Thermodynamics and statistical physics

Trace energy and entropy through processes, derive thermodynamic properties from microscopic states, and predict diffusion, heat capacity and equilibration.

Read the subject in depth

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

- **Statistical mechanics from microscopic counting to material response**

  Count states, construct ensembles and derive macroscopic predictions. Use fluctuations, time dependence and finite sample size to test how well a thermodynamic limit describes a measurement.

- **Counting, correlations and thermal measurements**

  Calculate an observable from an explicitly countable microscopic model. Determine which finite-size, sampling or instrumental effect changes a measured equilibrium or relaxation law.

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.

Compare the ensemble construction with the particle-counting problems. The applications to electron and phonon gases supply a second derivation of the low-temperature heat capacities. Compare the choice of potential when passing to coexistence. Torsten Fließbach · Statistische Physik, 6th edition

1. **Reservoirs, useful work and mixtures**

   Keep track of which reservoir fixes temperature, chemical potential or pressure. Close the energy and entropy balances before inferring the useful work or the coexisting phase fraction.

2. **Heat stored in quantum degrees of freedom**

   Build the single-particle spectrum before taking the thermal average. Compare acoustic branches, magnetic ordering and paired electrons, and trace how a radiometer samples a photon reservoir.

3. **Irreversible heat transfer and fluctuating signals**

   Measure a transport coefficient using a steady gradient and a decay. Derive the fluctuation that survives in equilibrium and work out which part of it a finite-bandwidth instrument measures.

4. **Ensembles, diffusion and finite data**

   Contrast conservation along a microscopic orbit with apparent relaxation after coarse sampling. Test a diffusion law and infer an unknown thermal state from a sequence of noisy observations.

How the course fits together

1. **Chapters 1–4**

   Describe a system by its equilibrium state, then keep accounts of energy exchanged as heat and work. Entropy selects accessible changes. Counting microscopic states explains why two bodies in contact

approach a common temperature.

2. **Chapters 5–8**

   Build a partition sum from permitted energies to predict averages. Particle statistics change the occupancy of those levels. Following moving particles and repeated collisions then connects equilibrium averages to transport rates.

3. **Chapters 9–12**

   Apply state counting to photons, lattice vibrations and chemical composition. Finally add a finite transfer rate to the equilibrium description and predict both the direction of change and the time taken to relax.

Contents

1. Chapter 1 · Equations of state and response functions
2. Chapter 2 · Heat, work and the limits of a cycle
3. Problem sheet 1 · State variables and energy conversion
4. Chapter 3 · Thermodynamic potentials and phase boundaries
5. Chapter 4 · Counting states and defining temperature
6. Problem sheet 2 · Potentials, phase lines and counting
7. Chapter 5 · Partition functions and thermal fluctuations
8. Chapter 6 · Indistinguishable particles and quantum gases
9. Problem sheet 3 · Populations and quantum statistics
10. Chapter 7 · Pressure and effusion from molecular motion
11. Chapter 8 · Diffusion from a random walk
12. Problem sheet 4 · Molecular leaks and diffusing tracers
13. Chapter 9 · Thermal radiation in a cavity
14. Chapter 10 · Vibrations and the Debye heat capacity
15. Problem sheet 5 · Radiation and vibrations as mode gases
16. Chapter 11 · Chemical equilibrium and reaction extent
17. Chapter 12 · Entropy production and relaxation
18. Problem sheet 6 · Equilibrium, kinetics and irreversibility
19. Chapter 13 · How a finite heat bath changes a fluctuation
20. Chapter 14 · Counting occupied sites on a surface
21. Chapter 15 · How a pair potential changes the gas pressure
22. Chapter 16 · Why two coexisting phases share a tangent
23. Chapter 17 · Extracting a length from thermal spin fluctuations
24. Chapter 18 · Why only a narrow shell of a Fermi gas carries heat
25. Chapter 19 · Mean-field ordering, response and the critical point
26. Chapter 20 · Design a thermometer with a countable energy spectrum
27. Problem sheet 7 · Can a fluctuation reveal attraction on a surface?
28. Problem sheet 8 · Count atoms in a trap and predict the onset of condensation
29. Statistical mechanics from microscopic counting to material response
30. Counting, correlations and thermal measurements
31. Reservoirs, useful work and mixtures
32. Heat stored in quantum degrees of freedom

33. Irreversible heat transfer and fluctuating signals
34. Ensembles, diffusion and finite data
35. Applied study
36. Notation and assumptions
37. Further reading

Chapter 1 · Equations of state and response functions

Use compressibility and thermal expansion to predict what a pressure gauge will record when a material is heated in a closed vessel.

FROM FIRST PRINCIPLES

Derive ideal-gas pressure from impacts on a wall

Pressure is the normal force on a wall per unit area. Consider N classical molecules of mass m in a cubic box of side L. Neglect intermolecular forces and molecular volume. For an elastic collision at a wall perpendicular to the x axis, a molecule transfers normal momentum $2m|v_x|$. It visits the same wall once in a round trip of duration $2L/|v_x|$. Its time-averaged force there is

$$F_{x,1} = \frac{2m|v_x|}{2L/|v_x|} = \frac{mv_x^2}{L}.$$

Sum all contributions and divide by the wall area L^2. With $V = L^3$, equilibrium isotropy gives

$$\langle v_x^2 \rangle = \langle v^2 \rangle / 3.$$

Classical thermal equilibrium assigns mean energy $k_B T/2$ to each quadratic velocity component. These inputs give

$$p = \frac{Nm}{V}\langle v_x^2 \rangle = \frac{Nm}{3V}\langle v^2 \rangle,$$
$$pV = Nk_B T.$$

This equation of state connects three directly measurable quantities at fixed particle number. The kinetic calculation also identifies where a real gas can differ: intermolecular forces transfer momentum, and dense molecules occupy an appreciable fraction of the volume.

DERIVATION AND APPLICATIONS

Connect a sound-speed measurement to two compressibilities

A slowly compressed sample has time to exchange heat with its surroundings. A small sound pulse usually compresses a fluid element faster than heat can diffuse across a wavelength. The elastic response relevant to sound is then nearly adiabatic. Consider a homogeneous single-phase fluid at equilibrium and a small-amplitude pulse. Assume negligible viscosity and thermal diffusion over one oscillation, so the specific entropy of a moving element remains constant.

Let v be volume per unit mass, c_p the specific heat at constant pressure, $\alpha = v^{-1}(\partial v/\partial T)_p$, and $\kappa_T = -v^{-1}(\partial v/\partial p)_T$. From the specific Gibbs differential $dg = -s\,dT + v\,dp$, equality of mixed derivatives gives

$$(\partial s/\partial p)_T = -(\partial v/\partial T)_p.$$

The entropy differential determines a fluid element's temperature change under isentropic compression,

$$ds = \frac{c_p}{T}\,dT - v\alpha\,dp = 0, \qquad \left(\frac{\partial T}{\partial p}\right)_s = \frac{Tv\alpha}{c_p}.$$

Insert this derivative into the volume response $dv/v = \alpha\,dT - \kappa_T\,dp$. The resulting isentropic compressibility is

$$\kappa_s = -\frac{1}{v}\left(\frac{\partial v}{\partial p}\right)_s = \kappa_T - \frac{Tv\alpha^2}{c_p} = \kappa_T\frac{c_v}{c_p}.$$

The last step uses the heat-capacity difference for a simple compressible material. For $\rho = 1/v$, a density perturbation gives

$$p' = (\partial p/\partial \rho)_s \rho'.$$

Linear momentum balance and continuity then combine into a wave equation with speed c,

$$\partial_t^2 \rho' = c^2 \nabla^2 \rho',$$
$$c^2 = \left(\frac{\partial p}{\partial \rho}\right)_s = \frac{1}{\rho\kappa_s}.$$

Adiabatic heating adds a pressure response on compression for an ordinary expanding fluid. The two compressibilities coincide where thermal expansion vanishes.

WORKED EXAMPLE

Predict the speed in an ideal gas

Change the inputs and investigate this derivation.

For dry model air at $p = 1.00 \times 10^5\,\mathrm{Pa}$, take $\rho = 1.20\,\mathrm{kg/m^3}$ and $c_p/c_v = 1.40$. Since $\kappa_T = 1/p$, sound propagates at $\sqrt{1.40p/\rho} = 342\,\mathrm{m/s}$. An isothermal model gives $289\,\mathrm{m/s}$. Comparing the two speeds measures the square root of the heat-capacity ratio under these gas assumptions.

SECOND APPLICATION

Infer a heat-capacity ratio for a liquid

Change the inputs and investigate this derivation.

An ultrasonic pulse travels at $1480\,\mathrm{m/s}$ through a liquid of density $998\,\mathrm{kg/m^3}$. It gives

$$\kappa_s = 1/(998 \times 1480^2) = 4.57 \times 10^{-10}\,\mathrm{Pa^{-1}}.$$

A slow isothermal experiment gives $4.65 \times 10^{-10}\,\mathrm{Pa^{-1}}$. Their ratio yields $c_p/c_v = 1.017$, which explains why the choice of compressibility produces a much smaller acoustic correction for this liquid.

Try it · Sound compressibility

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

Density (kg/m³) **1.2**

Heat-capacity ratio cp/cv **1.4**

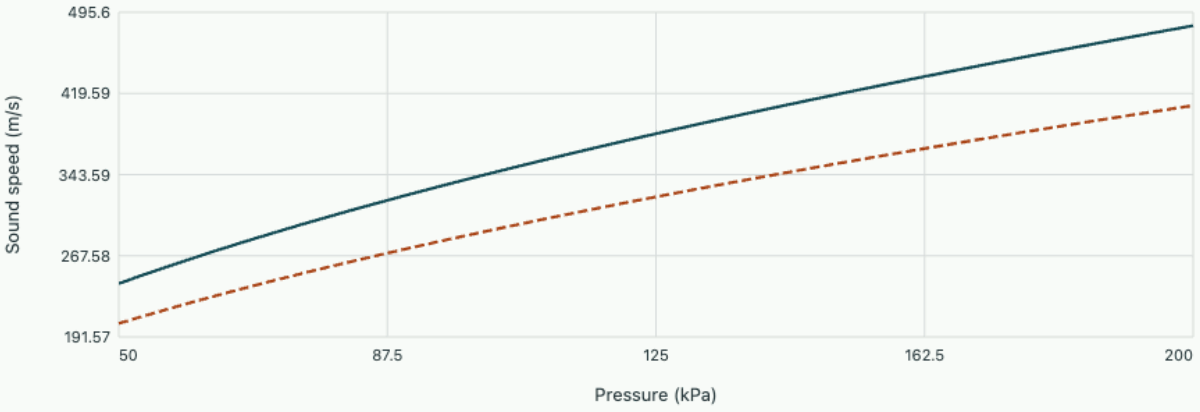

Adiabatic response Isothermal response

For ideal gas κT=1/p, κs=κT/γ. At 100 kPa the speed is 341.57 m/s. A liquid instead requires measured compressibility.

Decide which variables specify a state

An equilibrium sample of one simple phase and fixed composition can be described by pressure, volume and temperature related by an equation of state. A piston experiment can prescribe pressure while changing volume; a sealed rigid vessel instead fixes volume. The distinction matters because the same temperature change can lead to expansion in one apparatus and a large pressure increase in the other. Assume quasistatic measurements so that a single pressure and temperature describe the sample.

An equation of state is a constitutive statement. The ideal-gas relation $pV = nRT$ treats particles as dilute enough for their mutual interactions to have negligible effect on the measured pressure. An incompressible-liquid approximation often gives an adequate volume balance at a chosen pressure but can make a constrained heating problem ill posed if applied literally. Both small thermal expansion and small but finite compressibility may be needed to compute the pressure response.

Describe a neighbourhood of equilibrium

Treat volume as a differentiable function $V(T, p)$ near a reference state. Its total differential contains the two independent slopes $(\partial V/\partial T)_p$ and $(\partial V/\partial p)_T$. Define the two responses at their respective fixed controls,

$$\alpha = \frac{1}{V}\left(\frac{\partial V}{\partial T}\right)_p,$$
$$\kappa_T = -\frac{1}{V}\left(\frac{\partial V}{\partial p}\right)_T.$$

Ordinary mechanically stable material contracts when pressure increases, which gives positive isothermal compressibility. Inserting both definitions into the volume differential gives

$$\frac{dV}{V} = \alpha\, dT - \kappa_T\, dp.$$

Both terms are dimensionless. Expansion has units of inverse kelvin, while compressibility has units of inverse pascal. For fixed volume the pressure slope is

$$(\partial p/\partial T)_V = \alpha/\kappa_T.$$

This relation is a consequence of using the same volume differential. It does not assume a particular molecular model and it remains informative for a liquid for which the ideal-gas equation would give a poor answer.

For an ideal gas at fixed amount, differentiating $V = nRT/p$ gives

$$\alpha = 1/T,$$
$$\kappa_T = 1/p.$$

Their ratio is p/T, as follows independently from $p = nRT/V$. At 300 K and 100 kPa, warming a sealed ideal gas by 3.0 K raises the pressure by 1.0 kPa. At fixed pressure instead the relative volume increase is one percent.

Background · Why do slow compression tests and sound measure different moduli?

For a sealed ideal gas in thermal contact with a bath, a sufficiently slow compression allows heat to escape and fixes T. Then pV is constant. During a short, nearly reversible acoustic oscillation in the bulk there can be too little time for heat to leave the compressed fluid element. For this adiabatic limit pV^{γ} is constant. Differentiate at the same initial pressure.

$$K_T = -V\left(\frac{\partial p}{\partial V}\right)_T = p,$$

$$K_S = -V\left(\frac{\partial p}{\partial V}\right)_S = \gamma p.$$

At one bar and $\gamma = 1.4$ the moduli are 100 and 140 kPa. Each measures the response at the same equilibrium state, under a different thermal constraint. In a narrow tube heat can diffuse to the wall during a cycle. The appropriate limit depends on whether the thermal diffusion length is small or large compared with the tube radius.

Separate intensive, extensive and measured responses

Doubling the amount of material at the same temperature and pressure doubles volume but leaves the two normalised coefficients unchanged. Such intensive coefficients can be compared across samples. A vessel's own volume, however, can respond to both pressure and temperature. If its expansion coefficients are appreciable, the measured pressure slope describes the coupled sample and vessel. One cannot simply substitute the unconstrained expansion of the sample into a rigid-vessel formula.

To see this, let the vessel's fractional volume response be $\alpha_v\,dT + C_v^{\mathrm{vessel}}dp$, where its mechanical compliance is positive for an internal pressure increase. Equal fractional changes of vessel and sample give

$$(\alpha - \alpha_v)dT = (\kappa_T + C_v^{\mathrm{vessel}})dp.$$

A compliant wall reduces the pressure rise, and a vessel that expands just as much as the sample can nearly cancel the temperature signal. Here the symbol for vessel compliance denotes a mechanical quantity and has different units from heat capacity.

Compare liquid and container data

Take a representative liquid with $\alpha = 2.5 \times 10^{-4}$ K⁻¹ and $\kappa_T = 5.0 \times 10^{-10}$ Pa⁻¹ near its operating point. In an ideally rigid container, a 2.0 K increase gives $\Delta p \simeq 1.0$ MPa. The unconstrained volume change at fixed pressure would be only 0.050%. This combination explains how a small fractional expansion can produce a significant pressure response when the fluid is only weakly compressible.

If the container has $\alpha_v = 5.0 \times 10^{-5}$ K⁻¹ and $C_v^{\mathrm{vessel}} = 3.0 \times 10^{-10}$ Pa⁻¹, the same warming gives about 0.50 MPa. Inverting a measured slope without the vessel data could thus give a misleading estimate of compressibility. These are local linear predictions. Over a broad pressure or temperature range the

coefficients must be evaluated along the path, and crossing a phase boundary can invalidate the single-phase differential altogether.

Convert an isobaric calorimeter reading into an isochoric one

A sample can absorb different amounts of heat per kelvin depending on which variable is held fixed. For a fixed amount of a simple, single-phase material in equilibrium,

$$C_V = T(\partial S/\partial T)_V,$$
$$C_p = T(\partial S/\partial T)_p.$$

Apply the chain rule to $S(T,V)$ along the constant-pressure path. Subtracting the two derivatives leaves the effect of expansion. From $dF = -S\,dT - p\,dV$ equality of mixed derivatives gives

$$(\partial S/\partial V)_T = (\partial p/\partial T)_V.$$

Therefore

$$C_p - C_V = T\left(\frac{\partial p}{\partial T}\right)_V \left(\frac{\partial V}{\partial T}\right)_p = \frac{TV\alpha^2}{\kappa_T}.$$

Here α is the volume expansion coefficient and κ_T the isothermal compressibility. The last equality uses $dV/V = \alpha dT - \kappa_T dp$. For a mechanically stable material $\kappa_T > 0$, so the difference is nonnegative even for a material that contracts when warmed. For one mole of ideal gas, inserting $\alpha = 1/T$, $\kappa_T = 1/p$ and $pV = RT$ gives

$$C_p - C_V = R.$$

At 300 K a solid has molar volume 1.20×10^{-5} m³/mol, volumetric expansion coefficient 6.00×10^{-5} K⁻¹ and bulk modulus $B_T = 1/\kappa_T = 100$ GPa. The molar difference is $TV_mB_T\alpha^2 = 1.30$ J/(mol K). A constant-pressure measurement of 25.0 J/(mol K) thus implies $C_{V,m} = 23.7$ J/(mol K). For a cubic solid the linear expansion coefficient is approximately one third the volume coefficient. Using it directly in this formula would underestimate the correction ninefold. This conversion matters when comparing a tabulated calorimetric result with a vibrational model calculated at fixed volume.

Practise on sheet 1 · Test a calorimeter conversion near zero expansion. Further reading · University of Edinburgh · Junior Honours Thermodynamics · Topic 9, “How to find the difference in heat capacities, CP – CV”.

Predict the temperature change of a throttled gas

Some pressure changes take place in flowing rather than closed matter. Consider a well-insulated porous plug that sustains a pressure difference between two steady streams of the same single-phase gas. There is no shaft work and inlet and outlet speeds and heights have negligible differences in kinetic and gravitational

energy. Flow work at each end is already included in the molar enthalpy $h = u + pv$. The steady energy balance then gives

$$h_{\mathrm{out}} = h_{\mathrm{in}}.$$

A pressure drop through the restriction can nevertheless change temperature if the gas is nonideal.

From $dh = T\,ds + v\,dp$ and the Maxwell relation $(\partial s/\partial p)_T = -(\partial v/\partial T)_p$, write the enthalpy differential in terms of T, p. The coefficient of dT at fixed pressure is the molar heat capacity c_p. Setting $dh = 0$ gives the Joule–Thomson coefficient

$$\mu_{\mathrm{JT}} = \left(\frac{\partial T}{\partial p}\right)_h = \frac{T(\partial v/\partial T)_p - v}{c_p} = \frac{v(T\alpha - 1)}{c_p}.$$

For an ideal gas $\alpha = 1/T$, so an isenthalpic throttle produces no temperature change. Isentropic expansion in a turbine is a different process. A throttle's pressure drop is irreversible; constant enthalpy follows here from the steady energy balance and does not imply constant entropy.

Estimate a small temperature drop from gas data

Near a chosen operating point let the low-pressure molar volume be $v = RT/p + B(T)$, retaining a second virial coefficient $B = b - a/(RT)$. Here a and b characterize an approximate real-gas equation of state. The ideal-volume term cancels in $T(\partial v/\partial T)_p - v$. The remaining numerator is $TB' - B = 2a/(RT) - b$. At fixed local heat capacity, attractions and excluded volume therefore compete in determining the sign of the temperature response.

Take $a = 0.140$ Pa m⁶/mol², $b = 4.00 \times 10^{-5}$ m³/mol, $T = 300$ K and $c_p = 29.0$ J/(mol K). Then $2a/(RT) = 1.12 \times 10^{-4}$ m³/mol and $\mu_{\mathrm{JT}} = 2.49 \times 10^{-6}$ K/Pa, or 0.249 K/bar. A pressure change of -2.0 bar predicts temperature change -0.50 K if the properties remain close to these values. The leading model changes sign at $T = 2a/(Rb) = 842$ K. This inversion estimate should not be extrapolated into a dense or multiphase region. A proper finite pressure drop integrates the measured local coefficient along an isenthalp. Recording both enthalpy balance and inlet-to-outlet entropy increase offers an independent test of whether unwanted heat transfer through the apparatus affected the result.

The entropy change of the throttled stream can be checked without a calorimeter. Along an isenthalp the Gibbs relation gives

$$ds = -v\,dp/T.$$

Since molar volume and temperature are positive, any pressure decrease raises the molar entropy. Even when an ideal gas shows zero temperature change, its pressure decrease leaves a positive entropy difference. This reveals precisely why absence of observed cooling is insufficient evidence for thermodynamic reversibility. A pressure reading, inlet temperature and appropriate volume data together supply a sign test independently of the small temperature slope used to estimate the Joule–Thomson coefficient.

An independent heat-leak estimate is particularly important near inversion, where a small parasitic heat flow can dominate the intrinsic temperature change.

Find the mechanically unstable part of an isotherm

Use the reduced van der Waals equation $p = 8t/(3v - 1) - 3/v^2$. At fixed temperature a stable homogeneous phase must have positive isothermal compressibility, which requires a negative slope on this graph. Locate any interval in which the model reverses that slope.

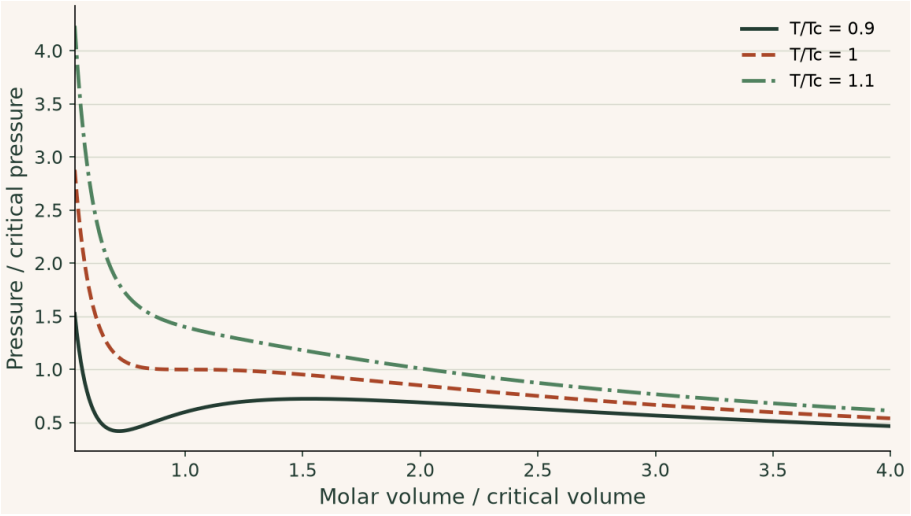

Below the critical temperature, larger pressures along the rising segment of the formal loop correspond to larger volumes. This implies negative isothermal compressibility. The homogeneous state is mechanically unstable there, and equilibrium coexistence replaces part of the curve with a constant-pressure tie line.

Try it · Vdw

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

Reduced temperature T/Tc **0.9**

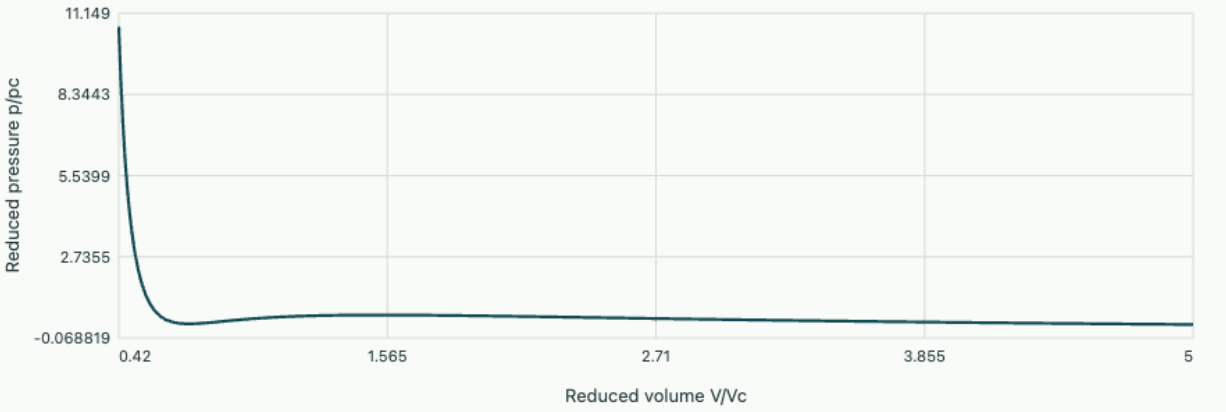

Homogeneous van der Waals isotherm

The derivative is $-24Tr/(3Vr-1)^2 + 6/Vr^3$. A positive derivative means negative compressibility. Below Tc equilibrium requires a coexistence construction; the displayed loop is the homogeneous equation.

▶ **Read values from the plot**

Study the full derivation

Pair correlations, the virial gas and a pressure measurement · Statistical mechanics from microscopic counting to material response

A dilute nonideal gas can display almost ideal pressure while its particles strongly avoid one another at short distance. Connect a two-particle interaction to pressure, internal energy and a measurable pair correlation. Then use the resulting integral equations to ask which features of a potential can actually be inferred from a limited set of pressure readings.

Chapter 2 · Heat, work and the limits of a cycle

Track changes of internal energy around a gas cycle and derive the temperature bound on the efficiency of a reversible engine.

FROM FIRST PRINCIPLES

Derive the reversible engine efficiency from two balances

A working substance in a cyclic engine returns to its initial state. Its internal energy is a state function, so the energy gained from a hot reservoir must equal the useful work plus the heat delivered to a cold reservoir. Let Q_h, Q_c be positive heat magnitudes and W positive work delivered by the engine. Energy conservation gives

$$W = Q_h - Q_c,$$
$$\eta = \frac{W}{Q_h} = 1 - \frac{Q_c}{Q_h}.$$

To determine the smallest possible heat rejection, apply the second law. Reservoirs at constant absolute temperatures $T_h > T_c$ lose entropy Q_h/T_h and gain entropy Q_c/T_c respectively. The engine itself returns to its original entropy. The total entropy production is therefore

$$S_{\mathrm{gen}} = \frac{Q_c}{T_c} - \frac{Q_h}{T_h} \geq 0, \qquad \eta \leq 1 - \frac{T_c}{T_h}.$$

A reversible cycle attains equality. Friction or heat transfer across finite temperature differences produces entropy and reduces efficiency at fixed reservoir temperatures. The ratio uses absolute temperatures because reservoir entropy changes are defined by reversible heat divided by thermodynamic temperature.

DERIVATION AND APPLICATIONS

Derive the efficiency of an idealized four-stroke gas cycle

A compression ratio predicts an efficiency for a model cycle only after the heat-transfer paths are specified. Consider a fixed amount of ideal gas with constant molar heat capacities and $\gamma = C_p/C_V > 1$. In the Otto model the gas undergoes reversible adiabatic compression from state 1 to 2, heat addition at fixed volume from 2 to 3, reversible adiabatic expansion from 3 to 4, then heat rejection at fixed volume back to 1. Use positive magnitudes Q_h, Q_c for the absorbed and rejected heat.

For either adiabat the first law gives

$$nC_V dT = -p\,dV = -nRT\,dV/V.$$

Separation and integration give

$$TV^{\gamma-1} = \text{constant}.$$

Introducing $r = V_1/V_2 = V_4/V_3 > 1$ sets both adiabatic temperature ratios,

$$\frac{T_2}{T_1} = \frac{T_3}{T_4} = r^{\gamma-1}.$$

The two fixed-volume steps do no boundary work. Their respective heats are internal energy changes, and cyclic energy conservation converts the heat ratio to efficiency,

$$Q_h = nC_V(T_3 - T_2),$$
$$Q_c = nC_V(T_4 - T_1),$$
$$\eta = 1 - \frac{Q_c}{Q_h}.$$

Substitute $T_3 = r^{\gamma-1}T_4$ and $T_2 = r^{\gamma-1}T_1$. Both temperature differences contain the same factor $T_4 - T_1$, which cancels and leaves

$$\eta = 1 - r^{1-\gamma}.$$

The gas absorbs and rejects heat over intervals of temperature. Its efficiency therefore depends on the chosen paths even with reversible adiabatic strokes. At $r = 1$ the pressure-volume loop collapses and produces no net work. Increasing the compression ratio also raises the temperature reached before heat addition, which can invalidate a constant-γ model.

