\sqrt{2}$, half of that selected beam now enters the lower z output, or one quarter of the original flux. With a second z analyser directly after the first, this lower output would have been empty. A change of measurement basis has observable consequences for the subsequent statistics.

How much geometry survives in a spectrum?

Inside a two-dimensional hard-wall region the same reasoning gives $-\nabla^2\phi_n = k_n^2\phi_n$ with $\phi_n = 0$ on its boundary. A membrane fixed along its edge obeys the same spatial equation with $k_n^2 = \rho_A\omega_n^2/\mathcal{T}$, where ρ_A is mass per area and $\mathcal{T}$ tension per unit length. This correspondence gives a numerical way to find the quantum energy levels of a freely moving particle confined within corresponding walls. Two distinct polygons can have identical lists of these eigenvalues. Mode shapes and local probabilities can still differ. Rigid-wall acoustic pressure modes illustrate how changing the boundary condition changes the spectrum. Standing waves in a closed room.

Exercises on measurement and scaling

1. Predict a local probability

At preparation time, what fraction of the equal, in-phase superposition in the well lies in its left half? What happens after half the beat period?

▼ **Integrate the interference term**

Each individual eigenstate has probability $1/2$ on the left. The cross-term integral is $\int_0^{L/2} \phi_1\phi_2\,dx = 4/(3\pi)$, using the difference of two cosines for the product of sines. Thus $P_{\text{left}} = 1/2 + 4/(3\pi) \simeq 0.924$ initially. Half a beat later the relative phase has changed by π and the probability is about 0.076. Both energy populations are still $1/2$.

---

2. Identify the shifted spectral line

Double the width of the well while keeping the electron mass fixed. How do the lowest energy difference and the superposition's beat period change?

▼ **Relate line spacing and time scale**

Each energy scales as L^{-2}, so the gap falls from 1.128 to 0.282 eV. The beat period scales as the inverse gap and grows from 3.67 to 14.7 fs. Scaling every coordinate of a two-dimensional hard-wall model by the same factor produces the same energy scaling.

ONESIMULATION LABORATORY

Turn spatial eigenvalues into confinement energies

Compare recorded OneSimulation runs on this page. To test your own inputs, download the notebook and model files and run the model in an activated OneSimulation environment.

The isospectral drums notebook solves the Dirichlet eigenproblem on two polygons. It reports `lam`, the dimensionless squared wave number $\lambda_n = k_n^2 \ell^2$. Its length unit is 0.1 m for the drum. To interpret a geometrically similar quantum well, assign that unit a quantum length ℓ in metres. Read the shapes and eigenvalues from the spatial solve; quantum time evolution and measurement probabilities for a chosen state follow analytically. Dirichlet spectra for two different outlines.

Run through the notebook's results, then add this cell. It converts all ten computed eigenvalues into electron energies. `sizes[-1]` selects the supplied 0.005 m mesh.

```
from scipy.constants import hbar, electron_mass, electron_volt

def energies_eV(shape, length_unit_m):
    lam = runs[shape, sizes[-1]]["lam"]
    k_squared = lam / length_unit_m**2              # 1/m²
    return hbar**2 * k_squared / (2 * electron_mass * electron_volt)

print(energies_eV("drum A", 1e-9))
print(energies_eV("drum B", 1e-9))
print(energies_eV("drum A", 2e-9))
```

1. **Baseline.** Assign 1 nm to the length unit in shape A. Check that the lowest dimensionless value agrees with 2.537944 within 0.01%, giving E_1 near 0.09670 eV. The 1 nm unit sets the geometric scale of the polygon.
2. **Change the outline.** Select shape B at the same scale. Compare the ground energies within 0.01% and inspect the first-mode maps. Equal eigenvalues can coexist with different spatial distributions.
3. **Change the confinement length.** Keep shape A and set its quantum length unit to 2 nm. Check $E_1(2\,\mathrm{nm})/E_1(1\,\mathrm{nm}) = 1/4$ and verify the ratio for an excited state. No new mesh is required for this exact uniform scaling.