WORKED EXAMPLE

Close a cycle with measured endpoint temperatures

Change the inputs and investigate this derivation.

For one mole with $r = 8.00$, $\gamma = 1.40$, $T_1 = 300\,\mathrm{K}$ and $T_3 = 1600\,\mathrm{K}$, the adiabatic factor is 2.297. Hence $T_2 = 689\,\mathrm{K}$, $T_4 = 696\,\mathrm{K}$, and $C_V = R/0.40 = 20.8\,\mathrm{J/(mol\,K)}$. The gas absorbs $18.9\,\mathrm{kJ}$ and rejects $8.24\,\mathrm{kJ}$, yielding $10.7\,\mathrm{kJ}$ of work and efficiency 0.565.

SECOND APPLICATION

Infer a compression ratio for a different working gas

Change the inputs and investigate this derivation.

A model monatomic gas has $\gamma = 5/3$. Achieving efficiency 0.600 requires

$$r = (1 - 0.600)^{-1/(\gamma-1)} = 3.95.$$

Starting at $280\,\mathrm{K}$, its temperature at the end of compression is

$$280/(1 - 0.600) = 700\,\mathrm{K}.$$

A cycle whose hottest temperature were $650\,\mathrm{K}$ could not realize the specified heat-addition stroke with this initial state.

Try it · Otto

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

Heat-capacity ratio γ **1.4**

Compression ratio **8**

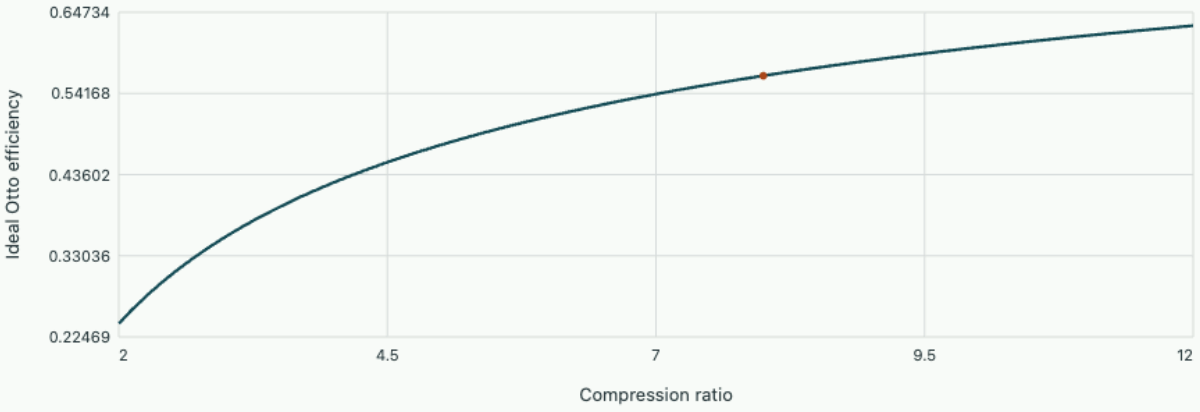

Air-standard Otto cycle Selected ratio

Efficiency 1−r^(1−γ) = 0.56472. At initial 300 K, adiabatic compression reaches 689.22 K. Constant heat capacities and reversible adiabats are assumed.

Start from this chapter. Write the first law for each insulated block and add the balances. Use the conserved weighted mean to predict the final temperature before changing the bridge conductance.

OneSimulation laboratory · Separate the equilibrium temperature from the relaxation time

Put signs on a physical boundary

Enclose a gas inside a movable cylinder. When the gas expands, it can raise an external weight. When heated at fixed volume, it can increase its temperature while doing no boundary work. To fix signs, take heat Q into the gas as positive and work W done by the gas as positive. The first law is

$$\Delta U = Q - W.$$

Internal energy is a state property; the other two quantities describe transfer during a process. A cyclic device returns to its initial state, so its net work equals its net heat input.

For a quasistatic expansion, a piston moving distance dx under external pressure p_{ext} receives work $p_{\mathrm{ext}}A\,dx = p_{\mathrm{ext}}dV$. In mechanical equilibrium p_{ext} approaches the gas pressure, giving the usual area under a pressure–volume curve. A suddenly expanding gas is not generally uniform or in mechanical equilibrium. Extending the equilibrium pressure formula to an unconstrained free expansion would predict work when no piston or other boundary force actually receives it.

Follow an ideal gas along two paths

For a calorically ideal gas, $dU = nC_{V,m}dT$ with constant molar heat capacity. Reversible isothermal expansion from V_1 to V_2 therefore gives zero internal energy change and $Q = W = nRT\ln(V_2/V_1)$. Revise the process so that the gas is thermally insulated but still expands reversibly. Then $nC_{V,m}dT = -p\,dV = -nRT\,dV/V$. Integrating and introducing $\gamma = C_{p,m}/C_{V,m} = 1 + R/C_{V,m}$ gives

$$TV^{\gamma-1} = \text{constant}.$$

Temperature now falls because the gas supplies work from its energy store.

The two paths can pass through a common starting state but they generally end at different pressures and temperatures for a given final volume. If one wants to compare work between identical endpoints one must specify a connecting step as well. This prevents a common misreading of a pressure–volume plot: an isotherm and an adiabat through the same point have different slopes and cannot serve as alternate paths to the same later point without additional processes.

Construct a reversible two-reservoir cycle

Let the gas expand isothermally at T_h, then expand adiabatically down to T_c, compress isothermally there and return adiabatically to the original state. If the first isotherm changes volume by a factor r, the two adiabatic relations force the cold compression to change volume by the reciprocal factor. The hot heat input is $Q_h = nRT_h\ln r$. The magnitude of heat expelled during cold compression is $Q_c = nRT_c\ln r$. Thus

$$W = Q_h - Q_c,$$
$$\eta = \frac{W}{Q_h} = 1 - \frac{T_c}{T_h}.$$

The working gas provides one explicit construction. More generally a reversible cycle exchanges entropy Q_h/T_h with the hot reservoir and returns exactly that entropy to the cold one. For an irreversible cyclic engine between the same two reservoirs, the reservoir entropy change $Q_c/T_c - Q_h/T_h$ is nonnegative. This gives the same expression as an upper bound on efficiency. The temperatures must be absolute and must refer to the reservoirs actually exchanging heat.

Test a proposed engine

An engine receives 1.50 kJ per cycle from a 450 K reservoir and rejects heat to a 300 K sink. Reversibility permits at most 0.500 kJ of useful work. A measured cycle delivering 0.360 kJ must reject 1.140 kJ by the

first law. The working fluid's entropy returns to its original value, while the two reservoirs together gain $1140/300 - 1500/450 = 0.467$ J/K per cycle.

The energy and entropy balances answer different design questions. Energy accounting gives the amount of rejected heat. Entropy generation explains why some of the available work has been lost relative to the reversible limit. Raising the running frequency can increase power even at lower efficiency, since power depends on work per unit time. The reversible limit describes a thermodynamic bound, and achieving a useful power output additionally depends on finite rates of heat transfer and mechanical motion.

Extract work from two finite heat reservoirs

The usual Carnot limit assumes two fixed reservoir temperatures. Two finite bodies instead change temperature as a reversible engine transfers energy between them. Consider insulated hot and cold blocks with constant heat capacities C_h, C_c. A cyclic engine operating between them can deliver work to an external mechanical store. The block temperatures vary slowly relative to internal equilibration, and the reversible engine draws and rejects each small amount of heat at its corresponding instantaneous block temperature.

If an infinitesimal amount of entropy passes through the engine reversibly, the hot block loses exactly the amount gained by the cold block. Integrating until both reach a common final temperature T_f yields

$$C_h \ln(T_f/T_{h,0}) + C_c \ln(T_f/T_{c,0}) = 0,$$

$$T_f = T_{h,0}^{C_h/(C_h+C_c)} T_{c,0}^{C_c/(C_h+C_c)}.$$

The logarithms contain temperature ratios, so the product expression is dimensionally a temperature. Once temperatures coincide a cyclic engine has no thermal difference left from which to extract work. Conservation of energy gives

$$W_{\max} = C_h(T_{h,0} - T_f) + C_c(T_{c,0} - T_f).$$

Any irreversible entropy production raises the attainable common temperature relative to the reversible result at the same initial conditions, leaving more energy in the blocks and less exported as work.

Compare a reversible engine with passive contact

Let $C_h = C_c = 1.00$ kJ/K, $T_{h,0} = 400$ K and $T_{c,0} = 300$ K. The reversible common temperature is the geometric mean $\sqrt{400 \times 300} = 346.41$ K. The engine receives 53.59 kJ from the hot block and transfers 46.41 kJ to the cold block; their difference is 7.18 kJ of work. The overall work fraction of extracted hot-side heat is 0.134. At the beginning the instantaneous reversible efficiency is $1 - 300/400 = 0.250$, falling to zero at the end. Applying that initial efficiency throughout would overestimate the work.

Now put the same blocks directly in thermal contact without an engine. Energy conservation gives the arithmetic mean of 350 K and zero work. The positive entropy produced is

$$1000 \ln[350^2/(400 \times 300)] = 20.6\,\mathrm{J/K}.$$

The final temperature exceeds the reversible result by 3.59 K for each block, precisely accounting for 7.18 kJ retained as thermal energy. In the limit of a very small initial temperature difference the geometric and

arithmetic means coincide to first order. Recoverable work thus starts at second order in the initial difference even though the heat passed between the blocks is first order, explaining why small temperature contrasts are poor resources for mechanical work.

The arithmetic and geometric means also reveal how sensitive the possible work is to initial temperature calibration. For equal heat capacities, write initial temperatures as $T_m + \delta$ and $T_m - \delta$. The reversible common value is

$$\sqrt{T_m^2 - \delta^2} \simeq T_m - \delta^2/(2T_m).$$

Multiplying the difference from the arithmetic mean by the combined heat capacity gives

$$W_{\max} \simeq C\delta^2/T_m.$$

Thus a small absolute error in a temperature difference can have a doubled relative effect on the predicted maximum work. Measuring both blocks with matched instruments matters more here than obtaining many digits for the mean temperature alone.

Study the full derivation

Work fluctuations and equilibrium free-energy differences · Statistical mechanics from microscopic counting to material response

Repeated rapid changes of a laboratory control can give a broad distribution of work even when every run follows the same prescribed protocol. Derive why a particular exponential average retrieves an equilibrium free-energy difference, why the average work is generally larger, and why a finite experiment can miss the rare trajectories that make the equality true.

Problem sheet 1 · State variables and energy conversion

Apply the local material response of Lecture 1 and the path-dependent energy balances of Lecture 2.

Correct a pressure thermometer

A sealed cell warms by 4.0 K and the pressure increases by 0.80 MPa. The vessel has expansion coefficient 4.0×10^{-5} K⁻¹ and pressure compliance 2.0×10^{-10} Pa⁻¹. The liquid has expansion coefficient 2.0×10^{-4} K⁻¹. Infer its compressibility and predict the pressure rise in a truly rigid cell.

▼ **Require the two volumes to coincide**

The observed slope is 2.0×10^{5} Pa/K. Equality of sample and vessel volume changes gives $\kappa_T = (\alpha - \alpha_v)/(dp/dT) - C_v^{\mathrm{vessel}} = 6.0 \times 10^{-10}$ Pa⁻¹. In a rigid cell the rise is

$$4\alpha/\kappa_T = 1.33\,\mathrm{MPa}.$$

Treating the original vessel as rigid would have given an apparent compressibility 1.0×10^{-9} Pa⁻¹, overstating the intrinsic response.

Compare isothermal and adiabatic expansion

One mole of monatomic ideal gas at 400 K expands reversibly to twice its initial volume. Find the final temperatures, work and heat for an isothermal path and a thermally insulated path. For which path is the final pressure lower?

▼ **Track work and internal energy**

Isothermally $T_f = 400$ K and $Q = W = RT\ln 2 = 2.31$ kJ. Adiabatically $\gamma = 5/3$ gives

$$T_f = 400 \times 2^{-2/3} = 252\,\mathrm{K}.$$

The work is the drop in internal energy, $W = (3R/2)(400 - 252) = 1.85$ kJ, with $Q = 0$. Both paths end at the same volume, so the colder adiabatic path also gives a lower final pressure. They are distinct final states, despite the same expansion ratio.

Diagnose a refrigerator claim

A cyclic refrigerator takes 600 J from a reservoir at 270 K and gives heat to a room at 300 K. Find the minimum work. A prototype claims to use 50 J per cycle. Test the claim by computing total reservoir entropy change.

▼ **Apply both laws to one full cycle**

At reversible operation the equality $Q_h/300 = 600/270$ gives $Q_h = 666.7$ J and minimum work $Q_h - Q_c = 66.7$ J. The claimed input gives $Q_h = 650$ J by energy conservation. Reservoir entropy would

change by $650/300 - 600/270 = -0.0556$ J/K per cycle, incompatible with the two-reservoir cyclic description. A hidden heat source or an incorrect energy measurement would need to be included.

Test a calorimeter conversion near zero expansion

At 280 K a liquid has $V_m = 2.00 \times 10^{-5}$ m³/mol, $\kappa_T = 5.00 \times 10^{-10}$ Pa⁻¹ and $\alpha = -2.00 \times 10^{-4}$ K⁻¹. Estimate $C_{p,m} - C_{V,m}$ and infer $C_{V,m}$ if $C_{p,m} = 75.0$ J/(mol K). What happens at a temperature where α crosses zero?

▼ **Track the derivative and its sign**

The result is $280(2.00 \times 10^{-5})(2.00 \times 10^{-4})^2/(5.00 \times 10^{-10}) = 0.448$ J/(mol K). Thus $C_{V,m} = 74.6$ J/(mol K). The negative expansion sign disappears because it occurs in both factors of the chain rule. At the zero crossing the local difference vanishes for finite positive compressibility. The capacities themselves need not vanish.

Revisit chapter 1 · Convert an isobaric calorimeter reading into an isochoric one.

Explore these relations

- For the pressure response, compare the distinction between isothermal and adiabatic compressibility; retain the separate vessel-compliance correction in this sheet. Open the live calculation.
- For heat and work, inspect the energy balance on an ideal cycle. Its prescribed paths differ from the individual isothermal and adiabatic expansions in the question. Open the live calculation.

Chapter 3 · Thermodynamic potentials and phase boundaries

Select the appropriate potential for a controlled experiment and infer the slope of a coexistence line from latent heat and volume change.

FROM FIRST PRINCIPLES

Derive a coexistence slope from equality of chemical potentials

At fixed temperature and pressure a closed single-component system lowers its Gibbs energy when matter moves into a phase with lower molar chemical potential. Transfer an infinitesimal amount dn from phase 1 to phase 2. Along the coexistence line, either direction is possible in equilibrium, so

$$dG = (\mu_2 - \mu_1)\, dn = 0,$$
$$\mu_1(T, p) = \mu_2(T, p).$$

Moving along the line must preserve that equality. The Gibbs differential for one mole gives

$$d\mu = -s_m\, dT + v_m\, dp,$$

with molar entropy s_m and volume v_m. Equating the differentials of the two phases leaves

$$(v_{m,2} - v_{m,1})\, dp = (s_{m,2} - s_{m,1})\, dT.$$

The latent heat per mole for reversible conversion $1 \to 2$ is $L_m = T(s_{m,2} - s_{m,1})$. Therefore the slope is $dp/dT = L_m/[T(v_{m,2} - v_{m,1})]$ for a nonzero volume jump. Its sign depends on both latent heat and volume change. This relation requires measured or modeled properties of the two phases at the same coexistence point.

Background · Helmholtz free energy

$$F = -k_B T \ln Z$$

At fixed particle number and temperature, differentiating with respect to volume yields minus the pressure. The ensemble specifies which variables to hold fixed.

DERIVATION AND APPLICATIONS

Calculate how a curved interface changes the equilibrium vapour pressure

A small liquid droplet in its own vapour reaches equilibrium at a different vapour pressure from a flat liquid surface. Consider a spherical droplet of radius r at uniform temperature T, with surface tension γ independent of size. Let p_v and p_l be the exterior vapour and interior liquid pressures. A virtual change of

radius changes area and volume by $dA = 8\pi r\,dr$ and $dV = 4\pi r^2 dr$. Balancing the pressure work against the surface free energy gives

$$(p_l - p_v)dV = \gamma\,dA,$$
$$p_l = p_v + \frac{2\gamma}{r}.$$

Use a flat interface at the same temperature as a chemical-potential reference. Both phases there have pressure p_∞. The difference of their chemical potentials is zero for the flat reference and for the equilibrium droplet. Integrating $d\mu = v_m\,dp$ along an isotherm for each phase then gives, for ideal vapour and nearly incompressible liquid of molar volume v_l,

$$RT\ln(p_v/p_\infty) = v_l(p_l - p_\infty) = v_l(p_v - p_\infty) + \frac{2\gamma v_l}{r}.$$

For a dilute vapour the molar volume greatly exceeds the liquid value. When $p_v - p_\infty$ is also small compared with the curvature-induced pressure jump, its contribution on the right is negligible. The Kelvin relation follows,

$$\ln(p_v/p_\infty) \simeq \frac{2\gamma v_l}{rRT}.$$

A convex liquid surface raises the required equilibrium vapour pressure. The result tends continuously to flat coexistence at large radius. At molecular radii the continuum description, including the assumption of a bulk surface tension, requires independent justification.

WORKED EXAMPLE

Estimate the supersaturation over a droplet

Change the inputs and investigate this derivation.

For water-like material at $298\,\mathrm{K}$ take $\gamma = 0.0720\,\mathrm{N/m}$, $v_l = 18.0\times10^{-6}\,\mathrm{m^3/mol}$, and $r = 20.0\,\mathrm{nm}$. The exponent is 0.0523, giving $p_v/p_\infty = 1.054$. The mechanical pressure jump is $7.20\,\mathrm{MPa}$. Even with $p_\infty = 3.17\,\mathrm{kPa}$, the exterior pressure increase is only about $170\,\mathrm{Pa}$, verifying the term dropped in obtaining the Kelvin relation.

SECOND APPLICATION

Infer the radius that can coexist with a given vapour

Change the inputs and investigate this derivation.

A different liquid has $\gamma = 0.0250\,\mathrm{N/m}$ and $v_l = 100\times10^{-6}\,\mathrm{m^3/mol}$ at $300\,\mathrm{K}$. A chamber stabilized at $p_v/p_\infty = 1.020$ matches radius $r = 2\gamma v_l/[RT\ln(1.020)] = 101\,\mathrm{nm}$. A smaller droplet has a higher required vapour pressure and tends to evaporate in that chamber. A larger one tends to grow, so matching the pressure selects an unstable balance if the vapour pressure is held fixed.

Try it · Kelvin droplet

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

Surface tension (N/m) **0.072**

Temperature (K) **300**

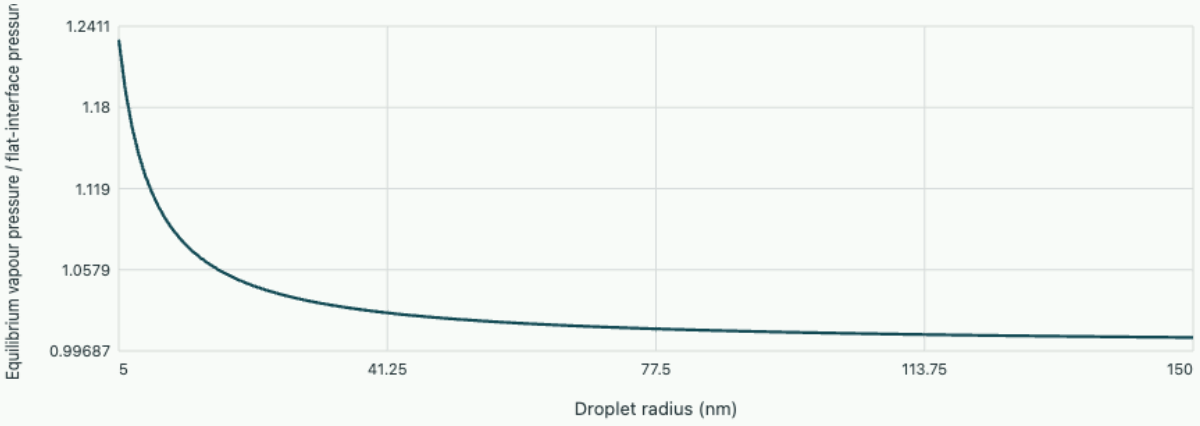

Kelvin curvature shift

Liquid molar volume fixed at 18 cm³/mol. Curvature raises the liquid chemical potential; molecular-scale droplets require corrections beyond this continuum model.

Change which variables are controlled

Consider a closed simple compressible system that can exchange heat with a reservoir and work by moving a boundary. In equilibrium its fundamental energy differential for fixed amount is $dU = T\,dS - p\,dV$. This expression makes entropy and volume natural independent variables, yet laboratory experiments more frequently control temperature and pressure. A change of thermodynamic potential exchanges one independent variable for its conjugate without losing any equilibrium information.

For the Helmholtz free energy $F = U - TS$, differentiation gives

$$dF = -S\,dT - p\,dV.$$

Its slope with respect to volume at fixed temperature is minus the pressure. For the Gibbs free energy $G = U + pV - TS$, the differential is $dG = -S\,dT + V\,dp$. At fixed temperature and pressure a spontaneous change in a closed sample with only pressure–volume work tends to lower its total G. The

external reservoir can exchange heat during that change, so a fall in sample energy by itself would not be the correct selection criterion.

Extract a measurable cross derivative

The entropy of a gas is harder to measure directly than its thermal expansion. Since $G(T, p)$ is a state function with commuting mixed derivatives inside a smooth single phase,

$$-(\partial S/\partial p)_T = (\partial V/\partial T)_p.$$

This is one Maxwell relation. It converts an entropy response to a derivative accessible from volume measurements. The equality applies to the same material and the same choice of fixed variables on both sides.

For one mole of ideal gas $(\partial V/\partial T)_p = R/p$. Compressing isothermally therefore changes its entropy by $\Delta S = -R\ln(p_2/p_1)$. Doubling pressure gives -5.76 J/(mol·K), in agreement with an independent reversible-heat calculation. The Maxwell relation obtains the state difference without requiring the actual compression process to be reversible. Only an equilibrium path used to evaluate it has to connect well-defined states.

Derive a coexistence-line slope

For one pure substance the molar Gibbs energy is the chemical potential. At two-phase equilibrium both phases must have the same chemical potential; otherwise transferring matter to the phase of lower value decreases total G. Label the phases 1 and 2. Along their coexistence curve a small change in T and p changes their chemical potentials equally. Substituting $d\mu = -s_m dT + v_m dp$ for each yields

$$\frac{dp}{dT} = \frac{s_{m,2} - s_{m,1}}{v_{m,2} - v_{m,1}} = \frac{L_m}{T(v_{m,2} - v_{m,1})}.$$

Here L_m is the molar latent heat for the transition from phase 1 to 2. The slope can be negative for melting when the liquid has smaller molar volume than the solid. For vaporisation the volume difference is usually positive and large at modest pressure. A pressure reading and a thermometer together can therefore constrain a latent heat through measurements of the slope and molar volumes.

Background · How can pressure favour melting?

At coexistence the molar Gibbs free energies of solid and liquid agree. Increasing pressure at fixed temperature increases each free energy by its molar volume times the pressure increment. The phase with the smaller molar volume gains the thermodynamic advantage. For ordinary ice near its melting point, melting reduces the molar volume by about 1.6 cm³/mol and requires about 6.0 kJ/mol of heat.

$$\frac{dT_m}{dp} = \frac{T_m(v_l - v_s)}{\Delta h_{\mathrm{fus}}} \simeq -7.3 \times 10^{-8}\ \mathrm{K/Pa}.$$

A 1 MPa increase lowers the equilibrium melting temperature by about 0.073 K near 273 K. This is a local prediction along the coexistence curve. Extrapolating it to hundreds of megapascals would require the pressure dependence of the volumes and latent heat and the possible appearance of another solid phase.

Estimate a boiling temperature near atmospheric pressure

For vaporising water near 373 K use $L_m = 40.7$ kJ/mol, a vapour volume of 0.0306 m³/mol and a much smaller liquid volume. The coexistence slope is approximately $40700/[373(0.0306)] = 3.57$ kPa/K. An increase of pressure by 10 kPa therefore raises the boiling temperature by roughly 2.8 K over this limited interval. The dimensional check is J/(m³·K), equivalent to Pa/K.

Over a larger interval, approximate the vapour as ideal, neglect the liquid volume and treat L_m as constant. The slope becomes $d\ln p/dT = L_m/(RT^2)$, giving $\ln(p_2/p_1) = -(L_m/R)(1/T_2 - 1/T_1)$. This form clarifies why a straight line may appear on a logarithmic-pressure versus reciprocal-temperature plot. Curvature in measured data can carry information about temperature-dependent latent heat and departures from ideal-gas behaviour. Close to a critical point the assumptions behind the simple integrated expression cease to be reliable.

For a substance whose solid is less dense than its liquid, melting can have a negative coexistence slope even though the latent heat of melting is positive. In the Clapeyron relation the sign then comes from the volume difference. Increasing pressure at fixed temperature favours the phase taking less volume. This sign check connects the derivative calculation to an independently measurable density contrast.

Choose the equilibrium pressure inside a two-phase region

A nonmonotone model isotherm can predict several volumes at one pressure below the critical temperature. The middle segment with $(\partial p/\partial v)_T > 0$ is mechanically unstable. Removing that segment alone does not identify the two coexisting phases. At the coexistence pressure p_* both endpoints have equal molar chemical potential. Write molar Gibbs energy as $\mu = f + p_* v$, where f is molar Helmholtz energy on this isotherm. Since $df = -p(v)\,dv$ at fixed temperature, equality of endpoint values implies

$$0 = \mu_g - \mu_l = -\int_{v_l}^{v_g} p(v)\,dv + p_*(v_g - v_l).$$

Geometrically, the areas above and below the horizontal line p_* cancel. The same condition makes that line's negative slope the common tangent of $f(v)$ at v_l and v_g. Between these endpoint volumes the equilibrium material changes the relative amounts of liquid and gas at fixed pressure.

As a local model take $p = p_0 + A(x - x^3)$ with $x = (v - v_0)/b$, $A > 0$, $v_0 = 2.00 \times 10^{-4}$ and $b = 3.00 \times 10^{-5}$ m³/mol. At $p = p_0$, the outer roots $x = \pm 1$ have negative slopes, and the excess pressure is odd in x, verifying equal areas exactly. The stable molar volumes are 1.70×10^{-4} and 2.30×10^{-4} m³/mol. At average volume 2.12×10^{-4} m³/mol, conservation of volume gives a gas mole fraction $(2.12 - 1.70)/(2.30 - 1.70) = 0.700$. The spinodals, where the original slope vanishes, instead occur at $x = \pm 1/\sqrt{3}$. They delimit local stability. They are not the coexistence volumes. This cubic is a local representation of the loop and should not be extrapolated to very large or negative volumes.

Practise on sheet 2 · Correct a trial coexistence pressure from an area residual. Further reading · University of Cambridge · David Tong, Statistical Physics · §5.1.1 Phase Equilibrium (Maxwell construction).

Why a phase can persist beyond the coexistence line

Equality of bulk chemical potentials determines coexistence, yet a sample can remain liquid below its equilibrium freezing temperature. Forming a small solid nucleus introduces an interface whose free-energy cost can initially outweigh the bulk driving force. Consider a spherical nucleus of radius r in a pure, uniformly undercooled liquid. Assume an isotropic solid–liquid interfacial energy $\gamma > 0$ and a positive bulk free-energy advantage per volume Δg_v for the solid. Ignore elastic accommodation, walls, impurities and curvature dependence of γ.

The gain in stable bulk material is proportional to the volume while the cost is proportional to area. Relative to uniform liquid the work is

$$\Delta G(r) = 4\pi r^2 \gamma - \frac{4\pi}{3} r^3 \Delta g_v.$$

Differentiating gives the nonzero stationary size $r_* = 2\gamma/\Delta g_v$. The second derivative there is negative, so the critical nucleus lies at a barrier rather than at a stable droplet size. Smaller nuclei reduce their free energy by shrinking, and larger ones reduce it by growing. Substituting the radius into the original expression gives

$$\Delta G_* = 16\pi\gamma^3/(3\Delta g_v^2).$$

Near coexistence, the difference of the two bulk free energies changes as the entropy-of-fusion difference times undercooling. With latent heat per unit volume L_v, use $\Delta g_v \simeq L_v(T_m - T)/T_m$ for a small fractional undercooling.

Determine how an undercooling changes the barrier

For an illustrative melt take $T_m = 400$ K, $T = 380$ K, $L_v = 3.00 \times 10^8$ J/m³ and $\gamma = 0.0300$ J/m². The bulk drive is 1.50×10^7 J/m³. Critical radius is 4.00 nm and the barrier is approximately 2.01×10^{-18} J. Since $k_B T = 5.25 \times 10^{-21}$ J at 380 K, this is roughly 383 thermal energy units, a formidable homogeneous

nucleation barrier. This comparison alone is insufficient for an exact waiting time, which also requires a kinetic prefactor and the sample volume.

Doubling undercooling to 40 K doubles the approximate bulk drive, halves the critical radius to 2.00 nm and reduces the energy barrier by four. The nucleation probability per attempt can thus vary by many orders of magnitude while the equilibrium transition temperature remains fixed. As undercooling tends to zero, both critical radius and barrier diverge in this ideal capillarity model. A container wall or an impurity can reduce the interface area needed and initiate crystallization much earlier. This distinction between the coexistence curve and observed onset explains why reducing the cooling rate or changing surface preparation can change a freezing record.

The spherical calculation exposes a distinction between a critical nucleus and a nucleus for which the total formation work has already become negative. Set the formation work itself, rather than its derivative, to zero. Apart from zero radius this gives $r = 3\gamma/\Delta g_v$, larger than the critical radius by a factor of one and a half. A cluster just above the critical radius can still have positive total formation work relative to the uniform liquid, yet growth reduces its free energy. Using the zero-work radius as a kinetic threshold would therefore overestimate the size needed for downhill growth.

Estimate the boiling temperature under reduced pressure

Anchor a liquid’s saturation curve at 373.15 K and 101.325 kPa. For a dilute ideal vapour, negligible liquid volume and approximately constant molar heat of vaporisation, integrate $d\ln p/dT = L/(RT^2)$. The three curves test sensitivity to the assumed latent heat.

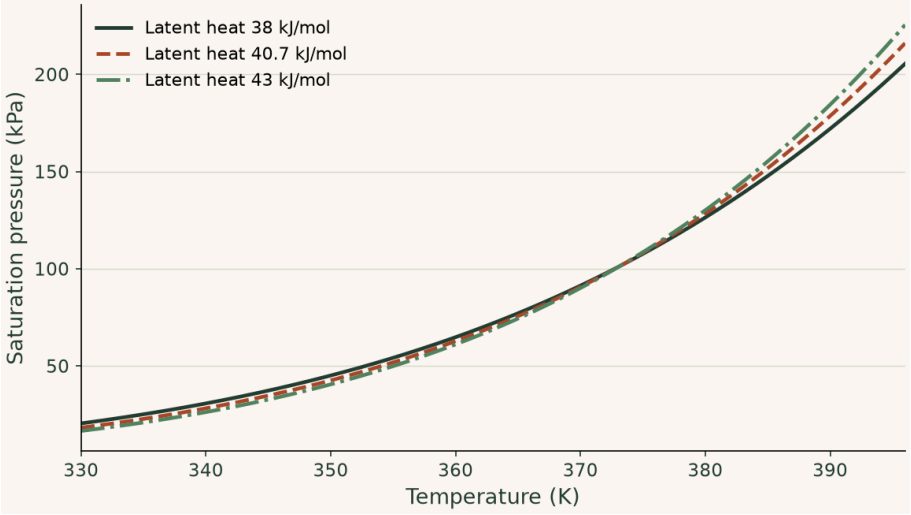

For 40.7 kJ/mol the predicted pressure at 350 K is about 42.5 kPa. Uncertainty in latent heat matters increasingly far from the anchor point. This local model needs new property data before extrapolation toward a critical point.

Try it · Clausius

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

Molar latent heat (kJ/mol) **40.5**

Anchor boiling temperature (K) **374**

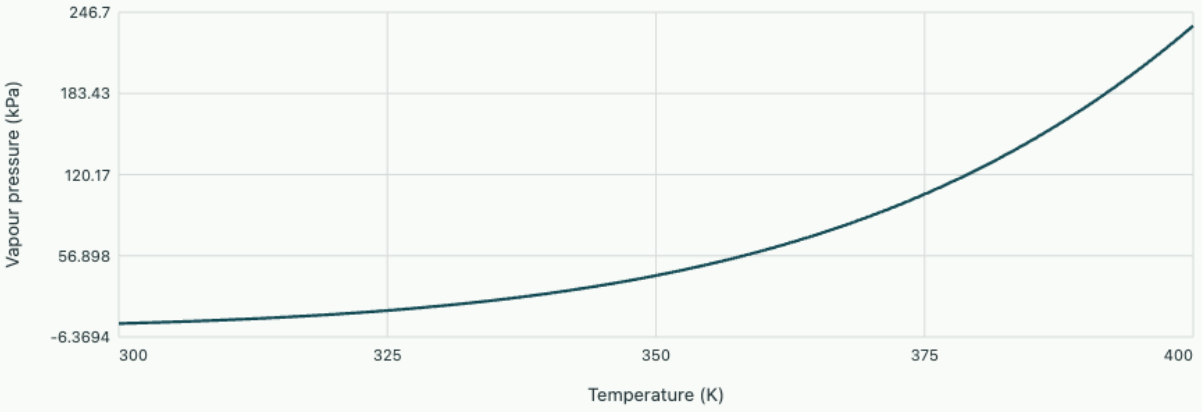

Constant-latent-heat approximation

Pressure is anchored to 101.325 kPa at 374 K. Ideal vapour, negligible liquid molar volume and constant latent heat are assumed.

▶ **Read values from the plot**

Study the full derivation

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.

Chapter 4 · Counting states and defining temperature

Use a model of discrete energy sharing to derive thermodynamic temperature and distinguish a most probable configuration from one microscopic state.

FROM FIRST PRINCIPLES

Derive thermal equilibrium by counting possible energy divisions

For an isolated system with a fixed energy range, assume all accessible microscopic states in that range have equal probability. Its multiplicity $\Omega(E)$ counts those states. Entropy is

$$S = k_B \ln \Omega;$$

the logarithm makes entropy additive when independent state counts multiply. Put systems A and B into weak thermal contact so they share a fixed total energy without exchanging particles or doing work. The number of states compatible with an energy E_A is

$$\Omega_{\text{total}}(E_A) = \Omega_A(E_A)\Omega_B(E_{\text{total}} - E_A).$$

The most probable energy division maximizes this product. For large systems with smooth entropy curves, take its logarithm and differentiate. A gain of A is a loss of B, so at an interior maximum

$$0 = \frac{d(S_A + S_B)}{dE_A} = \left(\frac{\partial S_A}{\partial E_A}\right)_{V,N} - \left(\frac{\partial S_B}{\partial E_B}\right)_{V,N}.$$

The common quantity is the inverse thermodynamic temperature,

$$1/T = (\partial S/\partial E)_{V,N}.$$

This definition relates a counting problem to the observed absence of net energy exchange at equilibrium. Positive heat capacities give a locally stable maximum and set the width of the energy-sharing fluctuations.

Background · Multiplicity and entropy

$$S = k_B \ln \Omega$$

Count microstates compatible with the macroscopic constraints. Distinguishable particles, indistinguishable occupations and a changed energy window need distinct counts.

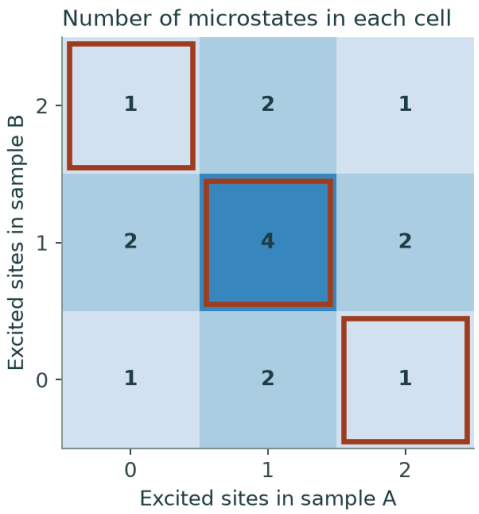

A counting diagram for a finite reservoir

Prepare the four equivalent sites independently with the same excitation probability strictly between zero and one. Keep only records with two excitations. Every such microscopic pattern then has the same probability. Only the outlined diagonal survives this constraint. The middle cell represents two choices in A times two in B. The probabilities for A to contain zero, one or two excitations are therefore 1/6, 4/6 and 1/6. Treating the three macroscopic outcomes as equally likely would give a different experiment.

Explore the related parameter study beside the computed figure.

DERIVATION AND APPLICATIONS

Find the width of the most probable energy division

An isolated pair of subsystems can have sharply defined total energy while the energy in each part fluctuates. Let $E_A + E_B = E_{\mathrm{tot}}$, and assume weak coupling, rapid internal equilibration and equal statistical weights of accessible microstates. For two macroscopic subsystems with smooth entropies, the probability density of a particular division of energy is proportional to its state count,

$$P(E_A) \propto \exp\left(\frac{S_A(E_A) + S_B(E_{\mathrm{tot}} - E_A)}{k_B} \right).$$

At its interior maximum $d(S_A + S_B)/dE_A = 1/T_A - 1/T_B = 0$. The temperatures there equal T. To examine nearby divisions, differentiate again. Each subsystem at fixed volume and particle number has $d(1/T)/dE = -1/(T^2 C)$, with heat capacity C. Squaring the chain derivative of E_B makes the two curvature contributions add,

$$\frac{d^2(S_A + S_B)}{dE_A^2} = -\frac{1}{T^2}\left(\frac{1}{C_A} + \frac{1}{C_B} \right).$$

For positive heat capacities the maximum is stable. Write $\delta E = E_A - E_A^*$ relative to the most probable energy. A quadratic expansion of the entropy about the maximum yields a Gaussian with variance

$$P(\delta E) \propto e^{-\delta E^2/(2\sigma_E^2)},$$

$$\sigma_E^2 = k_B T^2 \frac{C_A C_B}{C_A + C_B}.$$

The finite partner reduces the width because it must undergo exactly the opposite energy fluctuation. For $C_B \gg C_A$ the canonical variance $k_B T^2 C_A$ is recovered. A temperature inferred from the subsystem energy has fluctuation scale σ_E / C_A when its heat capacity is locally constant. A macroscopic state's relative fluctuations shrink because heat capacity is extensive.

WORKED EXAMPLE

Estimate fluctuations of a small calorimeter

Change the inputs and investigate this derivation.

At $T = 300\,\mathrm{K}$, let $C_A = 1.00 \times 10^{-12}\,\mathrm{J/K}$ and $C_B = 3.00 \times 10^{-12}\,\mathrm{J/K}$. The effective capacity in the variance is $0.750 \times 10^{-12}\,\mathrm{J/K}$. With $k_B = 1.381 \times 10^{-23}\,\mathrm{J/K}$, the root-mean-square energy fluctuation is $9.65 \times 10^{-16}\,\mathrm{J}$, equivalent to $0.965\,\mathrm{mK}$ of inferred temperature in A.

SECOND APPLICATION

Determine whether the partner acts as a reservoir

Change the inputs and investigate this derivation.

Another apparatus measures an energy standard deviation for A of 0.800 times the value predicted for an infinite bath. The variance ratio is

$$0.640 = C_B/(C_A + C_B).$$

Thus

$$C_B/C_A = 0.640/0.360 = 1.78.$$

For $C_A = 2.00\,\mathrm{nJ/K}$ this implies $C_B = 3.56\,\mathrm{nJ/K}$. Using the standard-deviation ratio directly in the variance formula would substantially overestimate the partner capacity.

Try it · Finite reservoir variance

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

Partner / sample heat capacity 3

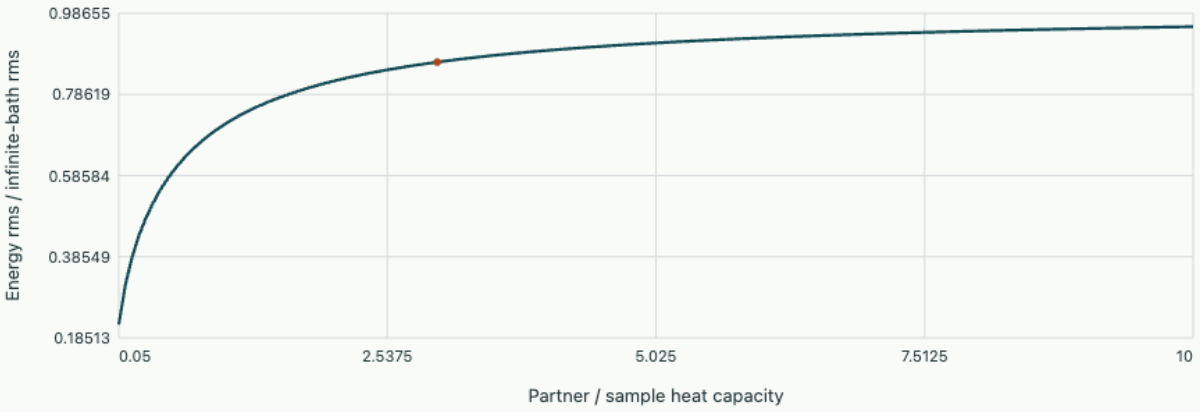

Two finite thermal bodies Selected capacity ratio

Variance ratio CB/(CA+CB) = 0.75. The canonical infinite-bath result is approached only when CB≫CA.

Start from this chapter. A small transfer δQ changes the total entropy by δQ(1/Tcold – 1/Thot). Follow that sign during the computed relaxation and find the temperature at which the derivative vanishes.

OneSimulation laboratory · Separate the equilibrium temperature from the relaxation time

State exactly what is counted

Consider N independent distinguishable vibrational modes, all of angular frequency ω, with a common energy quantum $\varepsilon = \hbar\omega$. Suppose they contain q excitation quanta in total. A microstate is an ordered list of occupation numbers $(q_1, \ldots, q_N)$ with nonnegative integers adding to q. Giving two quanta to mode 1 differs from giving the same two quanta to mode 2, while exchanging two identical quanta within one mode creates no new microstate. An additive ground-state energy plays no role in the counting below.

Represent a distribution by writing q dots and separating the modes with $N-1$ dividers. The locations of the dividers among all $q+N-1$ symbols determine the occupations, including a zero when two dividers coincide in succession. The multiplicity is therefore

$$\Omega(N, q) = \binom{q + N - 1}{q}.$$

It counts the microscopic realisations of a specified macroscopic energy, not different possible values of that energy.

Allow subsystems to exchange energy

Put two such solids into weak thermal contact but isolate the pair from the environment. For a fixed division q_A and $q_B = q - q_A$, weak coupling makes the joint multiplicity the product $\Omega_A(q_A)\Omega_B(q_B)$. Under the microcanonical hypothesis of equal probabilities for accessible microstates of the pair, this product sets the probability weight for an observed division. The most probable division maximises the product, though for finite systems other divisions retain nonzero probability.

As a concrete small system take $N_A = N_B = 2$ and $q = 4$. Each solid has multiplicity $q_i + 1$. The weights for $q_A = 0, 1, 2, 3, 4$ are $5, 8, 9, 8, 5$ with sum 35. Equal sharing is the most probable single division, yet it occurs only $9/35 = 0.257$ of the time. Any one specific occupation list still has probability $1/35$. The distinction between a microstate and a macrostate resolves the apparent tension between equal probabilities and a preferred energy division.

Obtain temperature from an entropy slope

Introduce $S = k_B \ln \Omega$ so that independent multiplicities produce additive entropy. For macroscopic N and q, use Stirling's approximation $\ln n! \simeq n \ln n - n$ and drop corrections small relative to N. Writing $U = q\varepsilon$ and holding N fixed gives

$$S/k_B \simeq (q + N)\ln(q + N) - q \ln q - N \ln N,$$

$$\frac{1}{T} = \frac{k_B}{\varepsilon} \ln(1 + N/q).$$

This derivative defines an equilibrium temperature from how quickly the number of ways to store energy grows. The entropy maximum of two systems in contact then requires equal inverse temperatures, even if the systems have different numbers of modes. Heat flows towards the subsystem gaining more entropy per added energy. The equilibrium criterion concerns equality of slopes; the two subsystems need not contain the same total energy.

Solving the derivative for energy per mode yields

$$U/N = \varepsilon / [\exp(\varepsilon / k_B T) - 1].$$

At high temperature the exponential in the denominator expands to first order, giving energy near $k_B T$ for one harmonic mode. At low temperature only a small fraction of modes carry excitations. The energy and heat capacity above their ground values are then strongly suppressed.

Connect the count with a thermometer

For a set of modes with $\varepsilon/k_B = 240$ K, an observed mean occupation $q/N = 0.50$ gives

$$T = 240/\ln 3 = 218\,\mathrm{K}.$$

Packing the same excitations into twice as many identical modes changes q/N to 0.25 and gives

$$T = 240/\ln 5 = 149\,\mathrm{K}.$$

Equal quantities of excitation energy therefore need not indicate equal temperatures. The reservoir of available modes matters.

These estimates assume the modes have a common frequency and have explored accessible distributions sufficiently for the statistical description to apply. A newly driven single mode could contain nearly all the energy while a thermometer responds differently from the occupation-based estimate. As the sample size increases the spread of typical energy per mode narrows, but finite-time preparation and weak internal relaxation still determine whether equilibrium counting is an appropriate description of a measurement.

The entropy curvature also determines whether sharing energy is stable. Near the maximum, expand the combined entropy to second order in a small transfer between subsystems. A negative summed curvature makes the probability fall on either side of the most likely energy. For ordinary positive heat capacities, each entropy curvature is negative, so the two bodies return statistically towards equal temperature after a small redistribution.

A bounded energy spectrum reverses the entropy slope

An assembly of independent localized two-state units supplies a finite system in which the sign of temperature can be read directly from multiplicity. Take N units with low energy zero and high energy ε. If n are excited, total energy is $n\varepsilon$ and the multiplicity of that sharp-energy macrostate is $\binom{N}{n}$. It grows up to half occupancy and decreases when nearly all units are forced into the high state. This assumes the ensemble is internally equilibrated over the available states near its specified energy.

Use the entropy of a narrow energy shell $S = k_B \ln \Omega$. Adding one excitation changes its entropy by $k_B \ln[(N-n)/(n+1)]$. Dividing by the associated energy increment defines a finite-difference inverse temperature. For large N with neither level nearly empty, replacing $n+1$ by n gives

$$\frac{1}{T} \simeq \frac{k_B}{\varepsilon} \ln \frac{N-n}{n},$$

$$\frac{n}{N} = \frac{1}{1 + e^{\varepsilon/(k_B T)}}.$$

Ordinary positive temperatures have more units in the lower state. Approaching equal populations from below sends inverse temperature toward zero. Above half occupancy the entropy falls as energy increases, giving a negative shell temperature. Such an internally equilibrated population inversion requires a bounded spectrum and isolation from unbounded energy reservoirs on the observation timescale. It cannot describe an unrestricted equilibrium ideal gas with arbitrarily high translational energies.

Predict the direction of exchange

Consider $N = 10,000$ units with $\varepsilon/k_B = 120$ K. A preparation with 2500 excited units has $T \simeq 120/\ln 3 = 109$ K. A preparation with 7500 excited units has $T \simeq 120/\ln(1/3) = -109$ K. Join the inverted system through weak energy exchange to a large positive-temperature reservoir at 300 K. If the system loses positive energy δE to the reservoir, the combined entropy change is $\delta E(1/300 - 1/(-109)) > 0$. Net energy therefore leaves the inverted sample even though comparing the signed temperature numbers as ordinary magnitudes would suggest otherwise.

For $\delta E = 1.00 \times 10^{-20}$ J, the first-order change is about 1.25×10^{-22} J/K, or nine Boltzmann constants. This uses an exchange small relative to the original ensemble energy so that the temperature slopes stay nearly constant. At exactly half occupancy the inverse temperature is zero in the large-system limit. Pumping a system into an inverted population and allowing it to equilibrate internally are separate experimental steps. A measured majority of excited units alone is evidence for inversion; additional checks on internal energy exchange are needed before assigning a thermodynamic temperature to that preparation.

The preparation also shows why inverse temperature is often more convenient than temperature itself. Moving the upper-state fraction continuously from slightly less than one half to slightly more than one half sends temperature through positive and negative infinity, while inverse temperature crosses zero smoothly. A weak probe with the same small level spacing can test the slope through the relative probabilities of up and down transitions when equilibrated with the sample. A probe with unrestricted kinetic energy, in contrast, can drain the inversion before it becomes a useful thermometer. The compatibility of thermometer and sample therefore matters near a bounded-spectrum population inversion.

Why is a shared-energy macrostate overwhelmingly likely?

Twenty distinguishable sites each have two possible energies. For independent excitation probability p, the macrostate probability includes both its multiplicity $\binom{20}{n}$ and the probability of one such arrangement. Compare changing the energetic bias with changing the counting rule.

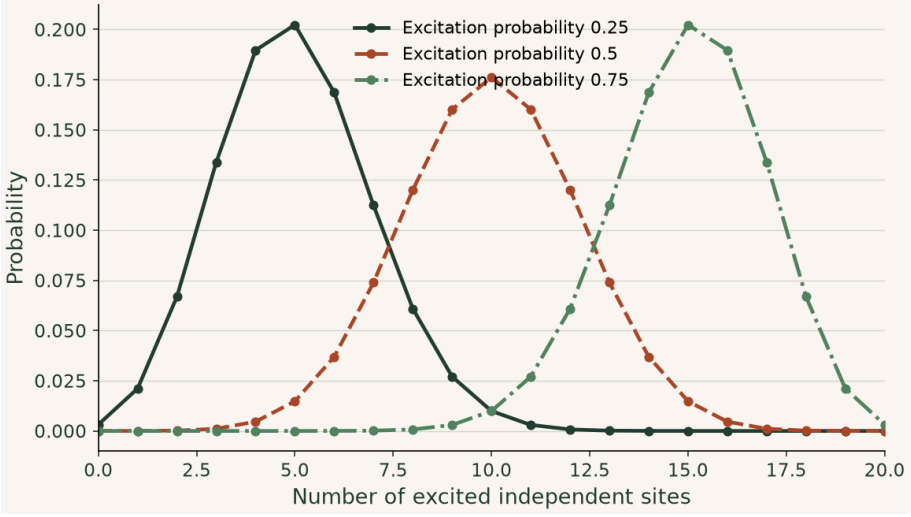

At p = 1/2 the central probability is about 0.176, whereas the all-excited state has probability less than one in a million. Each individual configuration remains equally probable at this p; it is the number of configurations that differs.

Try it · Physics lab counting limit

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

Independent sites N **20**

Excitation probability p **0.2**

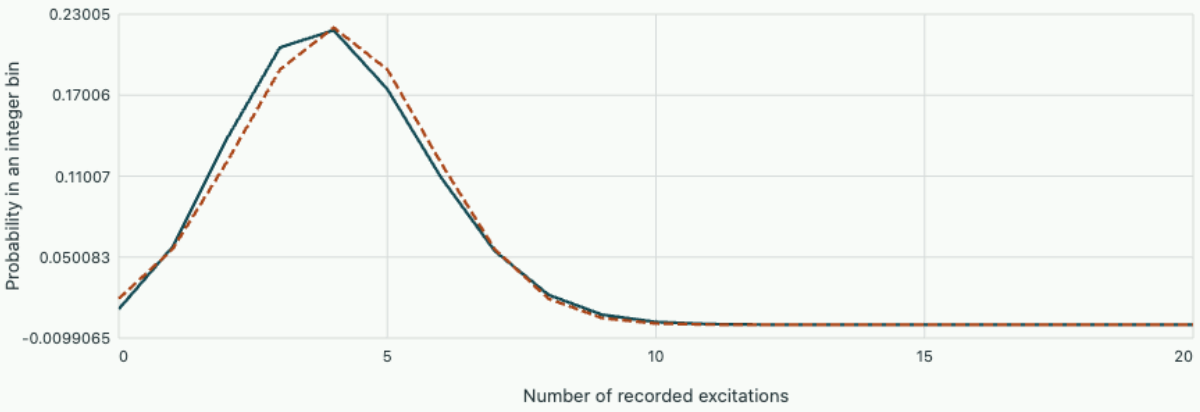

Mean 4, standard deviation 1.7889. Largest absolute bin error 0.017167 at count 2. Skewness 0.33541.

▶ **Read values from the plot**

Study the full derivation

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.

Problem sheet 2 · Potentials, phase lines and counting

Relate the equilibrium derivatives of Lecture 3 to the microscopic state counts in Lecture 4.

Infer a latent heat

At 350 K a pure liquid's vapour pressure is 80 kPa. Near this point the measured slope is 3.0 kPa/K. Estimate the latent heat under the ideal-vapour approximation and predict the local temperature increase for a 4.0 kPa pressure rise. Identify the neglected material quantity.

▼ **Rearrange the coexistence equation**

Neglecting liquid molar volume gives $v_g \simeq RT/p = 0.0364$ m³/mol. Then $L_m = T(v_g - v_l)\,dp/dT \simeq 350(0.0364)(3000) = 38.2$ kJ/mol. Locally $\Delta T \simeq 4.0/3.0 = 1.33$ K. A measured nonideal vapour molar volume or appreciable liquid volume would change the inferred heat; both appear in the exact Clapeyron slope.

Count excitations without assigning names to them

Three distinguishable modes share three identical quanta. Find the multiplicity, the probability that the first mode is empty and its mean occupation. Then explain why putting a label on every quantum would answer a different counting problem.

▼ **Enumerate by the first mode's occupation**

There are $\binom{5}{3} = 10$ states. For first occupancies 0, 1, 2 and 3, the numbers of ways to distribute the remainder over the other two modes are 4, 3, 2 and 1. The empty probability is $4/10$ and mean occupation is

$$(0 + 3 + 4 + 3)/10 = 1.$$

Symmetry also requires a mean of one per mode. Distinguishable quanta would attach extra multiplicities to assignments within these same occupation lists.

Read entropy from volume measurements

One mole of liquid at 300 K occupies 25 cm³ and has thermal expansion coefficient 6.0×10^{-4} K⁻¹. Approximate the isothermal entropy change when pressure is increased by 2.0 MPa. Predict what the sign would do for a material with negative thermal expansion.

▼ **Use the Maxwell relation at fixed temperature**

The derivative $(\partial V/\partial T)_p = \alpha V = 1.50 \times 10^{-8}$ m³/K for this sample. Hence $\Delta S \simeq -\alpha V \Delta p = -0.030$ J/K. The estimate treats the measured expansion response as nearly constant across the interval. For negative expansion the same pressure increase gives a positive entropy change. The sign follows the measured derivative and need not be inherited from an ideal-gas analogy.

Correct a trial coexistence pressure from an area residual

At a trial pressure the two outer roots of an isotherm differ by 6.00×10^{-5} m³/mol. Numerical integration gives $\int_{v_l}^{v_g} (p(v) - p_*)\, dv = +0.120$ J/mol. Derive the first-order pressure correction including the motion of the roots, and interpret its sign.

▼ **Differentiate the area with variable endpoints**

Both endpoint integrands vanish because the roots lie on the trial line. Only the derivative of $-p_*$ remains, giving $dA/dp_* = -(v_g - v_l)$. A Newton correction is

$$\delta p = -A / A' = +2.00 \times 10^{3}\,\mathrm{Pa}.$$

Raising the line reduces the positive area. New intersections must be located before re-evaluating the residual. The correction becomes poorly conditioned as the two coexistence volumes approach each other.

Revisit chapter 3 · Choose the equilibrium pressure inside a two-phase region.

Explore these relations

- For latent heat, vary the molar latent heat and compare the vapour-pressure slope. The plot uses a 101.325 kPa anchor, so transfer the logarithmic slope rather than that anchor pressure. Open the live calculation.
- For multiplicity, compare exact finite counts with the large-system approximation. Open the live calculation.