Numerical Dirichlet eigenvalues mapped to electron ground-state energies

Quantum well	Lowest eigenvalue λ_1	Mapped energy (eV)
Shape A, 1 nm unit	2.537961	0.09669586

Quantum well	Lowest eigenvalue λ_1	Mapped energy (eV)
Shape B, 1 nm unit	2.537966	0.09669606
Shape A, 2 nm unit	2.537961	0.02417397

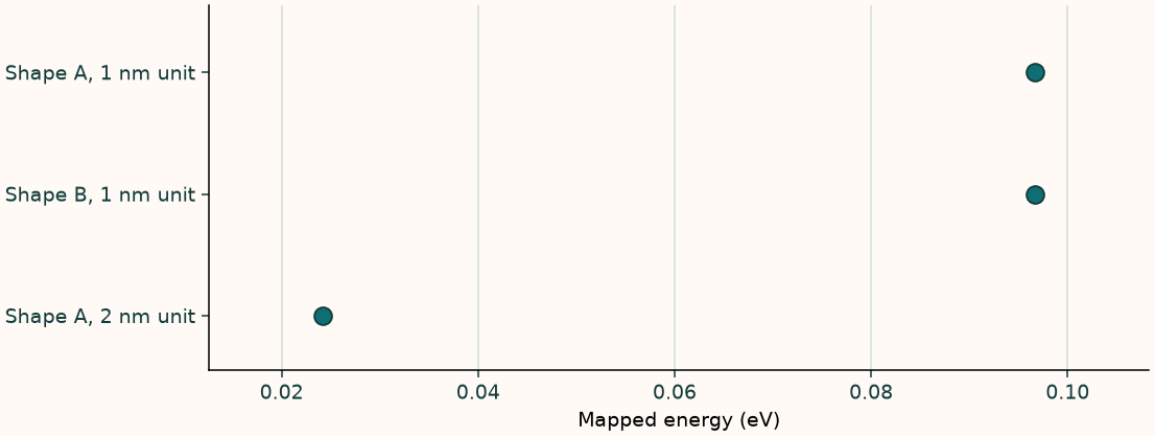

Both polygonal wells give nearly the same ground energy at the 1 nm length scale. Doubling the scale quarters the energy inferred from the spatial eigenvalue.

Download experiment files (ZIP)

Extract the files together and open `isospectral_drums.ipynb` in a local OneSimulation environment.

Problem sheet 6 · Which energy change really establishes a length change?

Notation and assumptions

Read each defining equation with its assumptions and units, then follow the link to the argument that gives the equation its meaning.

Commutator

$$[A, B] = AB - BA$$

Order matters. The expectation of the commutator controls a lower bound on the product of uncertainties when the state lies in the relevant operator domains.

Derivation · Operators and the quantum harmonic oscillator

Density operator

$$\rho = \sum_j p_j |\psi_j\rangle\langle\psi_j|$$

A physical state is positive and has trace one. Different statistical preparations can give identical density operators and measurement predictions.

Derivation · Density matrices and loss of phase

Probability current

$$j = (\hbar/m)\,\mathrm{Im}(\psi^* \partial_x \psi)$$

This expression assumes a scalar particle with no vector potential. Transmission probabilities require outgoing relative to incoming current, even when their wave numbers differ.

Derivation · Scattering and tunnelling in one dimension

Scattering length

$$k \cot \delta_0 = -1/a + O(k^2)$$

This sets the low-energy s-wave scattering scale for a short-range potential. Its sign carries phase information, and its size can greatly exceed the interaction range.

Derivation · Measure an interaction by its low-energy scattering length

State normalisation

$$\int |\psi(x)|^2 \, dx = 1$$

For a normalisable one-dimensional state, squared modulus is probability per length. Stationary incoming scattering states require a flux or continuum normalisation.

Derivation · Wave packets and probability current

Variational energy

$$E[\psi] = \langle \psi | H | \psi \rangle / \langle \psi | \psi \rangle$$

For an admissible nonzero trial state and a Hamiltonian bounded below, this is an upper bound on the lowest energy. A good energy can hide a less accurate state.

Derivation · Estimating a ground-state energy

Further reading

- Use the one-dimensional potential and scattering problems to practise matching wavefunctions, normalising flux and predicting measurement outcomes before evaluating the amplitudes. MIT · Quantum Physics I

Compare specific topics with the university sources