Chapter 5 · Partition functions and thermal fluctuations

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

$$Z(f) = \int e^{-\beta(H_0 - f x)} d\Gamma,$$

$$\frac{\partial \ln Z}{\partial f} = \beta \langle x \rangle_f.$$

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 $H_0 = k x^2/2$, 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**

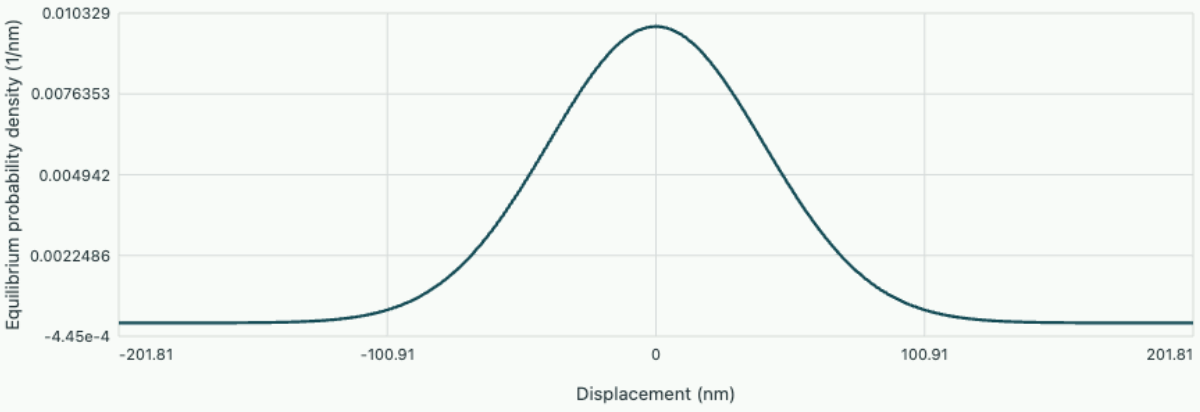

Harmonic equilibrium distribution

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.

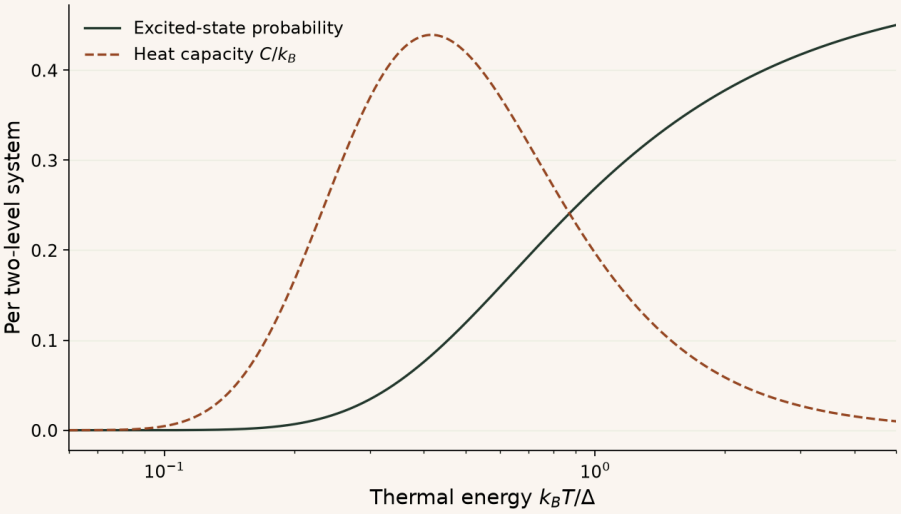

The two-level heat-capacity maximum marks the temperature range where changing the temperature most rapidly redistributes occupation between the levels.

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**

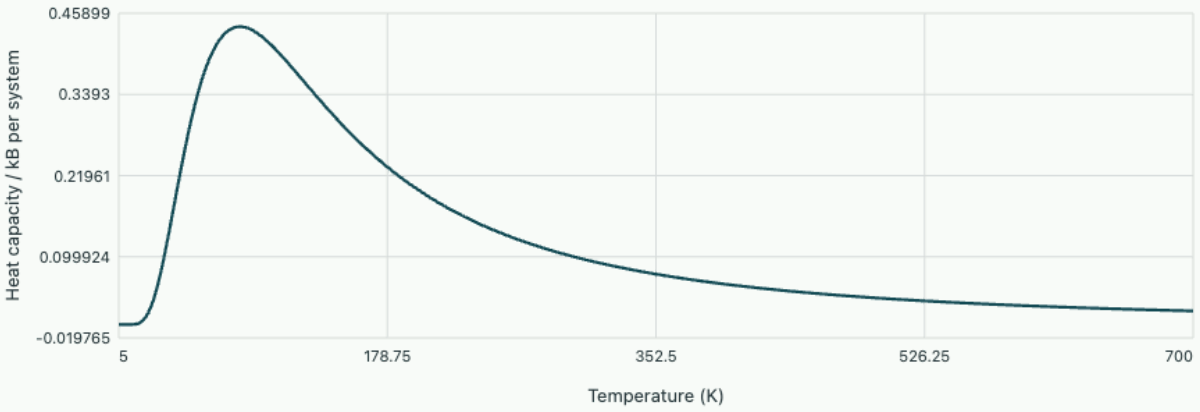

Two-level heat capacity

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)
```

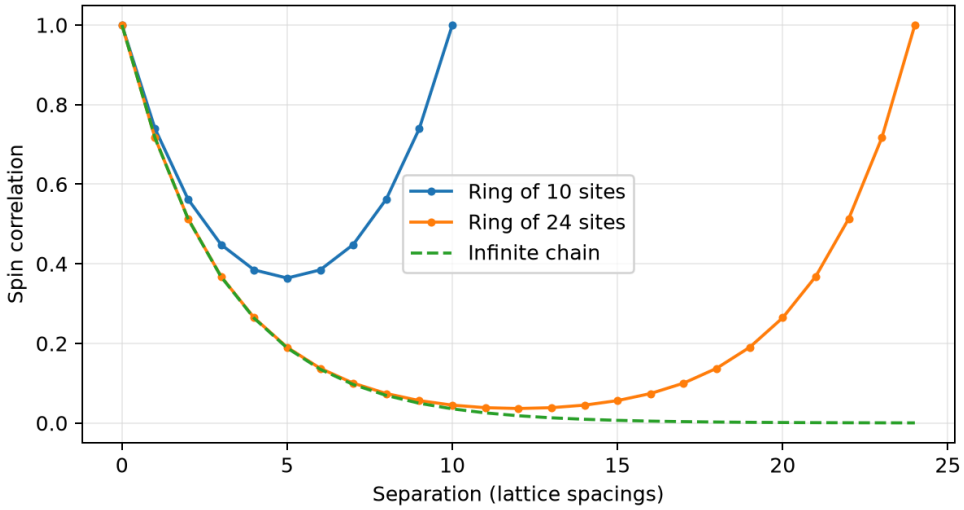

For J/(kBT) = 0.90 the antipodal correlation is 0.364 for ten spins and 0.0365 for 24. At a full turn the observable is the original spin squared.

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**

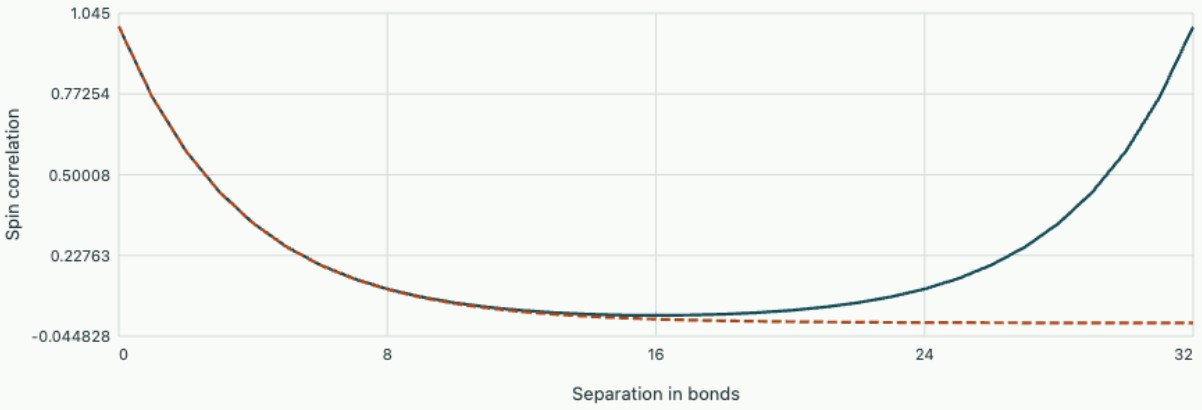

Closed finite ring Infinite chain

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.

Practise on sheet 3 · Infer a chemical potential difference from a fluctuating surface. Further reading · University of Cambridge · David Tong, Statistical Physics · §1.4.1 Grand Canonical Ensemble and §1.4.2 Grand Canonical Potential.

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.

Study the full derivation

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.

Chapter 6 · Indistinguishable particles and quantum gases

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

$$d[p/(k_BT)] = n\,dz/z.$$

Integration from the empty-gas limit gives the pressure series,

$$\frac{n\lambda_T^3}{g} = z + \frac{sz^2}{2^{3/2}} + O(z^3),$$
$$\frac{p\lambda_T^3}{gk_BT} = z + \frac{sz^2}{2^{5/2}} + O(z^3).$$

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,

$$\frac{p}{nk_BT} = 1 - \frac{s}{2^{5/2}}\frac{n\lambda_T^3}{g} + O(y^2).$$

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 $p/(nk_BT) \simeq 1.0113$. A measured classical pressure scale $nk_BT = 20.0\,\mathrm{pPa}$ 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

$$p/(nk_BT) = 0.9950.$$

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**

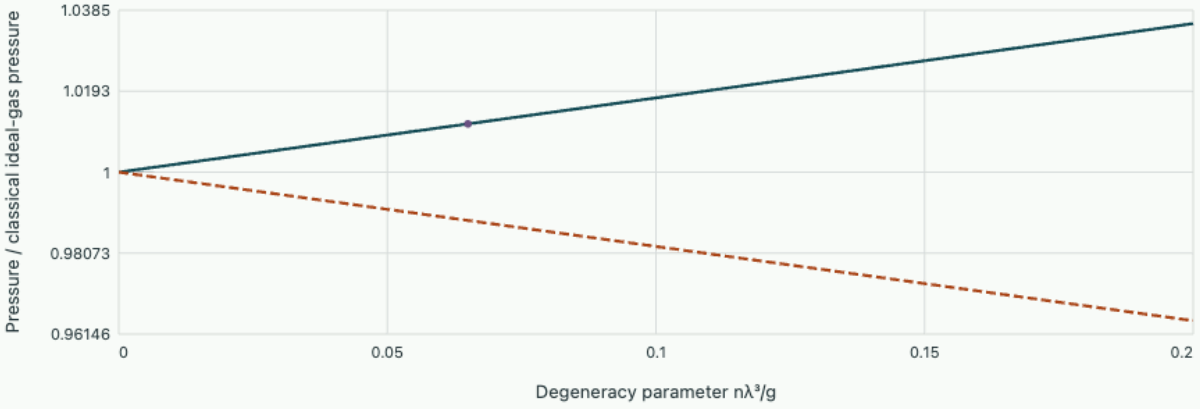

Fermi leading correction · Bose leading correction · Selected Fermi gas

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.

Study the full derivation

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.

Problem sheet 3 · Populations and quantum statistics

Use partition functions and single-state occupancies from Lectures 5 and 6 to diagnose thermal measurements.

A level with two ways to be excited

A unit has a nondegenerate ground state and two excited states of identical energy Δ. At some temperature the combined excited population is one half. Infer $\Delta/(k_B T)$, the energy variance and the heat capacity per unit at that temperature.

▼ **Distinguish a level's degeneracy from one state's probability**

The partition function is $1 + 2e^{-x}$ with $x = \Delta/(k_B T)$. Half occupation of the excited level group requires $2e^{-x} = 1$, so $x = \ln 2$. Individual excited states each have probability one quarter. Energy has two possible values with equal probabilities, hence variance $\Delta^2/4$. The fluctuation relation gives

$$C/k_B = (\ln 2)^2/4 = 0.120.$$

The degeneracy matters even though the two excited states carry the same energy.

Decide whether a dilute formula suffices

For a particular state,

$$(\epsilon - \mu)/(k_B T) = \ln 5.$$

Compute its Bose, Fermi and dilute-limit occupancies. Repeat for $\ln 101$ and decide whether the dilute estimate is accurate to 2% relative to both exact values.

▼ **Sum the single-state weights**

For the first state, Bose occupancy is $1/(5-1) = 0.25$, Fermi occupancy $1/(5+1) = 0.1667$ and the dilute approximation $1/5 = 0.20$. For the second the corresponding numbers are $1/100$, $1/102$ and $1/101$. The dilute estimate differs by $1/101 = 0.99\%$ relative to either exact result. Statistics matters when actual occupation per state is appreciable, not merely because particles are cold by an absolute temperature scale.

Compress an ideal electron gas

At zero temperature a spin-½ electron gas has Fermi energy 5.0 eV. Compress it to half its volume at fixed particle number. Find the changes in k_F, Fermi speed and Fermi energy. If the spin degeneracy could instead be reduced from two to one at the original volume, which of those scalings would repeat?

▼ **Recount the available momentum states**

State counting gives $N/V = g k_F^3/(6\pi^2)$ for degeneracy g. Doubling density scales k_F by $2^{1/3} = 1.260$. Since $v_F = \hbar k_F/m$, it gains the same factor. The energy grows by $2^{2/3} = 1.587$ to 7.94 eV. Removing one spin species at fixed density gives the same factors in this ideal counting exercise. A physical polarising field may also shift the energies of the species, which would require its own energetics.

Infer a chemical potential difference from a fluctuating surface

A surface has 500 independent single-occupancy sites at 300 K. Its mean population is 400. Infer $\mu - \epsilon$ and the predicted standard deviation. How would the variance change if the reservoir were adjusted to half occupation?

▼ **Invert the single-site probability**

With $p = 0.800$, the odds $p/(1-p) = 4 = e^{\beta(\mu-\epsilon)}$. Thus

$$\mu - \epsilon = k_B T \ln 4 = 35.8\,\mathrm{meV}.$$

The population variance is $500(0.8)(0.2) = 80$, giving standard deviation 8.94 particles. At half occupancy the variance rises to 125 despite a smaller mean. Correlated occupations or a finite reservoir can change the relation between mean and variance.

Revisit chapter 5 · Predict particle-number fluctuations in an open subsystem.

Explore these relations

- For the fluctuating surface, vary the chemical-potential bias and inspect occupation and number fluctuations. Open the live calculation.
- For the dilute-gas test, compare quantum corrections with the classical limit; the leading virial estimate applies only at small degeneracy. Open the live calculation.

Chapter 7 · Pressure and effusion from molecular motion

Use velocity distributions to predict gas pressure and determine why particles passing through a pinhole have a different average energy.

FROM FIRST PRINCIPLES

Count the particles that can reach a small opening

A small hole in a vessel of dilute equilibrium gas samples molecules moving toward its plane. Let n be number density, let $v_x > 0$ point outward, and let $f_x(v_x)$ be the normalized distribution of this velocity component. In a short time dt, only molecules starting within a layer of thickness $v_x dt$ can cross. Its volume is hole area A times that thickness. Summing over positive velocities gives the one-way number flux,

$$J = \frac{dN}{A\,dt} = n\int_0^\infty v_x f_x(v_x)\,dv_x.$$

At classical thermal equilibrium the component distribution is Gaussian with variance k_BT/m. Substitution and integration using the derivative of the exponential give

$$J = n\sqrt{\frac{m}{2\pi k_BT}}\int_0^\infty v_x e^{-mv_x^2/(2k_BT)}\,dv_x = n\sqrt{\frac{k_BT}{2\pi m}}.$$

This is the outgoing flux for an opening small compared with a mean free path and thin enough to avoid significant internal passage effects. Any gas outside can generate a reverse flux. The velocity factor in the integral explains why faster molecules are overrepresented among those leaving.

DERIVATION AND APPLICATIONS

Predict how an effusing gas mixture changes composition

A particle's contribution to the number flux through a small hole increases with its velocity normal to the opening. In a mixture at a common temperature this biases the escaping gas toward lighter species. Consider a rigid vessel of volume V, with hole area A much smaller than the wall area and hole diameter smaller than the molecular mean free path. Take the exterior pressure to be negligible. The walls maintain temperature T, and the reservoir re-equilibrates much faster than it empties. Thus the remaining species i has spatially uniform density N_i/V and a Maxwellian velocity distribution.

$$-\dot{N}_i = A\frac{N_i}{V}\int_0^\infty v_x\sqrt{\frac{m_i}{2\pi k_BT}}e^{-m_iv_x^2/(2k_BT)}dv_x = \frac{AN_i}{V}\sqrt{\frac{k_BT}{2\pi m_i}}.$$

The minus sign assigns positive outgoing flux to a decreasing inventory. Each species therefore has a constant decay rate in time under the isothermal assumption. Its time constant and its abundance ratio relative to another species are

$$\tau_i = \frac{V}{A}\sqrt{\frac{2\pi m_i}{k_B T}},$$

$$\frac{N_i(t)}{N_j(t)} = \frac{N_i(0)}{N_j(0)} e^{-t(1/\tau_i - 1/\tau_j)}.$$

At any instant the outgoing number ratio has a further velocity weighting compared with the gas left in the vessel,

$$\frac{|\dot{N}_i|}{|\dot{N}_j|} = \frac{N_i}{N_j}\sqrt{\frac{m_j}{m_i}}.$$

The ratio of molar masses may be used in this square root since the conversion to mass per particle cancels. A sample collected over a long time mixes successively different outlet compositions. A thick tube or a hole large enough for collective flow requires a different flow model.

WORKED EXAMPLE

Compare the outlet and reservoir compositions

Change the inputs and investigate this derivation.

Start with equal numbers of helium and neon atoms, taking molar masses 4.00 and $20.0\,\mathrm{g/mol}$. Initially the outlet helium-to-neon number ratio is $\sqrt{5} = 2.24$, giving a helium fraction $2.24/(1+2.24) = 0.691$. The initial reservoir helium fraction is 0.500. A detector reporting mass flow instead of particle count would multiply each contribution by its corresponding particle mass.

SECOND APPLICATION

Recover a heavy-gas residence time from a light tracer

Change the inputs and investigate this derivation.

In a separate chamber a helium tracer signal halves in $120\,\mathrm{s}$. For argon of molar mass $40.0\,\mathrm{g/mol}$ the predicted half-life is

$$120\sqrt{40.0/4.00} = 379\,\mathrm{s}.$$

After $240\,\mathrm{s}$, helium retains 0.250 of its initial number and argon retains $2^{-240/379} = 0.645$. With equal starting numbers the residual helium fraction is $0.250/(0.250+0.645) = 0.279$, showing the preferential depletion accumulated over time.

Why a leaking binary mixture changes its composition

A vessel containing two ideal gases is coupled to an empty chamber through a tiny hole. Both gases have the same temperature. Their mean speeds differ, so the gas remaining in the vessel, the gas passing through the

hole now and the gas accumulated downstream have different compositions.

Count the molecules that reach the opening

For a Maxwellian species i, the one-sided flux to a plane is $n_i \int_0^\infty v_z f_i(v_z)\, dv_z = n_i \sqrt{k_B T/(2\pi m_i)}$. Particles with larger positive normal velocity reach the opening from a thicker upstream slab in a fixed time. A raw count of half the particles would miss this velocity weighting. Assume the aperture is small compared with a mean free path locally, the vessel rethermalises rapidly compared with its depletion time, the temperature is maintained, and particles do not return from the collection chamber. For area A and reservoir volume V the number balance becomes

$$\frac{dN_i}{dt} = -k_i N_i,$$

$$k_i = \frac{A}{V}\sqrt{\frac{k_B T}{2\pi m_i}},$$

$$N_i(t) = N_i(0) e^{-k_i t}.$$

At any instant the light fraction in the vessel is $N_l/(N_l + N_h)$; the light fraction in the escaping beam is $k_l N_l/(k_l N_l + k_h N_h)$. Equal populations yield a larger instantaneous light fraction because $k_l/k_h = \sqrt{m_h/m_l}$. The amount collected over a long interval requires the integral of each flux. Starting from an empty collection chamber, that amount is $N_i(0) - N_i(t)$. Averaging the instantaneous fractions with equal time weights would over-represent the late, very weak beam.

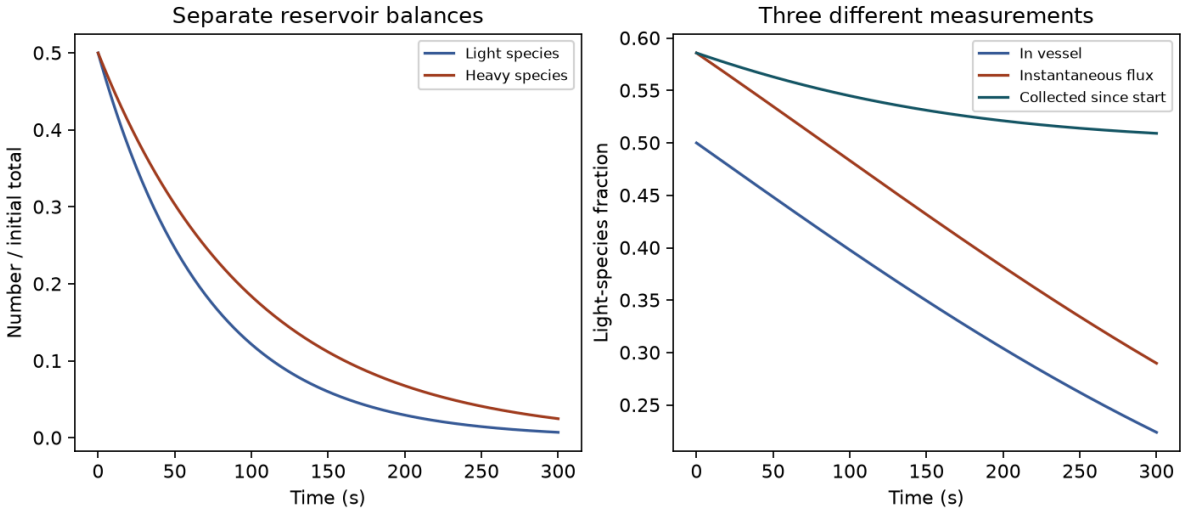

The escape flux initially favours light particles. At late times the reservoir has become enriched in the heavy species. The composition of the total collection approaches the original composition as the vessel empties.

Try it · Effusion mixture

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

Heavy / light molecular mass **4**

Initial light-species fraction **0.5**

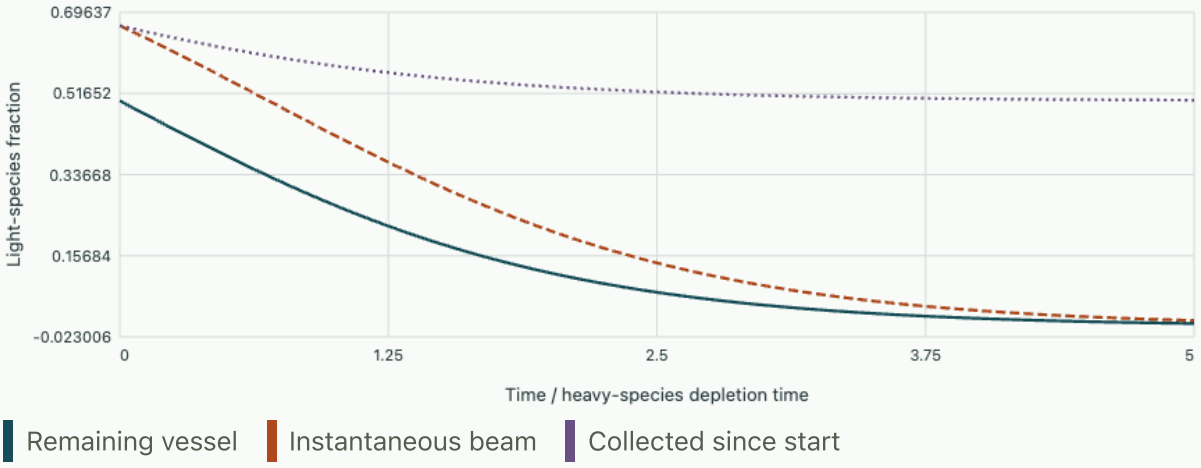

Light depletion rate is 2 times the heavy rate. Temperature is fixed and the aperture remains in the molecular-effusion regime.

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

Take equal starting populations, masses of 20 and 40 atomic mass units and a heavy-species depletion time of 100 s. Initially the outgoing light fraction is

$$\sqrt{2}/(1+\sqrt{2}) = 0.586.$$

At 100 s, the reservoir contains fractions $0.5e^{-\sqrt{2}}$ and $0.5e^{-1}$ of its original total in the light and heavy species. Its light fraction is therefore about 0.398, the instantaneous beam is about 0.483 light, and the collection since startup is about 0.545 light. These three compatible numbers describe different measurements at the same clock time.

Change the mass ratio in the notebook, then state what happens if the two masses approach equality. All three composition curves approach the initial fraction even though the total gas still leaks out. If the downstream chamber is closed and pressure builds up, molecules can return; the purely outward rate law

then fails. Similarly, a sufficiently large aperture supporting collisions and collective flow no longer provides this species-by-species effusive selection.

When does the escaping beam contain equal numbers of the two components in the default experiment? Does the gas remaining in the vessel also consist of equal numbers then?

▼ **Work through the solution**

Equate $k_l N_l(0) e^{-k_l t}$ and $k_h N_h(0) e^{-k_h t}$. For equal starting numbers, $t = \ln(k_l/k_h)/(k_l - k_h)$, about 83.7 s for the stated rates. The remaining ratio at that time is $N_l/N_h = k_h/k_l = 1/\sqrt{2}$, corresponding to a light fraction about 0.414. The faster rate compensates the light species' lower abundance in the outgoing beam.

Open editable notebook Download .ipynb

Pressure counts momentum flux

Consider a dilute gas at equilibrium with number density n, molecular mass m and no mean flow. Molecular diameters are small compared with intermolecular spacing. At a stationary wall an elastic specular collision reverses the normal component of velocity, transferring $2mv_x$ to the wall when $v_x > 0$. In time dt only molecules within a slab of thickness $v_x dt$ can reach it. Thus fast molecules strike the wall more frequently even at the same concentration.

With normalized velocity density $f(\mathbf{v})$ the pressure is $2nm \int_{v_x>0} v_x^2 f \, d^3v$. Reflection symmetry converts this to $nm\langle v_x^2 \rangle$. Isotropy makes each Cartesian variance a third of total mean square speed, giving $p = nm\langle v^2 \rangle/3$. Combining with the ideal gas equation gives

$$\tfrac{1}{2} m \langle v^2 \rangle = 3 k_B T/2.$$

Internal rotations and vibrations can store energy but do not directly carry translational momentum to the wall in this model.

Count particles crossing a hole

A small opening in a thin wall allows the crossing molecules to escape. Its diameter must be small compared with the mean free path, and the gas on the receiving side must be sufficiently dilute to neglect backflow. The outgoing flux per unit area is $J = n \int_{v_x>0} v_x f \, d^3v$. The positive-half integral of the Maxwellian normal component is $\sqrt{k_B T/(2\pi m)}$, equivalently

$$J = n\sqrt{\frac{k_B T}{2\pi m}} = \frac{n\bar{v}}{4},$$

$$\bar{v} = \sqrt{\frac{8 k_B T}{\pi m}}.$$

The factor of one quarter combines the half of velocity directions facing the hole with their angular projection. The emerging speed distribution is weighted by one extra power of speed relative to the equilibrium speed

distribution, since faster particles cross the hole more often. For a Maxwell gas it is proportional to $v^3 e^{-mv^2/(2k_BT)}$. Dividing its energy-weighted integral by its normalization gives mean translational energy $2k_BT$ per escaping molecule. A thermally isolated remaining gas can consequently cool as it leaks.

Estimate a molecular leak

Nitrogen of molecular mass 4.65×10^{-26} kg is at 300 K and pressure 1.00 Pa. The number density is $p/(k_BT) = 2.41 \times 10^{20}$ m⁻³ and mean speed is 476 m/s. A circular hole of diameter 20.0 µm has area 3.14×10^{-10} m². The emission rate JA is approximately 9.0×10^{12} molecules/s. This large count corresponds to only 4.2×10^{-13} kg/s. The small-hole condition must be checked independently using the collision mean free path.

At equal pressures and temperatures two pure gases have the same number density. Their effusion-rate ratio therefore depends only on inverse square root of molecular mass. If back pressure develops, apply the flux equation separately to both sides and subtract opposing fluxes. A pressure difference across a wide tube requires a flowing-gas model instead.

A pressure gauge attached to the receiving chamber reports a quantity affected by the pumping speed as well as by the leak. Weighing a collected condensate, measuring the source pressure decay in a known constant-temperature volume, or calibrating the pump independently can provide an absolute flux determination. Each method tests a different piece of the chain from measured pressure to inferred molecular rate.

A thermally isolated effusing gas cools as it loses particles

The speed distribution in an escaping beam differs from the distribution in its source. This selective loss can feed back on the source temperature. Consider a fixed, insulated vessel of volume V containing a dilute monatomic gas. A small aperture of area A lets molecules escape into vacuum. Assume collisions in the vessel establish a Maxwellian temperature much faster than particles leave, the aperture remains in the free-molecular regime, and thermal exchange with the vessel walls is negligible. These assumptions separate rapid internal equilibration from the slow evolution of temperature and particle number.

The gas stores $U = 3Nk_BT/2$, while the mean energy of each escaping molecule is $2k_BT$. During a small change $dN < 0$, the internal energy change is therefore

$$dU = 2k_BT\,dN.$$

Expanding the left side and collecting terms gives

$$dT/T = dN/(3N).$$

The resulting temperature law is

$$T/T_0 = (N/N_0)^{1/3}.$$

It predicts cooling because the outgoing beam carries more energy per particle than the source average, despite the absence of externally extracted mechanical work.

The effusive rate is

$$dN/dt = -(A/V)N\sqrt{k_B T/(2\pi m)}.$$

Define $\tau_0 = V/[A\sqrt{k_B T_0/(2\pi m)}]$. Substitution of the temperature law and integration give

$$\frac{N}{N_0} = \left(1 + \frac{t}{6\tau_0}\right)^{-6},$$

$$\frac{T}{T_0} = \left(1 + \frac{t}{6\tau_0}\right)^{-2}.$$

A fixed-temperature reservoir would instead exhibit a simple exponential decay in particle number. The initial logarithmic loss slopes agree, but cooling changes the later curve.

Predict pressure at half the initial gas inventory

Take argon of mass 6.63×10^{-26} kg initially at 300 K in a 1.00 litre vessel. A small hole of area 1.00×10^{-8} m² gives $\sqrt{k_B T_0/(2\pi m)} = 99.7$ m/s and $\tau_0 = 1003$ s. To reach $N/N_0 = 1/2$ requires

$$t = 6\tau_0(2^{1/6} - 1) = 737\,\mathrm{s}.$$

The gas temperature is then

$$300 \times 2^{-1/3} = 238\,\mathrm{K}.$$

Its pressure relative to the starting value is $(N/N_0)(T/T_0) = 0.397$ at fixed volume.

If a wall heater maintained 300 K instead, pressure at half inventory would be one half of its original value, reached after $\tau_0 \ln 2 = 695$ s. Monitoring both pressure and temperature can distinguish these cases. In practice chamber walls can have far greater heat capacity than a rarefied gas and reheat it rapidly. The isolated-gas prediction is therefore a limiting case that an experiment must establish, rather than something guaranteed merely by insulating the outer vessel.

The cooling law can be tested even if the absolute aperture area is uncertain. Plot $\ln(T/T_0)$ against $\ln(N/N_0)$ using simultaneous temperature and inventory records. The predicted slope is one third and does not contain the hole area, particle mass or vessel volume. The initial time derivative, by contrast, depends on all three. A correct logarithmic slope but an incorrect time constant therefore points toward an aperture calibration or transmission probability issue. A changing slope suggests thermal leakage, additional molecular energy stores or loss of internal equilibration during the measurement.

A seal or valve letting room-temperature gas enter the chamber also violates the assumed one-way balance and changes both logarithmic slopes.

Why does an effusing beam contain faster molecules?

In a dilute equilibrium gas the speed density contains a velocity-space factor u^2. Molecules crossing a small aperture are additionally weighted by their normal speed; the angular average leaves one extra power of speed. Normalise each distribution to unit area.

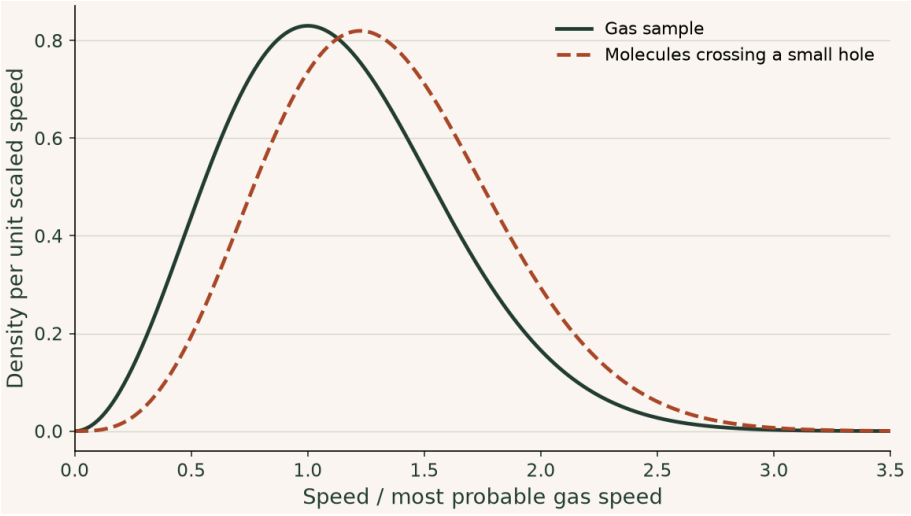

The gas distribution peaks at u = 1, the effusive distribution at $\sqrt{(3/2)} \approx 1.225$. A time-of-flight measurement of an effusing beam samples a flux-weighted population. Fitting it to the unweighted gas law biases a temperature estimate.

Try it · Speed law

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) 300

Molecular mass (atomic mass units) 28

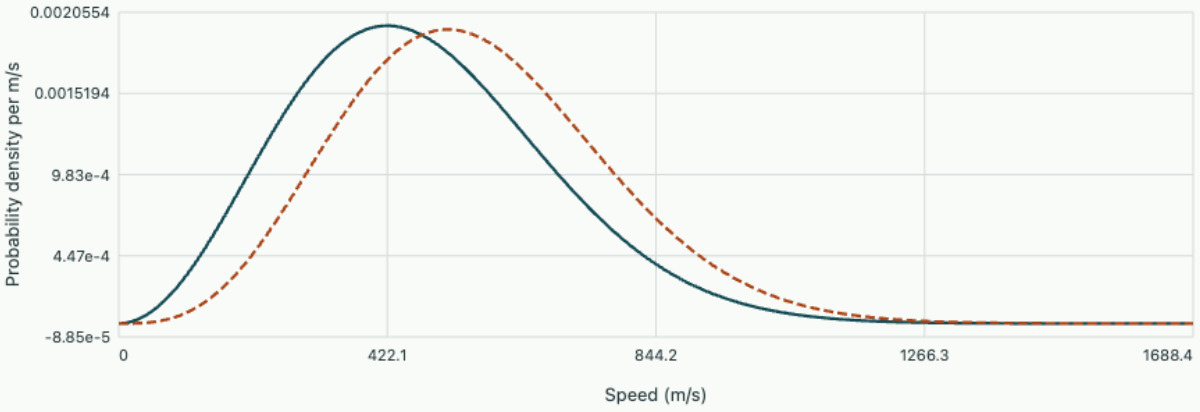

Gas particles Effusing flux

Gas mode 422.1 m/s; effusive mode 516.96 m/s. Both distributions integrate to one over all speeds.

▶ **Read values from the plot**

Study the full derivation

From binary collisions to the Boltzmann equation · Statistical mechanics from microscopic counting to material response

An initially anisotropic velocity distribution can relax even though every individual collision conserves energy and momentum. Describe a dilute gas with a phase-space density, derive the structure of its collision term and prove which distribution can stop producing entropy. Then examine the assumptions behind that proof and the measurements most likely to reveal their failure.

Chapter 8 · Diffusion from a random walk

Derive a transport coefficient from a microscopic model and separate a diffusion length from a relaxation time in a finite container.

FROM FIRST PRINCIPLES

Recover the diffusion equation by counting random steps

Consider independent particles on a line with sites spaced by ℓ. Every time interval τ, each particle takes one step left or right with equal probability. With no creation, loss or imposed drift, the expected concentration c at a site after a step equals the average of the earlier concentrations at its two neighbours,

$$c(x, t+\tau) = \tfrac{1}{2} c(x-\ell, t) + \tfrac{1}{2} c(x+\ell, t).$$

For a concentration changing slowly over one site and one step, expand both sides in Taylor series. The zeroth-order terms cancel and the two first spatial derivatives have opposite signs. The leading surviving change gives

$$\tau\, \partial_t c \simeq \frac{\ell^2}{2} \partial_x^2 c,$$

$$D = \frac{\ell^2}{2\tau}.$$

The resulting diffusion equation is

$$\partial_t c = D\, \partial_x^2 c.$$

Its coefficient has units of length squared per time. A constant concentration has zero curvature, so the mean concentration stays fixed even while individual particles keep moving. Unequal step probabilities would leave a first-derivative drift term; the model assumes the microscopic motion has no preferred direction.

DERIVATION AND APPLICATIONS

Derive a penetration depth for periodically driven diffusion

A periodically varying concentration at a surface penetrates a diffusive medium with a measurable delay and amplitude loss. Treat the medium as a uniform half-space $x > 0$ with constant diffusivity D, no advection and no reaction. At $x = 0$ impose $c = c_0 + C\cos(\omega t)$ with $C < c_0$. Far inside the sample the perturbation vanishes. Wait until the start-up transient has decayed at the observation positions. The diffusion equation is linear, so one can look for the real part of a single complex harmonic with a spatial attenuation factor,

$$c(x,t) - c_0 = \mathrm{Re}\left[C e^{i\omega t - q x}\right],$$

$$i\omega = D q^2.$$

There are two complex square roots. The condition at large x selects the root with positive real part. Its real and imaginary parts have the same magnitude,

$$q = (1+i)\sqrt{\frac{\omega}{2D}} = \frac{1+i}{\delta},$$

$$\delta = \sqrt{\frac{2D}{\omega}}.$$

Taking the real part gives the decaying oscillation and directly connects two independent measurements to the same transport coefficient,

$$c - c_0 = C e^{-x/\delta}\cos(\omega t - x/\delta),$$

$$\ln\frac{C(0)}{C(x)} = \omega\Delta t = \frac{x}{\delta}.$$

Here Δt is the delay of the corresponding peak, with the phase unwrapped across depth. Diffusion causes strong attenuation over a full phase cycle, so the apparent travel of a peak does not imply rigid translation of a material parcel. A reflecting far boundary would change this result once the sample thickness becomes comparable to δ.

WORKED EXAMPLE

Determine diffusivity from an amplitude measurement

Change the inputs and investigate this derivation.

A surface concentration oscillates with period $100\,\mathrm{s}$. At $200\,\mu\mathrm{m}$ depth its amplitude is 0.368 times the boundary amplitude. This gives $\delta \simeq 200\,\mu\mathrm{m}$, hence

$$D = \pi\delta^2/(100\,\mathrm{s}) = 1.26\times10^{-9}\,\mathrm{m}^2/\mathrm{s}.$$

The model also predicts a lag of one radian, corresponding to $100/(2\pi) = 15.9\,\mathrm{s}$.

SECOND APPLICATION

Check a transport coefficient with a phase measurement

Change the inputs and investigate this derivation.

For a second tracer the forcing period is $80.0\,\mathrm{s}$. A peak at $150\,\mu\mathrm{m}$ depth occurs $10.0\,\mathrm{s}$ after the corresponding surface maximum. Its phase lag is $\pi/4$, giving $\delta = 191\,\mu\mathrm{m}$ and $D = 1.43\times10^{-9}\,\mathrm{m}^2/\mathrm{s}$. The independent amplitude prediction is $e^{-\pi/4} = 0.456$ of the surface signal. A calibrated amplitude far from this fraction would challenge the assumed diffusion-only dynamics.

Try it · Diffusion skin

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

Diffusivity (10^{-9} m²/s) **1.2**

Forcing period (s) **100**

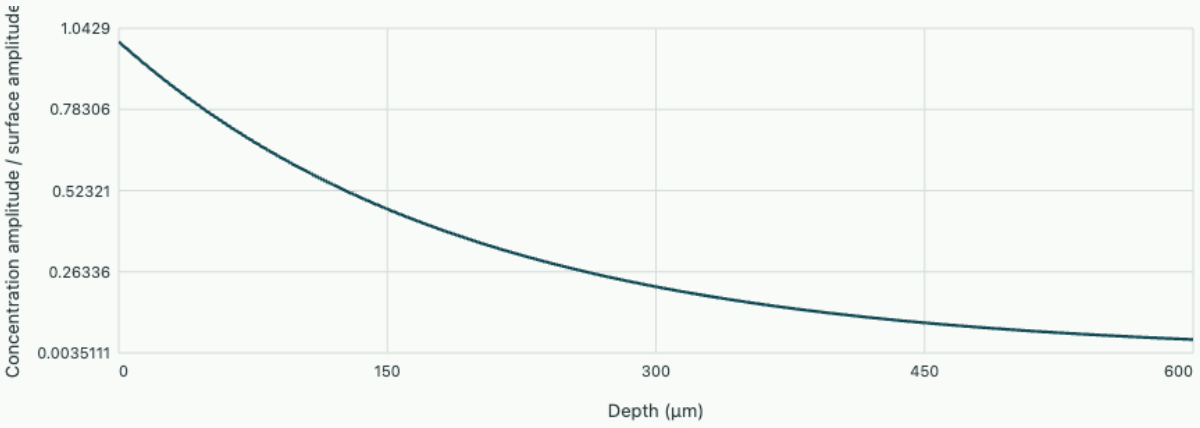

Oscillating boundary concentration

Penetration depth 195.44 µm; the phase lag is x/δ radians. Amplitude and phase supply independent checks of D.

A spreading cloud without a mean velocity

Release a narrow band of neutral tracer particles into a stationary uniform medium. Repeated collisions erase velocity memory. At times long compared with the collision time, labelled particles can spread while the mean tracer velocity remains zero. Start with an ideal symmetric random walk on a line. During a time step τ every particle moves a distance ℓ left or right with equal probabilities, independently of previous steps.

The concentration in cell x after a step is the mean of the two neighbouring concentrations before that step. Expanding the left side in τ and the right side in ℓ cancels all odd spatial derivatives, leaving

$$\partial_t c = D\,\partial_x^2 c,$$

$$D = \ell^2/(2\tau).$$

After $N = t/\tau$ independent steps the variance is $N\ell^2 = 2Dt$, confirming the same coefficient without a continuum expansion. The exact prefactor depends on which microscopic step model is used. In an isotropic gas the directional velocity correlations and distribution of free-flight lengths matter; an order-of-magnitude

estimate is $D \sim \overline{v}\ell_{\mathrm{mfp}}/3$. Substituting a fixed one-dimensional step length into this three-dimensional estimate would mix incompatible models.

Why a gradient drives current

Consider the cell boundary between x and $x+\ell$. The two crossing populations are proportional to the concentrations on their respective sides. Outgoing minus incoming probability current is therefore proportional to $c(x) - c(x+\ell)$. Taylor expansion gives Fick's law $j = -D\partial_x c$. Local conservation $\partial_t c = -\partial_x j$ recovers diffusion for constant D and remains useful when the transport coefficient varies in space. The continuum model assumes concentration changes little over one microscopic step.

In an unbounded domain the response to a narrow initial band is Gaussian at late times with variance increasing by $2Dt$. In a finite closed channel the final concentration is instead uniform. Reflecting ends impose zero gradient, so a prepared concentration modulation $c(x,0) = c_0 + a\cos(\pi x/L)$ relaxes through a single amplitude $a(t)$. Substitution into diffusion gives

$$a(t) = a(0)e^{-D\pi^2 t/L^2}.$$

Spatial wavelength matters quadratically because a broad concentration feature supplies only a weak gradient for the mass that must move.

Measure a coefficient two ways

Suppose a dyed band in water broadens from variance 0.250 mm² to 0.610 mm² in 300 s, with negligible flow. Taking the difference removes the initial width and gives $D = 0.360/(600) = 0.000600$ mm²/s, or 6.00×10^{-10} m²/s. In a 2.00 mm sealed channel, the longest nonuniform cosine mode then decays by $1/e$ in $L^2/(\pi^2 D) = 675$ s. The prediction connects an open-domain variance measurement to a finite-domain contrast decay.

Bulk drift shifts the measured band centroid and can be removed by centring each profile. Shear, by contrast, can spread the band even after centring. Independent checks of the variance law and of the finite-channel decay help distinguish molecular diffusion from unresolved fluid motion.

Over times short compared with a collision time, individual tracer trajectories are approximately straight and the displacement variance initially grows as the square of time. A fit to a diffusion coefficient should exclude this ballistic region. At long times finite walls can likewise limit the spread.

Add a weak force and recover the Einstein relation

A tracer concentration in a fluid can relax even when an external force produces a steady drift. Consider noninteracting overdamped particles at uniform temperature in one dimension. Define mobility b by the drift speed $v_d = bF$ under a constant, sufficiently weak mechanical force F. If concentration changes gently in space, the particle current combines a systematic drift and Fick transport:

$$j = bcF - D\,\partial_x c.$$

The two contributions can cancel at equilibrium even though individual particles remain in motion.

For a conservative force $F = -\partial_x U$, equilibrium concentration must have the Boltzmann form $c_{\mathrm{eq}} \propto e^{-U/(k_B T)}$. Differentiating it gives

$$\partial_x c_{\mathrm{eq}} = c_{\mathrm{eq}} F/(k_B T).$$

Requiring zero current for this equilibrium distribution for arbitrary weak potentials gives

$$D = b k_B T.$$

This is the Einstein relation for a dilute thermal tracer. A fluctuation quantity measured without an imposed force, its diffusion coefficient, determines a response quantity measured by applying a force. It presumes local thermal equilibrium, linear mobility and the same degrees of freedom in both measurements. Self-propelled particles, crowded media or substantial hydrodynamic coupling may invalidate this simple one-particle relation.

Use a sedimentation profile to weigh a colloid

Set the upward coordinate to z and let m_b be the buoyant particle mass after subtracting displaced liquid mass. Its downward force is $-m_b g$. Equilibrium then predicts $c(z) = c(0) e^{-z/\ell_g}$ with gravitational length $\ell_g = k_B T/(m_b g)$. Suppose an optical count at 298 K falls by a factor of four over 30.0 µm. The measured length is $\ell_g = 30.0/\ln 4 = 21.6$ µm. Thus

$$m_b = k_B T/(g \ell_g) = 1.94 \times 10^{-17}\,\mathrm{kg}.$$

This is a mass contrast with the solvent, so an independently known density difference is needed to infer particle volume.

An independent tracking experiment gives $D = 0.400$ µm²/s. The mobility is $D/(k_B T) = 9.72 \times 10^{7}$ m/(N s). The downward drift speed is $b m_b g = D/\ell_g = 0.0185$ µm/s. Without boundaries, a mean drift of one gravitational length would therefore take about 1170 s, while random displacement over that interval remains substantial for an individual particle. A large ensemble is needed to resolve the mean against diffusion. In the density-matched limit $m_b \to 0$, the sedimentation length diverges and the vertical profile becomes uniform even though measurements of diffusion and of mobility under a different externally applied force still yield finite values.

The mean-square displacement used in mobility inference should be computed after any steady drift has been removed. Over an interval Δt, a downward drift gives mean displacement $-v_d \Delta t$, while one-dimensional thermal variance is $2D\Delta t$. The raw mean-square displacement about zero contains the extra term $v_d^2 \Delta t^2$. Fitting that quantity to a linear relation at long times can overestimate diffusion and consequently distort the mobility comparison. Independent sedimentation and tracking measurements become especially valuable when the drift is too small to resolve in a single trajectory but emerges reliably in an ensemble average.

The statistical precision of the variance estimate also depends on the number of independent trajectories.

Study the full derivation

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.

Problem sheet 4 · Molecular leaks and diffusing tracers

Use microscopic counting to distinguish a mean speed, an escape rate and a diffusive spreading rate.

Compare gases in a molecular leak

Helium and argon are separately admitted at equal number density and temperature to the same tiny hole in a thin wall. With atomic masses 4.00 and 40.0 u, find their number-flux and mass-flux ratios. Determine which species is enriched in the emerging beam from an equimolar mixture.

▼ **Show the flux weighting**

The outgoing number flux is $J = n\sqrt{k_B T/(2\pi m)}$, so

$$J_{\mathrm{He}}/J_{\mathrm{Ar}} = \sqrt{40/4} = 3.16.$$

Multiplying by particle mass reverses the trend, giving helium-to-argon mass flux ratio 0.316. In an equimolar mixture the predicted helium number fraction among emerging atoms is

$$3.16/(1+3.16) = 0.760.$$

This conclusion uses the collisionless aperture regime. A bulk hydrodynamic flow would entrain both components differently.

Use variance without assuming a point release

The standard deviation of a dye profile grows from 0.300 to 0.500 mm in 200 s. The centroid simultaneously shifts by 0.400 mm. Infer the diffusivity and mean drift speed, then predict both observables after another 200 s under constant conditions.

▼ **Show the moment calculation**

Variance grows from 0.0900 to 0.250 mm², giving $D = (0.250 - 0.0900)/(2 \times 200) = 4.00 \times 10^{-4}$ mm²/s. The mean drift is

$$0.400/200 = 0.00200\,\mathrm{mm/s}.$$

After another 200 s the variance is 0.410 mm², or standard deviation 0.640 mm. The centroid is displaced by a total 0.800 mm from its starting point. A second moment about a fixed origin would also contain the squared centroid displacement and overestimate diffusion.

Choose a faster concentration test

A closed 3.00 mm channel supports initial concentration modulations proportional to $\cos(\pi x/L)$ and $\cos(3\pi x/L)$. For $D = 5.0 \times 10^{-10}$ m²/s find both half-decay times. Explain why an apparently similar

measurement using a sinusoid could violate the boundary condition.

▼ **Show the mode selection**

No-flux ends require $\partial_x c = 0$ at $0, L$, satisfied by the cosines. Substitution into diffusion gives

$$t_{1/2} = L^2 \ln 2 / (D n^2 \pi^2).$$

The fundamental takes 1264 s and the third harmonic 140 s. The shorter scale drives a larger flux gradient. A simple $\sin(n\pi x/L)$ instead has nonzero end slopes, corresponding to an exchange of material at the walls. A measured decay cannot be used with the same boundary model unless the prepared mode actually satisfies it.

Explore these relations

- For the molecular leak, set the heavy-to-light mass ratio to ten and the initial light fraction to one half, then compare the vessel and emerging beam. Open the live calculation.
- For the concentration test, vary diffusion and driving frequency and inspect penetration and phase delay. Open the live calculation.

Notebook practice · Why a leaking binary mixture changes its composition

Chapter 9 · Thermal radiation in a cavity

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**

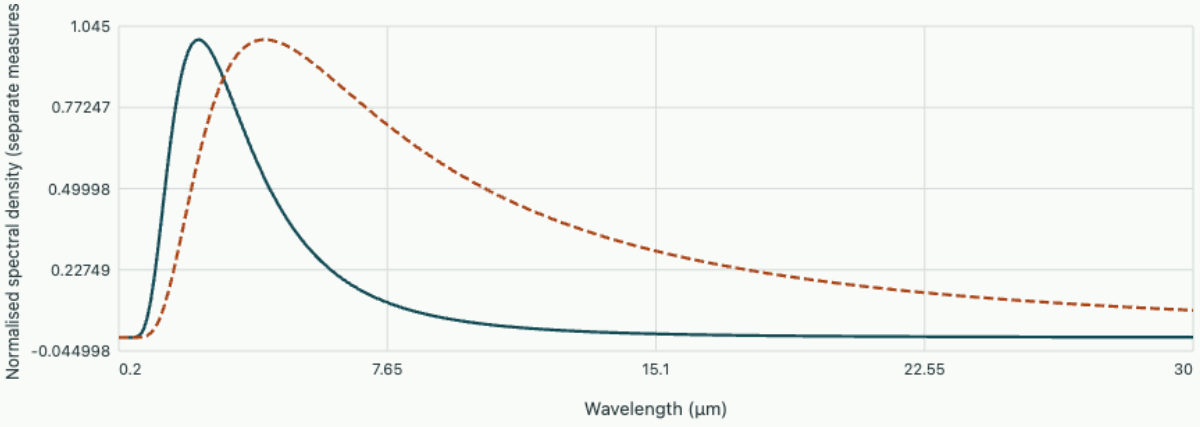

Equal-wavelength channels Equal-frequency channels, labelled by c/ν

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 $\omega^2/(\pi^2c^3)$. Hence

$$u_\omega(T) = \frac{\hbar\omega^3}{\pi^2c^3}\frac{1}{e^{\hbar\omega/k_BT}-1}.$$

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.

Study the full derivation

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.

Chapter 10 · Vibrations and the Debye heat capacity

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**

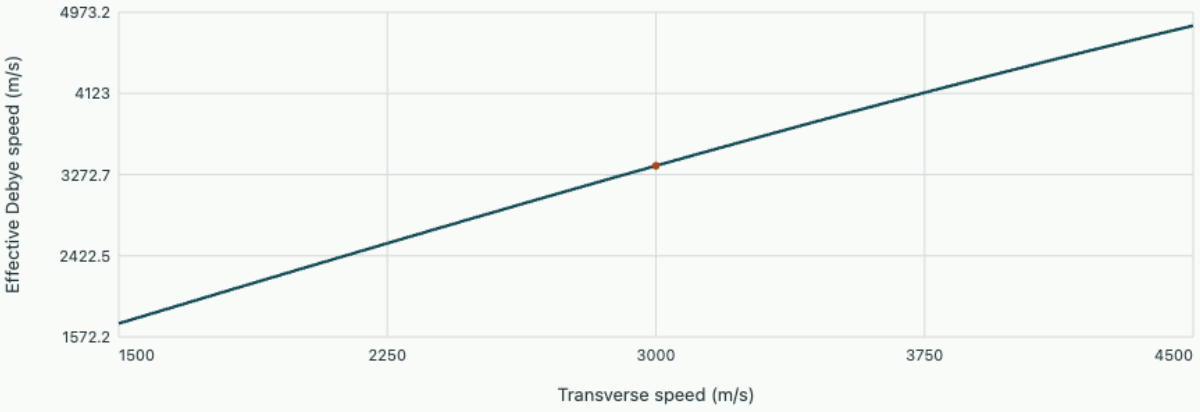

Three acoustic polarisations Selected material

$1/vD^3=(1/vL^3+2/vT^3)/3$. 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.

Practise on sheet 5 · Discriminate two low-temperature laws with a ratio. Further reading · University of Oxford · S. Simon, Lecture Notes for Solid State Physics · Chapter 2 · Specific Heat of Solids: Boltzmann, Einstein, and Debye · TUM / Walther-Meißner-Institut · R. Groß, Physik der kondensierten Materie I · §6.1.5 · Debye- und Einstein-Modell.

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 $g(\omega)\hbar\omega/[\exp(\hbar\omega/k_BT)-1]$ 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_BT 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

$$F_j = \hbar\omega_j/2 + k_BT \ln[1 - e^{-\hbar\omega_j/(k_BT)}].$$

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.

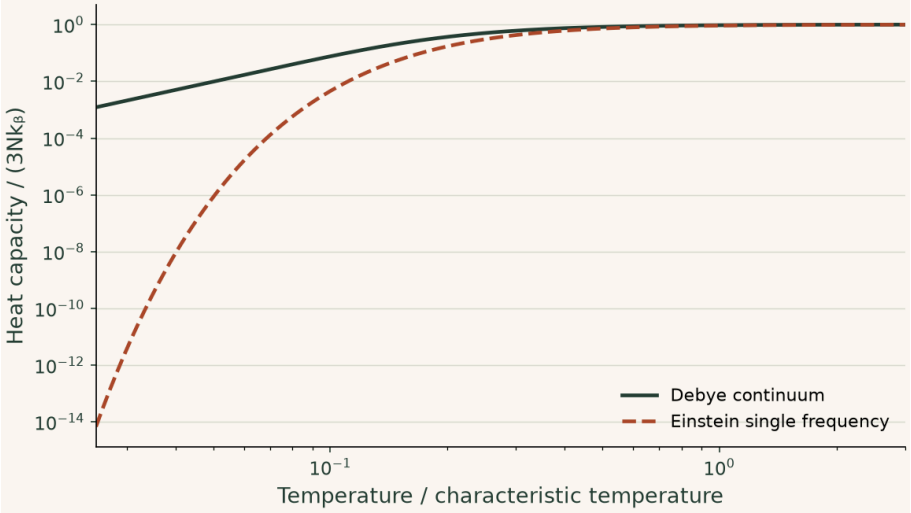

The single-frequency solid loses thermal excitations exponentially. The acoustic continuum still has low-frequency modes and its low-temperature curve follows a power law. Equal characteristic temperatures here set a comparison scale; fitting real data can give different values for the two models.

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**

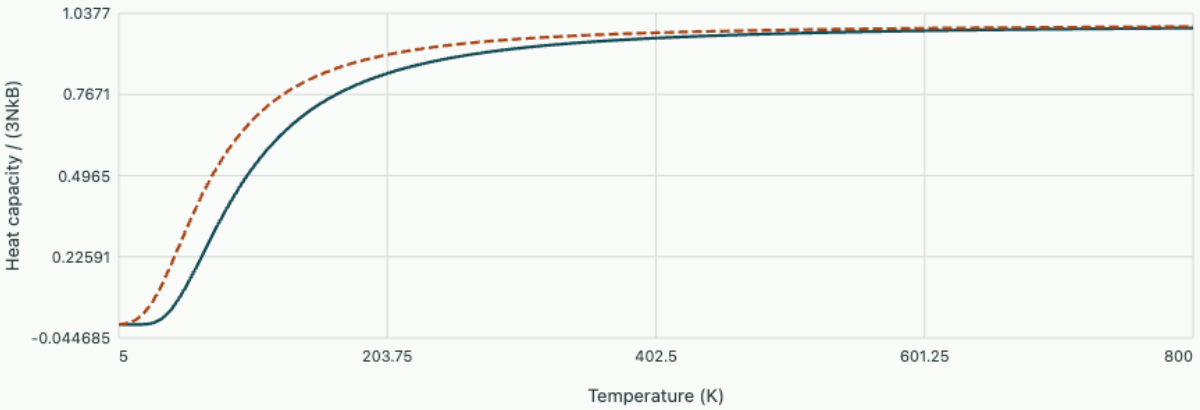

Einstein single-frequency model Debye acoustic continuum

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**

Study the full derivation

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.

Problem sheet 5 · Radiation and vibrations as mode gases

Connect equilibrium state counting to radiative transfer and distinguish a lattice heat capacity from electronic storage.

Two black surfaces

Two large parallel black surfaces at 350 K and 300 K face one another across vacuum. Calculate the net radiative flux. Linearize it about 325 K and assess the error. Use $\sigma = 5.6704 \times 10^{-8}$ W m⁻² K⁻⁴.

▼ **Show the transfer calculation**

Net flux is

$$\sigma(350^4 - 300^4) = 391.6\,\mathrm{W/m^2}.$$

The linear estimate $4\sigma(325)^3(50) = 389.3$ W/m² is about 0.6% low. For temperatures $T_m \pm \Delta T/2$, expanding the fourth powers gives $4T_m^3\Delta T + T_m(\Delta T)^3$. The omitted positive term explains the bias and shows how the relative error scales as $(\Delta T/T_m)^2$.

The missing calorimeter coefficient

Low-temperature molar measurements on a metal give $C/T = 4.00$ mJ mol⁻¹ K⁻² at 2.00 K and $C/T = 10.0$ mJ mol⁻¹ K⁻² at 4.00 K. Using $C = \gamma T + \beta T^3$, extract both coefficients and the atomic Debye temperature. Take $12\pi^4 R/5 = 1944$ J mol⁻¹ K⁻¹.

▼ **Show the separation of contributions**

A plot of C/T against T^2 has slope $\beta = (10 - 4)/(16 - 4) = 0.500$ mJ mol⁻¹ K⁻⁴ and intercept $\gamma = 4 - 0.5(4) = 2.00$ mJ mol⁻¹ K⁻². From $\beta = 1944/\Theta_D^3$ obtain

$$\Theta_D = (1944/0.000500)^{1/3} = 157\,\mathrm{K}.$$

The electronic term dominates below $\sqrt{\gamma/\beta} = 2.00$ K, provided the same low-temperature normal state persists.

Change the number of spatial dimensions

An idealized large sheet supports one acoustic branch with linear dispersion $\omega = vk$ in two dimensions. Derive the low-temperature power of the thermal energy and heat capacity. Explain why a freely suspended membrane with a flexural branch may behave differently.

▼ **Show the mode count**

The two-dimensional reciprocal-space annulus has area $2\pi k\,dk$, hence $g(\omega) \propto \omega$ for the linear branch. Integrating $g(\omega)\hbar\omega/[e^{\hbar\omega/k_BT} - 1]$ with $x = \hbar\omega/k_BT$ gives

$$U_{\mathrm{th}} \propto T^3,$$

$$C \propto T^2.$$

A quadratic flexural branch instead gives $k\,dk$ proportional to $d\omega$, so the density of states is approximately constant and its contribution is proportional to T in heat capacity. The temperature exponent diagnoses both dimensionality and dispersion.

Discriminate two low-temperature laws with a ratio

A mole of insulating crystal is measured at fixed volume at 40 K and 20 K. The recorded lattice heat capacities are 0.800 and 0.101 J/(mol K). Compare their ratio with a Debye low-temperature prediction and with an Einstein model of $\Theta_E = 240$ K. Explain what conclusion these two readings support.

▼ **Worked solution**

The measured ratio $C_V(20)/C_V(40) = 0.12625$ is close to the low-temperature acoustic prediction $(20/40)^3 = 0.125$. Applying the Einstein formula with $x = 6$ and 12 gives 2.237 and 0.02207 J/(mol K), a ratio of 0.00987. Even if one rescales its amplitude to match at 40 K, this proposed Einstein frequency predicts too sharp a fall.

The two measurements support an acoustic-mode contribution over this temperature interval. To establish the asymptotic exponent, take additional low-temperature measurements and examine C_V/T^3, including uncertainty in both the calorimetry and the subtraction of electronic or background heat capacity. The ratio test alone supplies neither the whole phonon spectrum nor a unique microscopic model.

Revisit chapter 10 · Test the single-frequency model before introducing a spectrum.

Explore these relations

- For mode counting and thermal radiation, compare the frequency and wavelength spectra and their temperature scaling. Open the live calculation.
- For the calorimeter laws, compare the Einstein and Debye heat capacities and their low-temperature limits. Open the live calculation.

Chapter 11 · Chemical equilibrium and reaction extent

Minimize Gibbs energy along a permitted reaction and predict a composition change when the system is compressed.

FROM FIRST PRINCIPLES

Build the reaction quotient from chemical potentials

At fixed temperature and pressure, a closed reacting mixture equilibrates by minimizing Gibbs energy subject to conservation of its constituent elements. Let ν_i be the balanced stoichiometric coefficient of species i, positive for a product and negative for a reactant. Advancing the reaction by a small extent $d\xi$ changes each amount by $dn_i = \nu_i d\xi$. The free-energy change per mole of reaction is then

$$\frac{dG}{d\xi} = \sum_i \nu_i \mu_i = \Delta_r G.$$

For an interior equilibrium this derivative vanishes. Chemical potentials can be written $\mu_i = \mu_i^\circ + RT \ln a_i$ using dimensionless activities relative to specified standard states. Collecting all the logarithms gives

$$\Delta_r G = \Delta_r G^\circ + RT \ln Q,$$

$$Q = \prod_i a_i^{\nu_i},$$

$$K = e^{-\Delta_r G^\circ/(RT)}.$$

The equilibrium composition satisfies $Q = K$. For an ideal gas, $a_i = p_i/p^\circ$; nonideal mixtures require activities measured or modeled for the mixture. If $Q < K$, a small forward reaction lowers G. This thermodynamic preference determines a direction near the present state while leaving the reaction timescale to kinetics.

Background · Chemical potential

$$\mu = (\partial F/\partial N)_{T,V}$$

This change in free energy per added particle controls exchange between subsystems. In a finite small system an actual finite difference may matter.

DERIVATION AND APPLICATIONS

Propagate an equilibrium constant when reaction heat capacity changes the enthalpy

An equilibrium constant quoted at one temperature does not determine the equilibrium composition at another unless the energy change of the reaction is also known. Consider a reversible reaction with a

dimensionless standard equilibrium constant $K(T)$ and unchanged standard-state choices. All reaction quantities below are per mole of reaction as written. At fixed standard pressure, the Gibbs differential $dG = -S\,dT$, together with $G = H - TS$, gives

$$d(G/T)/dT = -H/T^2.$$

Apply this identity to the stoichiometric differences of standard molar quantities,

$$\ln K = -\frac{\Delta_r G^\circ}{RT},$$
$$\frac{d\ln K}{dT} = \frac{\Delta_r H^\circ(T)}{RT^2}.$$

If heat capacities of products and reactants differ, the enthalpy changes with temperature. Over an interval where their stoichiometric difference $\Delta_r C_p^\circ$ can be treated as constant, integration of $dH^\circ = C_p^\circ dT$ gives

$$\Delta_r H^\circ(T) = \Delta_r H^\circ(T_1) + \Delta_r C_p^\circ (T - T_1).$$

Insert this linear enthalpy into the derivative for $\ln K$. The constant part integrates as $1/T^2$; the term proportional to temperature also produces a logarithm. The two-temperature relation is

$$\ln\frac{K_2}{K_1} = \frac{\Delta_r H^\circ(T_1)}{R}\left(\frac{1}{T_1} - \frac{1}{T_2}\right) + \frac{\Delta_r C_p^\circ}{R}\left[\ln\frac{T_2}{T_1} + \frac{T_1}{T_2} - 1\right].$$

Both terms vanish when the two temperatures coincide. For a small temperature interval the heat-capacity correction begins at second order, consistent with the local slope depending only on the current reaction enthalpy. Phase changes would interrupt the assumed smooth enthalpy law.

WORKED EXAMPLE

Estimate a change in an association equilibrium

Change the inputs and investigate this derivation.

Take $K_1 = 4.00$ at $300\,\mathrm{K}$, $\Delta_r H^\circ(T_1) = -20.0\,\mathrm{kJ/mol}$, and $\Delta_r C_p^\circ = 100\,\mathrm{J/(mol\,K)}$. At $360\,\mathrm{K}$ the enthalpy term in $\ln(K_2/K_1)$ is -1.336 and the heat-capacity term is $+0.188$. Thus

$$K_2 = 4.00 e^{-1.148} = 1.27.$$

The association remains exothermic throughout the interval, although its magnitude declines on heating.

SECOND APPLICATION

Infer enthalpy from equilibrium measurements

Change the inputs and investigate this derivation.

For a different reaction K rises from 0.500 at $290\,\mathrm{K}$ to 0.900 at $310\,\mathrm{K}$. If the reaction heat-capacity difference is negligible here, inversion gives

$$\Delta_r H^\circ = R\ln(0.900/0.500)/(1/290 - 1/310) = 22.0\,\mathrm{kJ/mol}.$$

The positive sign is consistent with the increased equilibrium constant on heating. At $300\,\mathrm{K}$ the same enthalpy predicts a local logarithmic slope of $0.0293\,\mathrm{K}^{-1}$, testable with a more closely spaced measurement.

Try it · Vanthoff

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

Reaction enthalpy at 300 K (kJ/mol) **-20**

Reaction heat-capacity difference (J/mol/K) **100**

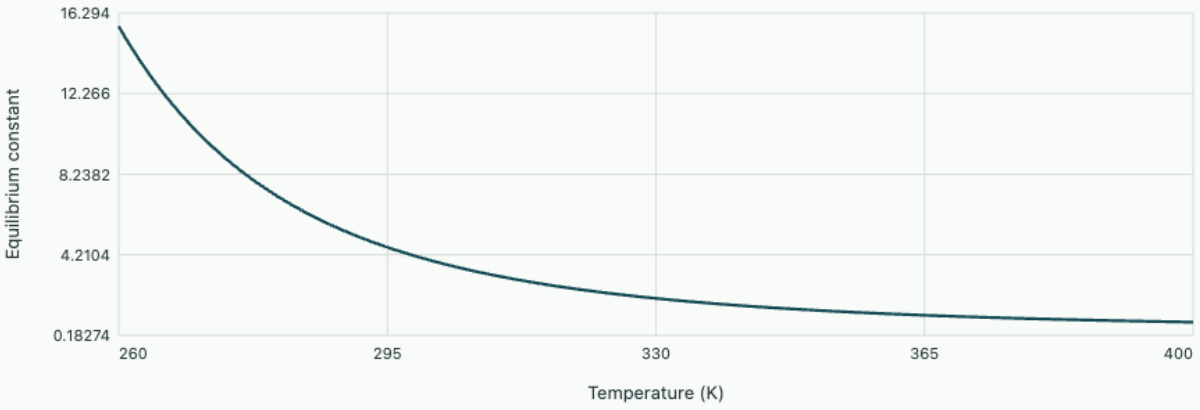

Integrated temperature-dependent enthalpy

K=4 at 300 K, ΔH(T)=ΔH(300)+ΔCp(T−300). Activities and standard states are held consistent.

A reaction has one constrained direction

A closed vessel can change particle species while conserving the total number of atoms of each element. At fixed temperature and pressure the equilibrium state minimizes Gibbs energy. Introduce a reaction extent ξ in moles and stoichiometric coefficients ν_i, negative for reactants and positive for products. A small reaction advance changes amounts by $dn_i = \nu_i\, d\xi$. The Gibbs differential at fixed T, p therefore gives

$$dG = \sum_i \mu_i \, dn_i = \Delta_r G \, d\xi,$$

$$\Delta_r G = \sum_i \nu_i \mu_i.$$

An interior minimum satisfies $\Delta_r G = 0$. A negative value means that a small positive advance lowers G. Reversibility of the microscopic chemistry is assumed for equilibrium to be accessible. Thermodynamic preference alone does not provide the time needed to reach it.

Activities and the equilibrium constant

For ideal gases $\mu_i = \mu_i^\circ(T) + RT \ln(p_i/p^\circ)$. The activity is the dimensionless ratio $a_i = p_i/p^\circ$. Inserting this expression into the reaction Gibbs energy gives

$$\Delta_r G = \Delta_r G^\circ + RT \ln Q,$$

with $Q = \prod_i a_i^{\nu_i}$. The equilibrium value of Q is $K = \exp[-\Delta_r G^\circ/(RT)]$. Standard pressure supplies the reference that makes the logarithm and the equilibrium constant dimensionless even when the total molecular count changes.

For the dimerization $2A \rightleftharpoons A_2$, begin with one mole of monomer units and define α as the fraction of units bound into dimers. Then $n_A = 1 - \alpha$, $n_{A_2} = \alpha/2$ and total gas amount is $1 - \alpha/2$ moles on this basis. Partial pressures follow from mole fractions times total pressure. Substitution into the reaction quotient gives

$$Q = \frac{\alpha(2 - \alpha)}{4(1 - \alpha)^2} \, \frac{p^\circ}{p}.$$

The ratio of activities is determined by both composition and total pressure. Compression at the original equilibrium composition immediately lowers Q, making the forward reaction favourable. The subsequent association reduces the number of gas particles while conserving monomer units.

A pressure change at fixed temperature

Suppose $K = 0.750$ at the chosen temperature. Let $y = 1 - \alpha$. Because $\alpha(2 - \alpha) = 1 - y^2$, solving $Q = K$ gives

$$y = [1 + 4Kp/p^\circ]^{-1/2}.$$

At $p = p^\circ$ this gives $\alpha = 0.500$, with one half of all units bound and only one third of gas particles being dimers. At $p = 4p^\circ$ the bound-unit fraction is

$$1 - 1/\sqrt{13} = 0.723.$$

The distinction between fractions of conserved units and fractions of gas molecules matters for interpreting a spectroscopic composition measurement.

A temperature change can also move equilibrium. Differentiating $\ln K = -\Delta_r G^\circ/(RT)$ and using the Gibbs–Helmholtz relation gives

$$d\ln K/dT = \Delta_r H^\circ/(RT^2).$$

For an exothermic association increasing temperature decreases K locally. Pressure changes at fixed temperature leave K itself fixed under the same standard-state convention. It is the composition-dependent quotient that adjusts to the imposed mechanical conditions.

The distinction between a chemical potential per mole and one per particle is essential in this derivation. Here partial pressures were combined with R because the extent was measured in moles. Molecular statistical mechanics instead uses k_B with energies per molecule, but both conventions give the same dimensionless ratio of equilibrium activities. A useful numerical check is to scale every initial amount and the container volume by the same factor at fixed temperature. Mole fractions and partial pressures stay fixed, so the equilibrium fraction is unchanged even though twice as many molecules may react.

Use a membrane to turn chemical equilibrium into pressure

Equality of chemical potential can constrain a species even when other species cannot move. Separate pure solvent at pressure p from a dilute ideal solution at pressure $p + \Pi$ by a membrane permeable only to solvent. Hold temperature fixed and assume a nearly incompressible solvent with approximately constant partial molar volume $\bar{v}_w$. The solvent can exchange across the membrane; solute cannot. Its equilibrium condition is therefore equality of solvent chemical potentials, even though total pressures and solute concentrations differ.

For an ideal solution the solvent chemical potential relative to pure solvent at the same pressure changes by $RT\ln x_w$, where x_w is the solvent mole fraction. Increasing pressure by Π raises the solvent chemical potential by $\int_p^{p+\Pi} \bar{v}_w\,dp' \simeq \bar{v}_w \Pi$. Equating the two sides gives

$$\Pi = -\frac{RT}{\bar{v}_w}\ln x_w \simeq \frac{RT}{\bar{v}_w} x_s \simeq RT c_s.$$

The last approximation assumes a small solute mole fraction x_s and solution volume close to the solvent volume, with c_s in moles of independently dissolved particles per cubic metre. The pressure difference is positive because mixing lowers the chemical potential of solvent. If pressure is released, solvent flows into the solution, modifying its concentration until a new balance is attained.

Estimate a hydrostatic head and a molar mass

At 298 K dissolve 2.00 g/L of a nonvolatile solute with molar mass 0.200 kg/mol. The mass concentration is 2.00 kg/m³, so $c_s = 10.0$ mol/m³. Predicted osmotic pressure is $8.314 \times 298 \times 10.0 = 24.8$ kPa. With solution density approximated by 1000 kg/m³, an open-column experiment would need height difference $\Pi/(\rho g) = 2.53$ m to balance it. The final concentration must be measured after the column rises, since transferred solvent can appreciably dilute a small sample.

Conversely, for the same 2.00 g/L preparation a pressure reading of 12.4 kPa would imply an apparent molar mass $RT(2.00)/12400 = 0.400$ kg/mol. Association of two formula units into one particle is one possible

explanation if the chemical analysis gives 0.200 kg/mol. Nonideal activities, leakage and concentration calibration can also change the inferred particle count. Repeating at lower concentrations tests whether Π/c_s approaches the ideal-gas constant times temperature. In the pure-solvent limit $x_w \to 1$ the chemical-potential difference vanishes so no osmotic pressure is required, despite the membrane still being present.

Concentration units can be a major source of error in this example. One mole per litre is a thousand moles per cubic metre, so the gas constant in joules per mole kelvin gives a pressure in pascals only after this conversion. For fully dissociated dilute electrolytes there may also be more than one independently mobile particle per formula unit. Conversely strong ion pairing or aggregation can reduce the count. Measuring a nonvolatile neutral reference solute of known molar mass in the same apparatus provides a way to test both the pressure scale and the effectiveness of the membrane.

Study the full derivation

Pair correlations, the virial gas and a pressure measurement · Statistical mechanics from microscopic counting to material response

A dilute nonideal gas can display almost ideal pressure while its particles strongly avoid one another at short distance. Connect a two-particle interaction to pressure, internal energy and a measurable pair correlation. Then use the resulting integral equations to ask which features of a potential can actually be inferred from a limited set of pressure readings.

Chapter 12 · Entropy production and relaxation

Derive a measurable rate of entropy production during contact between finite bodies and test it using temperature histories.

FROM FIRST PRINCIPLES

Derive entropy production during heat exchange

Connect two bodies in a thermally isolated enclosure. Assume each body reaches internal equilibrium quickly enough to have a single temperature, and no external work is done. Let $J > 0$ denote energy flowing per time from the hot body at T_h to the cooler one at T_c. Their internal energies change at equal and opposite rates. Using the equilibrium state relation $dS_i = dU_i/T_i$ for each body at fixed volume and composition gives

$$\dot{S}_{\text{total}} = -\frac{J}{T_h} + \frac{J}{T_c} = J\frac{T_h - T_c}{T_h T_c}.$$

This entropy production is positive while heat goes from hot to cold. To predict a relaxation time requires an empirical heat-flow relationship. If the contact has constant conductance G, so $J = G(T_h - T_c)$, and the heat capacities C_h, C_c stay constant, energy balance gives

$$\frac{d(T_h - T_c)}{dt} = -G\left(\frac{1}{C_h} + \frac{1}{C_c}\right)(T_h - T_c).$$

The temperature difference decays exponentially. The sign of the entropy change follows from thermodynamics, while the measured conductance controls how quickly equilibrium is reached.

DERIVATION AND APPLICATIONS

Connect the recovery of a two-state population to entropy production

A system can return to equilibrium exponentially even when its entropy production has a more complicated time dependence. Consider many independent units with two internal configurations A and B in contact with one heat bath at fixed temperature T. Each changes from A to B with rate k_+ and back with rate k_-, both positive and constant. Write p for the fraction in B. Conservation of population gives the net forward rate per unit and its stable fixed point,

$$\dot{p} = k_+(1 - p) - k_- p = -(k_+ + k_-)(p - p_*),$$
$$p_* = \frac{k_+}{k_+ + k_-}.$$

The solution is

$$p(t) = p_* + [p(0) - p_*]\exp[-(k_+ + k_-)t].$$

To relate that motion to thermodynamics, suppose the rates obey detailed balance with the bath. For an independently mixed ensemble the excess free energy per unit relative to its equilibrium population is

$$F(p) - F(p_*) = k_B T \left[p \ln \frac{p}{p_*} + (1-p) \ln \frac{1-p}{1-p_*} \right].$$

This expression combines the change in the two-state energy average with the entropy of distributing the units between A and B. Differentiate with respect to p, multiply by $\dot{p}$, and use $p_*/(1-p_*) = k_+/k_-$. Defining forward and reverse event rates $j_+ = k_+(1-p)$ and $j_- = k_- p$ gives the entropy production per unit,

$$\dot{S}_{\mathrm{total}} = -\frac{\dot{F}}{T} = k_B (j_+ - j_-) \ln(j_+/j_-) \geq 0.$$

The inequality follows because the larger positive flux also has the larger logarithm. At equilibrium the two one-way event rates can both be finite even though their difference and total entropy production vanish. Near equilibrium the thermodynamic driving force and net rate are both linear in $p - p_*$, so their product decays twice as rapidly in its exponential time dependence.

WORKED EXAMPLE

Compute the initial dissipation during recovery

Change the inputs and investigate this derivation.

Set $k_+ = 3.00\,\mathrm{s}^{-1}$, $k_- = 1.00\,\mathrm{s}^{-1}$ and prepare $p = 0.250$. The forward and reverse rates are 2.25 and $0.250\,\mathrm{s}^{-1}$. Initial entropy production is $2.00\,k_B \ln 9\,\mathrm{s}^{-1} = 4.39\,k_B/\mathrm{s}$ per unit. After $0.250\,\mathrm{s}$, one relaxation time, the population has risen to $0.750 - 0.500/e = 0.566$.

SECOND APPLICATION

Infer microscopic rates from a relaxation trace

Change the inputs and investigate this derivation.

A different ensemble reaches equilibrium fraction $p_* = 0.200$. The magnitude of a small population offset halves every $0.500\,\mathrm{s}$. The sum of the rates is $\ln 2/0.500 = 1.386\,\mathrm{s}^{-1}$; hence $k_+ = 0.277\,\mathrm{s}^{-1}$ and $k_- = 1.109\,\mathrm{s}^{-1}$. At equilibrium each one-way flux is $0.222\,\mathrm{s}^{-1}$ per unit. Event counting can test this prediction even when the mean population looks stationary.

Try it · Master entropy

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

Forward rate (1/s)	3
Reverse rate (1/s)	1
Initial upper-state probability	0.25

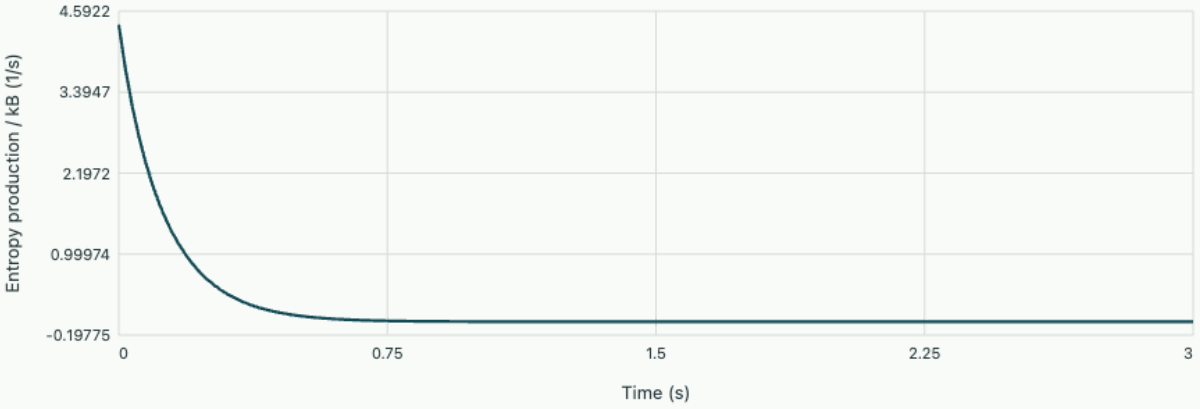

Two-state entropy production

Equilibrium probability 0.75; relaxation time 0.25 s. Production is nonnegative and vanishes at detailed balance.

Start from this chapter. Subtract the two energy balances to get an equation for the temperature difference. Use its exponential decay rate to decide which change of conductance or heat capacity can also change equilibrium.

OneSimulation laboratory · Separate the equilibrium temperature from the relaxation time

Two bodies with no external heat exchange

Connect two bodies through a weak thermal contact inside a perfectly insulated enclosure. Assume each equilibrates internally much faster than they exchange heat, so each has a single instantaneous temperature T_h or T_c. Take their heat capacities C_h, C_c as positive constants over the range, with no work or phase change. Let J denote heat flowing per second from the initially hot body into the initially cold one. Energy conservation gives

$$C_h \dot{T}_h = -J,$$

$$C_c \dot{T}_c = J.$$

This immediately fixes the conserved weighted mean temperature.

A body's entropy change for this model depends only on its successive internal equilibrium states, so $dS_i = C_i\, dT_i/T_i$. The bridge can carry irreversible heat while that state relation remains valid inside each nearly uniform body. Adding both entropy rates gives

$$\dot{S}_{\mathrm{total}} = J(1/T_c - 1/T_h).$$

Positive heat flow down the temperature difference produces positive entropy. Equilibrium sets the thermodynamic force $1/T_c - 1/T_h$ to zero. Thermodynamics constrains the sign of J, yet a constitutive heat-transfer law is still needed to predict how quickly temperatures change.

Identify a rate law experimentally

Suppose $J = G(T_h - T_c)$ with temperature-independent thermal conductance G. Subtracting the two temperature balances shows that $\Delta = T_h - T_c$ satisfies

$$\dot{\Delta} = -G(1/C_h + 1/C_c)\Delta.$$

The decaying difference and conserved weighted mean determine both temperature histories. The entropy-production rate becomes $G\Delta^2/(T_h T_c)$. Near final temperature T_f, the temperature difference decays exponentially while the entropy-production rate decays approximately twice as fast because it depends on the square.

The total entropy increase can also be obtained without integrating this rate. Integrating each state differential from its initial temperature to the common final value gives

$$\Delta S = C_h \ln(T_f/T_{h,0}) + C_c \ln(T_f/T_{c,0}).$$

This expression is independent of conductance so long as the enclosure is isolated and the final state is the same. Larger conductance produces entropy faster at early times but shortens the duration of the process. Replacing a sliding-contact resistance by a smaller one changes the temperature record even when the equilibrium entropy gain is identical.

Predict a quantity beyond the equilibrium temperature

Take $C_h = 500$ J/K, $C_c = 1500$ J/K, $T_{h,0} = 360$ K, $T_{c,0} = 300$ K and $G = 2.00$ W/K. Conservation gives $T_f = 315$ K. The initial heat current is 120 W, and $\dot{S}(0) = 120(1/300 - 1/360) = 0.0667$ W/K. The relaxation time is 187.5 s. The net entropy gain is $500 \ln(315/360) + 1500 \ln(315/300) = 6.42$ J/K to two decimal places.

Use the two-body temperature histories to estimate the initial entropy-production rate and sum changes in $C_i \ln T_i$ over the record. Doubling the conductance should double the initial rate and halve the measured

relaxation time. Extending both runs to the same final tolerance should recover the same total entropy gain. A discrepancy that decreases as the sampling interval shrinks points to time resolution in the estimated initial slope.

At a finite recorded time the entropy increase is smaller than its asymptotic endpoint value. Evaluate the state-function expression at the last recorded temperatures before attributing a shortfall to the numerical algorithm. The remaining difference measures entropy still to be produced during subsequent equilibration.

Couple heat transport to electrical current

A thermal bridge can also carry electric charge. In a simple thermoelectric generator a temperature difference produces an electromotive force through the Seebeck effect, while current carries Peltier heat at each end. Take an idealized device with constant effective Seebeck coefficient S, internal resistance R, open-circuit thermal conductance K and no additional contact losses. Let $T_h > T_c$ be externally fixed endpoint temperatures and choose positive current I as useful generator output. Connect a passive load R_L to the device.

The open-circuit voltage is $S(T_h - T_c)$. Current through the load is $I = S\Delta T/(R + R_L)$, giving useful power $P = I^2 R_L = IS\Delta T - I^2 R$. For a uniform element with constant properties, half its internal Joule heat flows to each end. The reversible heat carried per unit charge is ST by the Kelvin relation. Combining these contributions with Fourier conduction gives hot-side input and cold-side output,

$$\dot{Q}_h = ST_h I + K\Delta T - \tfrac{1}{2} I^2 R,$$

$$\dot{Q}_c = ST_c I + K\Delta T + \tfrac{1}{2} I^2 R.$$

Their difference is exactly the electrical output. With the sign conventions fixed, the total reservoir entropy production is $\dot{Q}_c/T_c - \dot{Q}_h/T_h$. The Peltier contributions cancel, leaving a positive term from conduction and another from Joule dissipation. This explicitly separates a reversible cross effect from two irreversible processes in the same device.

Find a matched-load operating point

Use $T_h = 350$ K, $T_c = 300$ K, $S = 200$ µV/K, $R = 0.100$ Ω and $K = 0.00200$ W/K. The open-circuit voltage is 0.0100 V. With matched load $R_L = 0.100$ Ω, current is 0.0500 A and useful power 0.250 mW. Conductive heat leakage alone is 0.100 W. At the hot end the Peltier contribution is 3.50 mW and half the internal Joule heating is 0.125 mW. Thus $\dot{Q}_h = 103.375$ mW and $\dot{Q}_c = 103.125$ mW. Their measured difference is much smaller than either heat current.

The operating efficiency $P/\dot{Q}_h = 0.00242$ is far below the reversible limit $1 - 300/350 = 0.143$. Maximizing $P = (S\Delta T)^2 R_L/(R + R_L)^2$ at fixed endpoint temperatures indeed gives $R_L = R$, but this choice need not maximize efficiency because the heat input also depends on current. With the circuit open the load receives zero power while $K\Delta T$ still crosses the element and produces entropy. A reliable laboratory test therefore records electrical power and both thermal flows separately instead of interpreting an open-circuit voltage as evidence of appreciable useful energy conversion.

The numerical heat currents allow a direct entropy audit. The reservoirs together gain $0.103125/300 - 0.103375/350 = 4.84 \times 10^{-5}$ W/K. Most of this is already explained by bare heat conduction, since $0.100(1/300 - 1/350) = 4.76 \times 10^{-5}$ W/K. The small remainder is the Joule contribution deposited at the two ends. This calculation uses reservoir rates with signs for heat leaving the hot bath and entering the cold bath. A sign error can suggest negative entropy production when the parameters instead describe an ordinary inefficient generator.

Recording the cold-end heat flow with its sign explicitly marked avoids the common confusion between heat entering the cold reservoir and heat entering the device.

Account for entropy during isolated thermal exchange

Two equal, temperature-independent heat capacities begin at 350 K and 250 K and exchange heat through a constant conductance. Energy fixes their common final temperature at 300 K. The temperature difference decays exponentially and each entropy change is $C \ln(T/T_i)$.

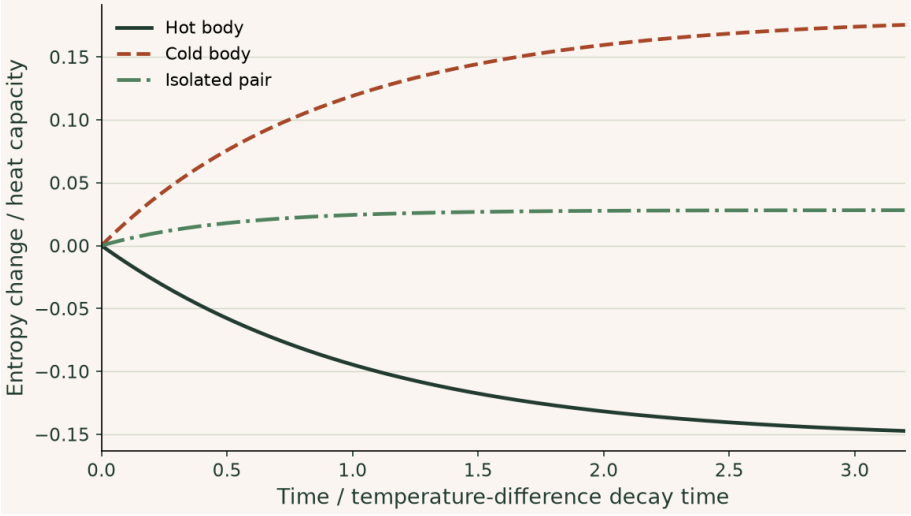

The hot body loses entropy while the isolated pair gains it. At equilibrium the total increase divided by one body's heat capacity is $\ln(90000/87500) \approx 0.0282$. Energy conservation and entropy production constrain different quantities.

Try it · Thermal entropy

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

Initial hot-body temperature (K)	350
Conductance (W/K)	4
Heat capacity of each body (J/K)	500

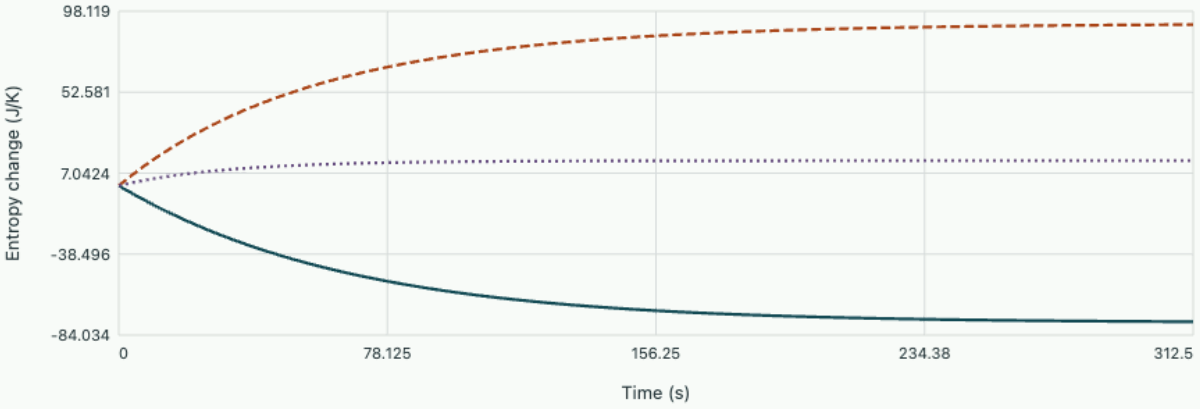

Equilibrium temperature 300 K; entropy increase 14.085 J/K. Two equal heat capacities, constant conductance and no external losses.

▶ **Read values from the plot**

Study the full derivation

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.

Problem sheet 6 · Equilibrium, kinetics and irreversibility

Separate a chemical equilibrium constant from a thermal relaxation coefficient and quantify entropy changes in a finite system.

Change the gas volume

At fixed temperature an ideal dimerization $2A \rightleftharpoons A_2$ has $K = 2.00$. At total pressure p° find the bound-unit fraction and dimer mole fraction. Immediately expand isothermally until all original partial pressures have halved. Find the instantaneous reaction quotient and the preferred direction.

▼ **Show the composition and direction**

With bound-unit fraction α, equilibrium gives

$$1 - \alpha = (1 + 4Kp/p^\circ)^{-1/2} = 1/3.$$

Hence $\alpha = 2/3$ and the dimer mole fraction $\alpha/(2-\alpha) = 1/2$. Because $Q = a_{A_2}/a_A^2$, halving both partial pressures doubles Q from 2 to 4 before chemical readjustment. Thus $\Delta_r G = RT\ln(Q/K) > 0$, favouring dissociation. Continuing mechanical readjustment depends on whether volume or pressure is subsequently held fixed.

Two equal heat capacities

Two insulated blocks each have $C = 200$ J/K and start at 320 and 280 K. Find the common temperature, total entropy gain and initial entropy-production rate for conductance 1.50 W/K. What is their temperature difference after 50.0 s?

▼ **Show the balance and rate**

The final temperature is 300 K. Entropy gain is

$$200\ln[300^2/(320\times 280)] = 0.891\,\mathrm{J/K}.$$

Initial current is $1.50(40) = 60.0$ W and initial entropy production is $60(1/280 - 1/320) = 0.0268$ W/K. With equal capacities $\tau = C/(2G) = 66.7$ s. The difference after 50.0 s is

$$40e^{-0.75} = 18.9\,\mathrm{K}.$$

The final entropy gain depends only on endpoints, while initial rate and time history depend on conductance.

Determine whether a drift is heat leakage

A trace from two strongly insulated bodies has equal capacities of 150 J/K. Initially the measured temperatures are 330 and 290 K. After the difference has vanished both read 309 K. Estimate the apparent

net heat exchange with the surroundings. Can increasing internal contact conductance explain it? Give an independent instrumental diagnosis.

▼ **Show the energy inventory**

For equal capacities the initial arithmetic mean is 310 K and must stay fixed under pure internal heat exchange. Final readings imply $\Delta U = 150(309 - 330) + 150(309 - 290) = -300$ J. Finite contact conductance can change the approach rate, but its internal current cancels in the total energy balance. A common thermometer drift of -1 K between beginning and end would create the same apparent loss. Calibration against an independent temperature reference, or an enclosure heat-flux measurement, distinguishes these possibilities.

Extract work from a finite hot body

A body of constant heat capacity 1.00 kJ/K cools from 500 K to 300 K. A large reservoir stays at 300 K. Find the maximum work recoverable, the heat delivered to the reservoir and the entropy generated by direct thermal contact. Then predict how much of the available work remains if cooling stops at 400 K.

▼ **Worked solution**

For a reversible sequence of small engines the instantaneous efficiency is $1 - T_0/T$, where T is the body's current temperature. The body supplies $-C\,dT$ on cooling, so $W_{max} = C[(T_i - T_0) - T_0 \ln(T_i/T_0)] = 46.75$ kJ. Applying the initial Carnot efficiency to the entire 200 kJ would overestimate the work because the hot-side temperature falls.

The remaining 153.25 kJ enters the reservoir. Its entropy gain is 153.25/300 = 0.5108 kJ/K, equal to the magnitude of $C \ln(300/500)$ lost by the body. With direct contact, the reservoir instead receives all 200 kJ and total entropy production is 0.6667 − 0.5108 = 0.1558 kJ/K. Multiplication by 300 K recovers the 46.75 kJ of lost work.

At 400 K the body still has $C[100 - 300 \ln(4/3)] = 13.70$ kJ available relative to the reservoir. The reversible work already extracted is therefore 33.05 kJ. Additivity of these two stages provides a check on both the logarithm and the choice of reference temperature.

Explore these relations

- For thermal relaxation and conservation, select equal capacities and inspect the common final temperature and exponential difference. The live temperatures start at 360 K and 300 K. Open the live calculation.
- For chemical equilibrium, vary reaction enthalpy and compare temperature dependence of the equilibrium constant; the separate pressure dependence in this sheet follows mass action. Open the live calculation.

Identify a heat capacity and a bridge conductance

Two insulated blocks start at 360 and 300 K. The cooler block has heat capacity 1500 J/K. The capacity-weighted equilibrium temperature is 320 K, and a fitted temperature difference decays by a factor e in 125 s. Infer the unknown hot-block heat capacity and the conductance. Predict the two observables after doubling that conductance in the downloadable model.

▼ Work through the solution

Energy conservation gives $C_h(360-320)=1500(320-300)$, hence $C_h=750$ J/K. Subtracting the two block balances gives

$$1/\tau = G(1/C_h + 1/C_c).$$

Therefore $G=[125(1/750+1/1500)]^{-1}=4.00$ W/K. Doubling G to 8.00 W/K gives 62.5 s and leaves the final temperature at 320 K. Equilibrium data alone would identify the heat-capacity ratio while leaving the conductance undetermined.

Use the OneSimulation laboratory to check your answer

Set the hot capacity to 750 J/K, cold capacity to 1500 J/K and conductance to 4 W/K. Check 320 K and 125 s, then double the conductance to 8 W/K. Open the corresponding live laboratory reference.

Chapter 13 · How a finite heat bath changes a fluctuation

A tiny sample exchanges heat with its mounting plate inside an insulated chamber. A thermometer samples the same mean temperature whether the plate is large or small. Can the time record of the sample energy reveal the size of the plate?

Count the states compatible with one reading

Call the two bodies A and B. Fix total energy U, volume and particle numbers. Assume weak thermal contact so energies add, although exchange occurs over the observation time. Divide energy into small equal bins wide enough to contain many microstates. If the sample has energy E, the plate must have $U - E$. Under the microcanonical equal-weight hypothesis the probability of this bin is proportional to the product of the two state counts. Introduce $S = k_B \ln \Omega$. Maximising probability is now equivalent to maximising a sum.

$$k_B \ln P(E) = S_A(E) + S_B(U - E) + \text{constant}.$$

Moving energy from B to A increases the first entropy with slope $1/T_A$ and decreases the second with slope $1/T_B$. At an interior maximum their temperatures therefore agree. This argument determines the position of the peak. Its width requires one more derivative. At fixed volume $dT/dE = 1/C$, so differentiating $1/T$ with respect to energy gives $-1/(T^2 C)$. Differentiating the plate's entropy twice brings two factors of minus one from its argument. Its curvature consequently enters with the same sign as the sample curvature.

$$\ln \frac{P(E_* + \delta E)}{P(E_*)} \simeq -\frac{(\delta E)^2}{2 k_B T^2} \left(\frac{1}{C_A} + \frac{1}{C_B} \right).$$

For positive heat capacities the peak is stable. Match this expression to the exponent of a Gaussian to identify $\sigma_E^2 = k_B T^2 C_A C_B / (C_A + C_B)$. The reciprocal capacities add because heating A also cools B. A genuinely macroscopic plate changes temperature so little that its reciprocal capacity disappears. The remaining formula then agrees with the canonical variance.

Background · Heat capacity

$$C_V = \text{var}(E)/(k_B T^2)$$

This fluctuation formula uses a canonical equilibrium distribution at fixed volume and a temperature-independent energy operator. A finite reservoir suppresses subsystem fluctuations.

Two calorimeters at the same temperature

Measure energies in $k_B T$ and capacities in k_B. Let A have capacity 60. An effectively infinite reservoir gives a standard deviation $\sqrt{60} = 7.75$. A plate with equal capacity gives

$$\sqrt{30} = 5.48.$$

A plate with capacity 20 gives

$$\sqrt{15} = 3.87.$$

Doubling the width would quadruple the variance. These distinctions matter if an energy histogram is used to estimate heat capacity. In the equal-capacity case, applying the canonical formula would yield 30 instead of the actual 60.

A second consequence concerns correlations. Within this isolated system a fluctuation of $+3k_BT$ in A forces a change of $-3k_BT$ in B. Their covariance equals $-\sigma_E^2$. The temperatures consequently fluctuate in opposite directions by amounts inversely proportional to each body's capacity. Repeated measurements must be spaced by the relaxation time if their scatter is to estimate an equilibrium variance with independent samples.

Compare two normalised histograms

Keep the sample capacity fixed while reducing the reservoir capacity. A narrower curve must grow taller to retain unit integrated probability. Locate the displacement at which its height is about 61 percent of its maximum. This displacement measures one standard deviation directly in the horizontal units. The quadratic expansion models the central peak. Energy bounds, strong coupling or a heat capacity varying substantially across the peak require the original entropy functions.

Open this experiment online

Try the calculation

Sample capacity C_a/k_β **60**

Plate capacity C_β/k_β **60**

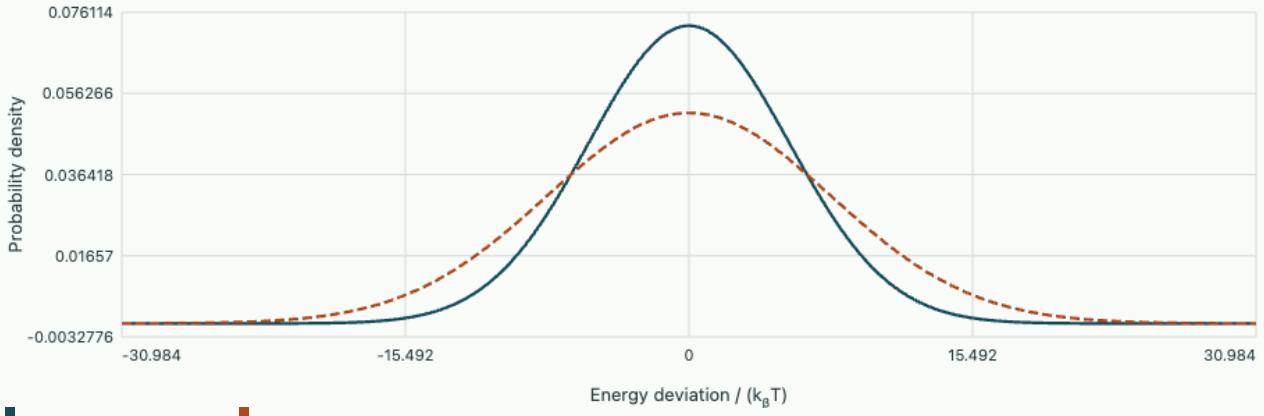

Standard deviations are 5.477 $k_\beta T$ for the finite plate and 7.746 $k_\beta T$ for the infinite bath.

Tutorial on inferring a heat capacity

A sample with heat capacity $90k_B$ has a measured standard deviation of $6k_BT$. Infer the plate capacity and predict the scatter when the original plate is replaced with two identical plates in thermal contact.

▼ **Work through the solution**

The measured variance in thermal units is 36. Write $1/36 = 1/90 + 1/c_B$ to obtain $c_B = 60$. Two plates give 120. The new variance is $90(120)/210 = 51.43$, so the standard deviation is $7.17 k_B T$. Increasing the plate capacity moves the result towards the canonical limit $\sqrt{90} k_B T$.

A calorimeter measures both bodies simultaneously. Its inferred temperature deviations have slope $\delta T_B / \delta T_A = -3$. Combined with an energy variance $15 (k_B T)^2$, determine both capacities.

▼ **Work through the solution**

Energy conservation gives $\delta T_B / \delta T_A = -C_A / C_B$, hence $C_A = 3 C_B$. The effective dimensionless capacity is $3 c_B / 4 = 15$. Thus $C_B = 20 k_B$ and $C_A = 60 k_B$. A covariance measurement and a histogram together identify both parameters where a width alone would give only their harmonic combination.

How would halving both heat capacities change their energy scatter and their temperature scatter at unchanged T?

▼ **Work through the solution**

The effective capacity scales by one half, so σ_E decreases by $1/\sqrt{2}$. Each temperature scatter is σ_E / C and increases by $\sqrt{2}$. Smaller thermal systems can exhibit smaller absolute energy excursions and larger temperature excursions simultaneously.

Further reading · Cambridge statistical physics, fundamentals and ensembles · Fließbach, Statistische Physik.

Study the full derivation

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.

Chapter 14 · Counting occupied sites on a surface

A surface with M identical sites exchanges particles with a gas. As the gas chemical potential rises, the coverage grows and eventually levels off. The same measurement also predicts the magnitude of number fluctuations, without estimating a time-dependent sticking rate.

Start with one site and a reservoir

An empty site has energy zero and population zero. An occupied site has energy $-\epsilon$ with $\epsilon > 0$, and population one. Assume each site can bind at most one particle, the sites do not interact, and both states have unit degeneracy. A reservoir that supplies heat and particles assigns a state weight proportional to $\exp[-\beta(E - \mu N)]$, where $\beta = 1/(k_B T)$. One may recover the combination $E - \mu N$ by expanding the reservoir entropy after removing energy E and particles N, using $T\,dS = dU - \mu\,dN$ at fixed volume.

$$\Xi_1 = 1 + e^{\beta(\mu+\epsilon)},$$

$$\theta = P(1) = \frac{1}{1 + e^{-\beta(\mu+\epsilon)}}.$$

At $\mu = -\epsilon$ empty and occupied are equally likely. A more negative chemical potential penalises bringing a particle out of the gas. Positive binding energy counteracts that penalty. The reference for both energies must be the same. For an ideal dilute gas at fixed temperature, changing the pressure from p_1 to p_2 changes the chemical potential by $k_B T \ln(p_2/p_1)$. Thus doubling the pressure doubles the odds $\theta/(1-\theta)$, even when the coverage itself grows only slightly.

Factorisation explains the noise

With independent sites every configuration has a product of one-site weights, so summing over all binary choices gives

$$\Xi = \Xi_1^{M}.$$

Alternatively group configurations by population N. There are $\binom{M}{N}$ ways to choose the occupied sites. Division by Ξ gives the binomial probability $P(N) = \binom{M}{N}\theta^N(1-\theta)^{M-N}$. Since a binary variable satisfies $n^2 = n$, its variance is $\theta - \theta^2$. Independent sites have zero cross-covariance. Adding their variances gives

$$\langle N \rangle = M\theta,$$

$$\operatorname{var} N = M\theta(1-\theta) = k_B T \left(\frac{\partial \langle N \rangle}{\partial \mu} \right)_T.$$

For $M = 100$ at half coverage the mean is 50 and the standard deviation is 5 particles. Raising the odds to 9 gives mean 90 and deviation 3. On the steep part of the isotherm small changes of reservoir conditions produce the largest response. Thermal noise is largest there for precisely the same reason. A broad distribution near a sharp steepening can also signal collective behaviour, but the independent-site prediction provides a quantitative comparison before invoking interactions.

Measure a shift in binding energy

Imagine comparing two surfaces in the same chamber. Their temperature and chemical potential agree. If the second binds particles more strongly by $k_B T \ln 4$, it multiplies the odds by four. Where the first has coverage $1/5$, its odds are $1/4$, so the second reaches exactly one half. Move the binding-energy slider and identify the chemical potential at the centre of the curve. The plot also displays the number variance per site on the same dimensionless scale. At very small coverage the two curves nearly coincide because θ^2 becomes negligible.

Open this experiment online

Try the calculation

Binding energy ε/(k$_β$T) 2

Chemical potential μ/(k$_β$T) -2

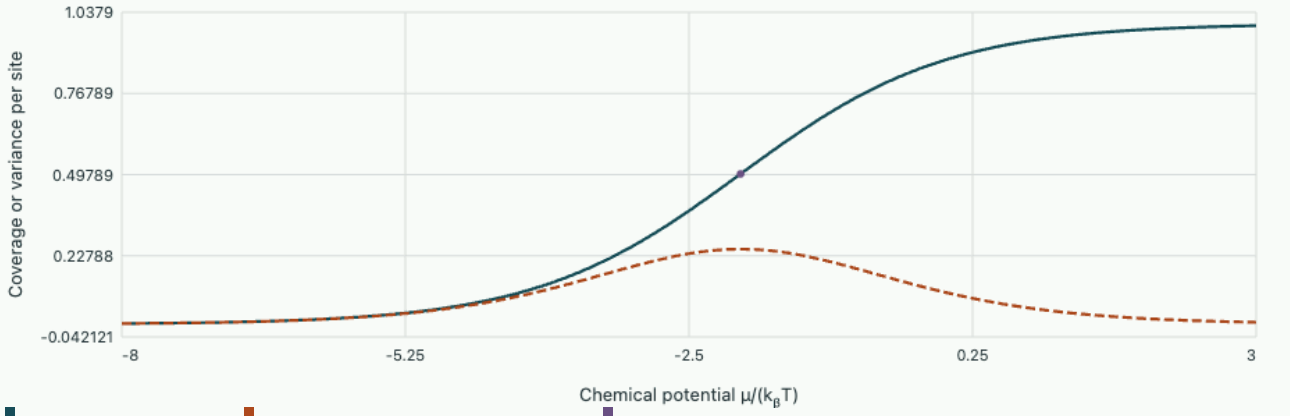

At this chemical potential, 100 sites have mean occupancy 50 particles and standard deviation 5 particles.

Tutorial on recovering an isotherm

At pressure p_0 a surface has mean coverage 0.75. How much must pressure change for coverage 0.90 at the same temperature? For 400 sites, calculate the initial and final root-mean-square fluctuations in coverage.

▼ **Work through the solution**

The odds change from 3 to 9, requiring $p = 3p_0$. The variance of the fractional coverage N/M is $\theta(1-\theta)/M$. Taking square roots gives 0.0217 initially and 0.0150 finally. These numbers measure instantaneous equilibrium fluctuations. The uncertainty of a time average additionally depends on how long successive readings remain correlated.

Infer two independent predictions from the probability of a completely empty three-site patch at half coverage. What fraction of frames show at least two occupied sites?

▼ **Work through the solution**

Empty probability is

$$(1/2)^3 = 1/8.$$

Two occupied sites have probability $\binom{3}{2}/8 = 3/8$, and three have probability $1/8$, giving one half for at least two. Non-binomial counts at fixed mean are a direct test of equivalence and independence of the sites.

Suppose every occupied site admits two internal states of equal energy while an empty site has only one. Determine the shift of the half-coverage chemical potential.

▼ **Work through the solution**

The one-site grand sum becomes $1 + 2e^{\beta(\mu+\epsilon)}$. Half coverage requires the occupied and empty weights to agree, so $\mu = -\epsilon - k_B T \ln 2$. Internal entropy shifts the measured midpoint without changing the exclusion of a second particle.

Further reading · Cambridge statistical physics, grand canonical ensemble and quantum gases.

Study the full derivation

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.

Chapter 15 · How a pair potential changes the gas pressure

A low-density gas can measure its own intermolecular forces. At a prescribed temperature the first deviation from ideal behaviour is a volume. Its sign tells us whether overlap exclusion or attraction makes the larger contribution.

Separate independent pairs from correlated pairs

Let N classical particles occupy volume V, with a total potential energy equal to a sum of short-range pair potentials $u(r_{ij})$. Momentum integrals give the ideal-gas factors. In the remaining positional integral write each pair's Boltzmann weight as $e^{-\beta u} = 1 + f(r)$. The auxiliary function $f = e^{-\beta u} - 1$ vanishes where particles do not interact. In a large vessel the positional integral with one f is $V^{N-1} \int f(\boldsymbol{r})\, d^3 r$. One particle of the pair may be placed anywhere and its neighbour's relative coordinate is sampled against f. There are $N(N-1)/2$ possible unordered pairs.

In the low-density expansion of $\ln Z$, independent repetitions of separate pairs cancel against products from the logarithm. The first connected correction is $N(N-1)\int f/(2V)$. This is a density expansion of the free energy per particle. The total correction need not be small for a macroscopic sample. Differentiating with $p = k_B T (\partial \ln Z / \partial V)_{T,N}$ supplies a minus sign from the reciprocal volume. In the thermodynamic limit with number density $\rho = N/V$,

$$\frac{p}{\rho k_B T} = 1 + B_2(T)\rho + O(\rho^2),$$

$$B_2 = -\tfrac{1}{2}\int f\, d^3 r = -2\pi \int_0^\infty f(r) r^2\, dr.$$

An impenetrable core gives $f = -1$. An attractive annulus gives $f > 0$. The sign of B_2 follows from their integrated volumes, including the spherical shell factor r^2.

Integrate a square well

Set $u = \infty$ for $r < d$, $u = -\epsilon$ for $d < r < \lambda d$, and $u = 0$ outside. The two nonzero contributions are constants integrated over spherical volumes. With $B_{\mathrm{core}} = 2\pi d^3/3$, obtain

$$B_2 = B_{\mathrm{core}}\{1 - (\lambda^3 - 1)[e^{\epsilon/(k_B T)} - 1]\}.$$

A sphere of diameter d has actual volume $\pi d^3/6$. The hard-core coefficient is four times that volume. The forbidden positions concern the centre of a second particle, and the factor of one half avoids counting an interacting pair twice. For $\lambda = 1.5$ and $k_B T/\epsilon = 2$, the ratio $B_2/B_{\mathrm{core}} = 1 - 2.375(e^{0.5} - 1) \simeq -0.541$. A dilute sample with $\rho B_{\mathrm{core}} = 0.02$ consequently has $p/(\rho k_B T) \simeq 0.9892$ at this order.

Locate the temperature at which the slope vanishes

Set $B_2 = 0$ and solve for the Boyle temperature. The result is $k_B T_B/\epsilon = 1/\ln[1 + 1/(\lambda^3 - 1)]$, or 2.846 for the preceding range. Vary the width of the well. A larger attraction volume can offset the hard core even when the temperature is high. At T_B the pressure curve has zero first density correction. Three-particle

contributions can still produce curvature. At very low temperatures pair binding can become strong enough that even a superficially low density fails to justify truncation.

Open this experiment online

Try the calculation

Outer radius / core diameter 1.5

Reduced temperature $k_\beta T/\varepsilon$ 2

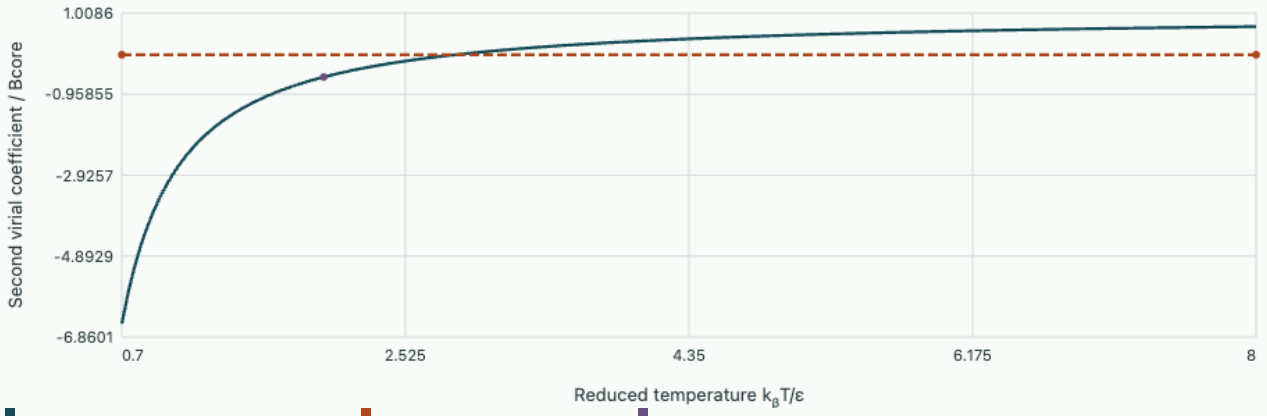

Hard core plus attraction | Zero density slope | Selected temperature

B_2/Bcore = -0.5407. The Boyle temperature is 2.846 ε/k_β.

Tutorial on inferring the range

In a square-well model the measured Boyle temperature is

$$k_B T_B/\epsilon = 2.$$

Determine λ, then predict the sign of the pressure correction when the temperature is doubled.

▼ **Work through the solution**

Rearrangement gives $\lambda^3 = 1 + 1/(e^{1/2} - 1) = 2.5415$, hence $\lambda \simeq 1.365$. Doubling the temperature decreases the attractive Boltzmann factor and leaves the hard-core term fixed. Substitution at reduced temperature 4 gives

$$B_2/B_{\text{core}} \simeq 1 - 1.5415(0.2840) = 0.562.$$

At sufficiently low density the pressure exceeds the ideal prediction.

Can a negative B_2 at one temperature fix how much of the molecular attraction comes from its depth and how much from its range?

▼ **Work through the solution**

The coefficient in the square-well model depends on the product $(\lambda^3 - 1)[e^{\epsilon/k_B T} - 1]$. Many combinations of width and depth give the same product. Measurements at two temperatures can constrain both parameters if the core diameter is independently known and the model is appropriate. Additional temperatures expose inadequacies in the assumed well shape.

Take two measured compressibility factors $Z = 0.9950, 0.9900$ at $\rho B_{\mathrm{core}} = 0.010, 0.020$ at the same temperature. Estimate B_2/B_{core} and check consistency of the truncation with these two measurements.

▼ **Work through the solution**

The slopes from the ideal limit are $(Z-1)/(\rho B_{\mathrm{core}}) = -0.50$ for both points. Their common value agrees with a leading pair correction and shows no detectable curvature at the stated precision. It provides no bound on higher-order terms outside that measured density interval.

Further reading · Cambridge statistical physics, classical gases and the cluster expansion.

Study the full derivation

Pair correlations, the virial gas and a pressure measurement · Statistical mechanics from microscopic counting to material response

A dilute nonideal gas can display almost ideal pressure while its particles strongly avoid one another at short distance. Connect a two-particle interaction to pressure, internal energy and a measurable pair correlation. Then use the resulting integral equations to ask which features of a potential can actually be inferred from a limited set of pressure readings.

Chapter 16 · Why two coexisting phases share a tangent

A material prepared with an intermediate composition can separate into two regions of different compositions. Its total composition is conserved. How can a free-energy graph predict both final compositions and the amount of each phase?

Minimise the free energy of a mixture

Work at fixed temperature and volume with a constant density of equivalent sites. Let $f(c)$ be the Helmholtz free energy per site of a spatially uniform state with fraction c of species B. If the two regions have compositions c_1, c_2 and the second contains fraction x of all sites, conservation requires

$$\bar{c} = (1-x)c_1 + xc_2.$$

Neglect interface energy per site in the macroscopic limit. The mixture free energy per site is then the weighted sum of the two f's, a point on the chord joining them.

Introduce a multiplier μ and minimise $(1-x)f(c_1) + xf(c_2) - \mu[(1-x)c_1 + xc_2 - \bar{c}]$. Varying c_1, c_2 when both regions are present gives

$$f'(c_1) = f'(c_2) = \mu.$$

This derivative measures the chemical work to replace an A particle by B. Varying x supplies the missing condition,

$$f(c_2) - f(c_1) = \mu(c_2 - c_1).$$

The common slope therefore equals the slope of the connecting chord. A straight line must touch the free energy at both endpoints and lie below competing homogeneous states. Equivalently $\int_{c_1}^{c_2} [f'(c) - \mu]\, dc = 0$. This equal-area condition is the composition version of Maxwell's construction. For liquids plotted as pressure versus specific volume, the conjugate slope acquires the sign from $p = -\partial F/\partial V$.

Find the coexistence and stability boundaries separately

Take $f(c) = A[(c-m)^2 - d^2]^2 + \mu_0 c$, with $A > 0$, $m = 0.5$ and $0 < d < 0.5$. Subtracting the linear term tilts all trial chords by the same amount. At $c = m \pm d$ the remaining term vanishes with zero derivative. Hence the common tangent is $\mu_0 c$, and the coexistence compositions are $m - d, m + d$. At $d = 0.30$ a sample of average composition 0.45 separates into 0.20 and 0.80 regions. Conservation gives $x = (0.45 - 0.20)/(0.80 - 0.20) = 5/12$ for the B-rich region.

Now perturb a uniform state into two equal parts at $c \pm \eta$. Their average energy differs from the uniform energy by $f''(c)\eta^2/2 + O(\eta^4)$. Here $f'' = 4A[3(c-m)^2 - d^2]$. Between $m - d/\sqrt{3}$ and $m + d/\sqrt{3}$ infinitesimal fluctuations reduce the energy. Between either of these limits and the nearest coexistence endpoint a uniform state resists small fluctuations but a sufficiently large rearrangement can still lower the energy. This accounts for metastability and for the distinction between spinodal and coexistence boundaries. The uniform model does not predict a nucleation time because interfaces carry no energy in this construction.

Get a phase fraction from a graph

The display subtracts $\mu_0 c$ and measures energies in A. Move the mean composition across the whole interval. Outside the coexistence window the stable composition is homogeneous. Inside, the endpoints stay put while the rich-region fraction changes linearly. Compare $c = 0.45$ to $c = 0.55$ for the same gap. Reflection symmetry exchanges the two phase fractions and leaves the free-energy saving unchanged.

Open this experiment online

Try the calculation

Composition half-gap d **0.3**

Mean B fraction **0.45**

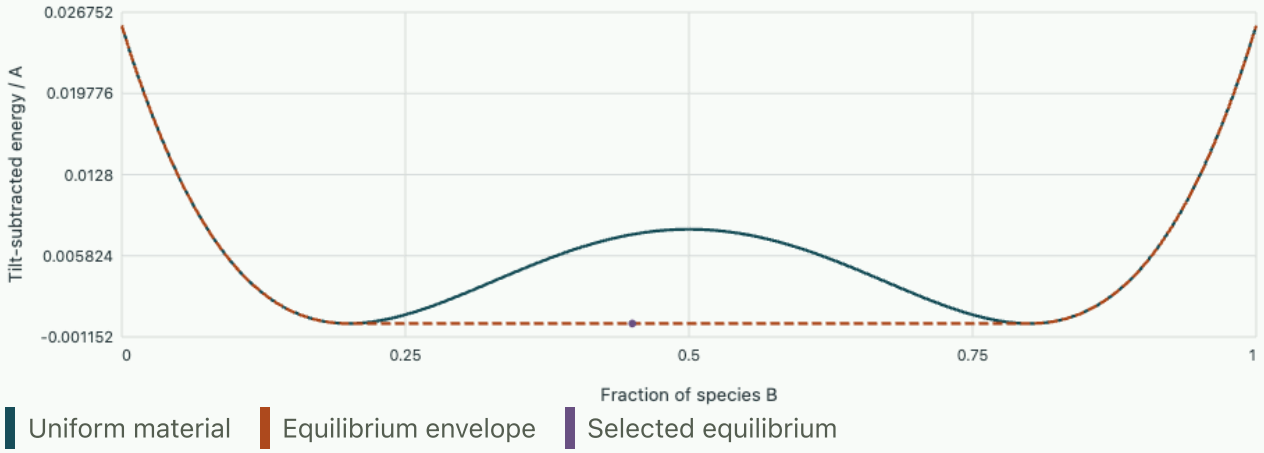

The B-rich phase occupies 41.67% of sites. Energy saving is 0.007656 A per site.

Tutorial on a composition quench

A homogeneous sample is prepared at $c = 0.26$ with $d = 0.30$. Does an infinitesimal composition fluctuation lower its energy? What fraction of the final mixture occupies the B-rich phase in equilibrium?

▼ Work through the solution

The lower spinodal lies at $0.5 - 0.3/\sqrt{3} = 0.327$, so the prepared state lies below it and has positive curvature. Nevertheless 0.26 lies between 0.20 and 0.80. Its equilibrium rich fraction is

$$0.06/0.60 = 0.10.$$

Local stability and the global free-energy minimum answer different experimental questions.

For $m = 0.50$, $d = 0.30$, compare free-energy gains of demixing at mean compositions 0.20 and 0.50 after subtracting the tilt. Report the gain in units of A.

▼ Work through the solution

At a coexistence endpoint the pure homogeneous material already lies on the common tangent and has zero gain. At the midpoint the homogeneous energy exceeds the tangent by $d^4A = 0.0081A$ per site. The equilibrium mixture uses equal amounts of the two compositions. The gain tends continuously to zero at either endpoint.

Would adding $b + mc$ to the free energy change the two selected compositions or spinodals? What would change in the tangent construction?

▼ **Work through the solution**

Conservation fixes the mean composition, so the new term contributes $b + m\bar{c}$ to every candidate mixture at that mean. The minimiser is unchanged. Second derivatives are unchanged, preserving spinodals. Both endpoint tangent slopes, and thus the chemical-potential difference assigned to coexistence, shift by m.

Further reading · Cambridge statistical physics, phase equilibrium and Maxwell construction.

Study the full derivation

Phase nucleation, metastability and waiting-time statistics · Statistical mechanics from microscopic counting to material response

A cooled liquid can remain liquid below its equilibrium freezing temperature and then transform abruptly. How can a bulk phase be thermodynamically favoured while a small embryo of it tends to shrink? Balance interfacial cost against bulk gain, convert a barrier into a probabilistic waiting-time prediction and test which measurements can distinguish surface nucleation from a bulk mechanism.

Chapter 17 · Extracting a length from thermal spin fluctuations

A one-dimensional row of magnetic moments favours local alignment. At low temperature a snapshot can contain a long row of identical spins. An exact calculation distinguishes a large correlated domain from a thermodynamic phase with permanent order.

Sum all configurations by multiplying a small matrix

Consider N classical spins $s_j = \pm 1$ on a ring with $s_{N+1} = s_1$. At zero applied field the energy is $-J \sum_j s_j s_{j+1}$ with $J > 0$. The Boltzmann weight factors into nearest-neighbour terms $e^{K s_j s_{j+1}}$, with $K = J/(k_B T)$. Construct the matrix $T_{ss'} = e^{K s s'}$, whose diagonal elements are e^{K} and off-diagonal elements e^{-K}. Multiplying T's automatically sums each intermediate spin over its two choices. Closing the first and last indices gives

$$Z_N = \operatorname{tr}(T^N).$$

The symmetric and antisymmetric vectors $(1,1)/\sqrt{2}$ and $(1,-1)/\sqrt{2}$ have eigenvalues $\lambda_+ = 2\cosh K$ and $\lambda_- = 2\sinh K$. Thus $Z_N = \lambda_+^N + \lambda_-^N$. At every finite positive temperature the ratio of eigenvalues is strictly below one. In the large-system limit the free energy per spin is $-k_B T \ln(2\cosh K)$. Differentiating with respect to β yields the energy per spin $-J \tanh K$.

Background · Connected correlation

$$C(r) = \langle s_i s_{i+r} \rangle - \langle s_i \rangle \langle s_{i+r} \rangle$$

Subtract the independent contribution before identifying a decay length. A periodic sample introduces a second propagation distance around the ring.

Insert two local measurements

Multiplying an individual configuration by a measured spin corresponds to inserting $S = \operatorname{diag}(1,-1)$ in the matrix trace. To measure two sites separated by r bonds use $\operatorname{tr}(S T^r S T^{N-r})/Z_N$. In the eigenbasis of T, S exchanges the symmetric and antisymmetric vectors. The numerator therefore contains one factor $\lambda_-^r \lambda_+^{N-r}$ and one with the eigenvalues exchanged. With $q = \lambda_-/\lambda_+ = \tanh K$,

$$\langle s_0 s_r \rangle_N = \frac{q^r + q^{N-r}}{1 + q^N},$$

$$\langle s_0 s_r \rangle_\infty = q^r = e^{-ra/\xi},$$

$$\xi = -\frac{a}{\ln q}.$$

Here a is the lattice spacing and the infinite-chain limit holds for fixed r. The second finite-ring term comes from the other way around the circle. At equal sites the full ratio equals one. In zero field spin reversal makes the mean spin zero in the thermal distribution. Thus this correlation already equals the connected correlation.

Compare a chain at two temperatures

At $k_BT = J$, $q = 0.7616$ and $\xi = 3.67a$. Spins ten bonds apart have correlation only 0.0657. At $k_BT = 0.5J$, $\xi = 27.30a$ and the corresponding correlation is 0.693. Lower temperature suppresses individual domain walls. Since $s_j s_{j+1}$ is +1 for alignment and –1 for a wall, its average equals $1 - 2P_{\text{wall}} = q$. At low temperature $P_{\text{wall}} \simeq e^{-2K}$ and $\xi/a \simeq e^{2K}/2$. Even rare walls eventually destroy alignment in a sufficiently long chain.

Compare the finite ring with the infinite-chain result. Change the ring length while leaving the temperature fixed. When the correlation length becomes comparable to the circumference, the long route contributes appreciably. Returning to the starting site must restore correlation one. A straight-line fit to the back half of a semilog finite-ring plot would consequently give a misleading decay length.

Open this experiment online

Try the calculation

Temperature $k_\beta T/J$ **1**

Sites around the ring **32**

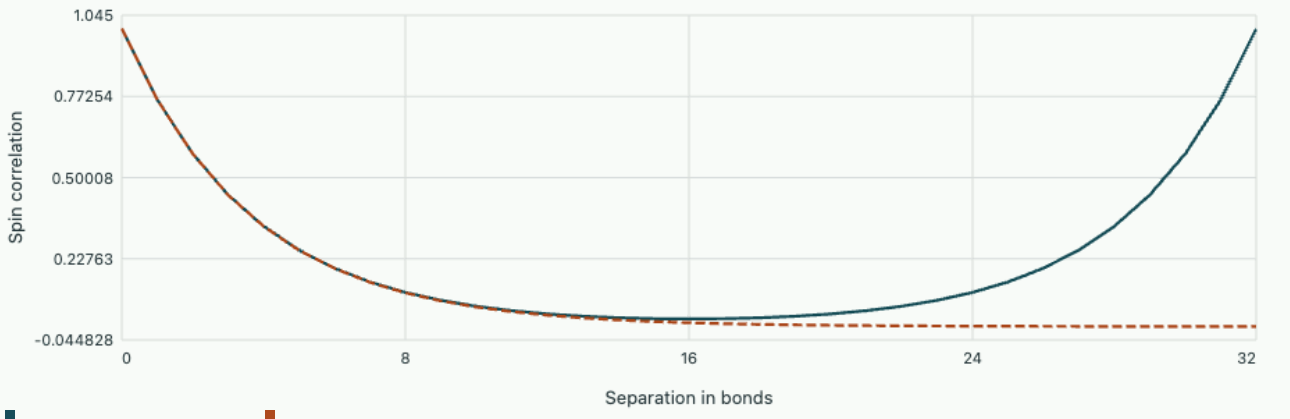

Infinite-chain correlation length 3.672 lattice spacings. The opposite site on the 32-site ring has correlation 0.02562.

Tutorial on estimating an interaction

In a very long zero-field chain the nearest-neighbour correlation is 0.6. Determine $J/(k_BT)$, the probability of a wall on a given bond and the separation at which correlation first falls below 0.1.

▼ **Work through the solution**

Invert $q = \tanh K$ to obtain

$$K = \tfrac{1}{2}\ln[(1+q)/(1-q)] = 0.6931.$$

The wall probability is

$$(1-q)/2 = 0.20.$$

Solve $r > \ln(0.1)/\ln(0.6) = 4.51$. At five bonds the correlation is 0.07776. An even number of walls can restore the original spin direction, so this correlation counts more than the probability of finding a completely wall-free interval.

At $k_B T = J$ a nearest-neighbour bond is aligned with probability p. Find p and its relation to the mean energy per bond.

▼ **Work through the solution**

The correlation is $2p - 1 = \tanh 1 = 0.7616$, so $p = 0.8808$. Multiplying $2p - 1$ by $-J$ gives $-0.7616J$. This directly connects counts of local defects in snapshots with the thermodynamic energy obtained from the partition function.

Why can an odd number of walls never occur on a closed ring? Would dropping the finite-ring correction give the correct wall count at arbitrary temperature?

▼ **Work through the solution**

Each wall reverses the spin when following the ring. Returning to the starting site requires an even number of reversals. The exact nearest-neighbour correlation is $(q + q^{N-1})/(1 + q^N)$. Thus the mean wall count is $N(1 - C_1)/2$. Approximating C_1 by q works when N substantially exceeds the correlation length; otherwise ring closure changes the result.

Further reading · Cambridge statistical physics, Ising chain and correlation functions · Cambridge statistical mechanics, phase transitions and transfer matrices.

Study the full derivation

Finite-size criticality and the evidence for a phase transition · Statistical mechanics from microscopic counting to material response

A finite magnet in zero field can reverse between two equally magnetised directions. Its exact mean magnetisation is then zero even below the temperature where a bulk material orders. Learn how to extract order, correlation length and critical exponents from finite measurements, and which apparent crossings or peaks would survive the thermodynamic limit.

Chapter 18 · Why only a narrow shell of a Fermi gas carries heat

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**

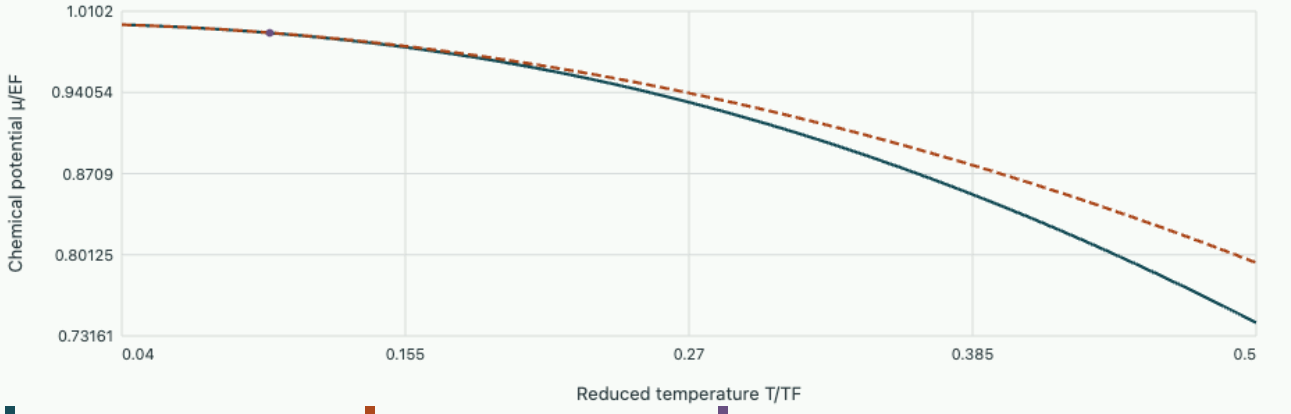

Fixed-number μ/EF = 0.9916, quadratic prediction 0.9918. Leading heat capacity $C_v/(Nk_\beta)$ = 0.4935.

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.

Study the full derivation

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.

Chapter 19 · Mean-field ordering, response and the critical point

Interactions can make a magnet favour a nonzero moment even when the applied field vanishes. Averaging the local exchange field gives a solvable model whose critical predictions can be compared with the exact one-dimensional chain.

Build the free energy from the count of spin states

Take N spins $s_i = \pm 1$, nearest-neighbour energy $-J\sum_{\langle ij\rangle} s_i s_j$ and Zeeman energy $-h\sum_i s_i$, where h has energy units. Denote mean spin by m. The fractions of positive and negative spins are $p_\pm = (1 \pm m)/2$. The binomial count and Stirling's approximation give entropy per spin $-k_B(p_+ \ln p_+ + p_- \ln p_-)$. Independent-neighbour mean field replaces each pair average by m^2. With z neighbours and each bond counted once, interaction energy per spin becomes $-zJm^2/2$.

$$f(m) = -\tfrac{1}{2} zJm^2 - hm + k_B T[p_+ \ln p_+ + p_- \ln p_-].$$

Differentiate at fixed T, h. The entropy contribution has derivative $k_B T \operatorname{atanh} m$, giving $m = \tanh[(zJm + h)/(k_B T)]$. This equation can have three real roots, so an iterative solver returning a root alone is insufficient. The equilibrium phase minimises f; the second derivative $k_B T/(1 - m^2) - zJ$ distinguishes local minima and maxima.

Expand about the loss of stability

Near zero magnetisation,

$$\operatorname{atanh} m = m + m^3/3 + \cdots.$$

Set $T_c = zJ/k_B$ and integrate the expanded derivative to obtain

$$f(m) - f(0) \simeq \tfrac{k_B}{2}(T - T_c)m^2 + \tfrac{k_B T}{12} m^4 - hm.$$

At zero field and $T > T_c$, $m = 0$ is stable. Below T_c, symmetry gives two minima whose leading magnitudes satisfy $m^2 \simeq 3(T_c - T)/T$. Thus the order parameter vanishes like $(T_c - T)^{1/2}$ close to T_c. At $T = T_c$ a small nonzero field instead gives

$$m \simeq [3h/(k_B T_c)]^{1/3}.$$

These powers describe this model's limiting behaviour. For a finite sample symmetry makes the unbiased ensemble average zero if it samples both minima. Spontaneous magnetisation refers to the thermodynamic limit with an infinitesimal field selecting a sign.

Compute a response, then compare with an exact counterexample

Above T_c the small-field stationary equation gives

$$\chi = \partial m/\partial h = 1/[k_B(T - T_c)].$$

Below T_c, differentiate the exact stationarity equation on the positive branch, obtaining $\chi = [k_B T/(1-m^2) - zJ]^{-1}$. Expanding near T_c there gives $\chi \simeq [2k_B(T_c - T)]^{-1}$, half the above-transition amplitude at equal distance. At $T = 0.8T_c$, the full zero-field solution is about $m = 0.7104$; the small-m estimate $\sqrt{3(1-0.8)/0.8} = 0.866$ is already visibly inaccurate.

Mean field predicts an ordering temperature even for a one-dimensional chain with $z = 2$. The exact finite-temperature correlation length in the earlier transfer-matrix solution is finite at all $T > 0$. Spatial fluctuations and domain walls invalidate the independent-neighbour closure in one dimension. The experiment plots the full mean-field free energy and its quartic approximation for the same normalised T, h. Find the minima before inspecting the calculated order parameter. Compare how a weak bias changes their relative depths.

Open this experiment online

Try the calculation

Temperature T/Tc **0.8**

Field h/(k_βTc) **0**

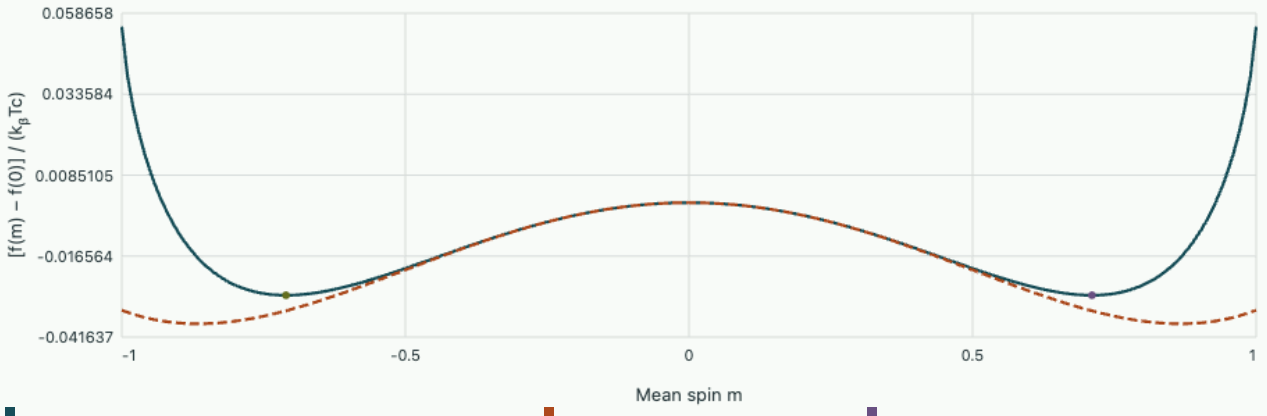

Full independent-neighbour free energy · Quartic approximation · Equilibrium minimum · Symmetry-related minimum

Equilibrium mean spin ±0.7104. Curvature at the minimum 0.6151 in units of k_βTc.

Tutorial on critical response

At $T = 1.10T_c$, what is the linear-response mean spin expected at $h/(k_B T_c) = 0.002$? At the critical point use the cubic equation to estimate m for the same field.

▼ **Work through the solution**

Above the transition $m \simeq 0.002/0.10 = 0.020$. At criticality $m \simeq (0.006)^{1/3} = 0.182$. The larger response signals the loss of the quadratic restoring term; linear extrapolation through T_c would give a divergent susceptibility.

At $h = 0$ and $T = 0.95T_c$, is $m = 0$ a possible output from a fixed-point solver? How should an equilibrium calculation prevent a wrong interpretation?

▼ **Work through the solution**

Starting the symmetric iteration at exactly zero returns zero repeatedly, even below T_c. Its free-energy curvature is $k_B(T - T_c) < 0$; it is a local maximum. Compare the free energies or bracket the two symmetry-related nonzero stationary points before reporting an equilibrium moment.

Further reading · Cambridge statistical physics, mean-field Ising and Landau theory.

Study the full derivation

Finite-size criticality and the evidence for a phase transition · Statistical mechanics from microscopic counting to material response

A finite magnet in zero field can reverse between two equally magnetised directions. Its exact mean magnetisation is then zero even below the temperature where a bulk material orders. Learn how to extract order, correlation length and critical exponents from finite measurements, and which apparent crossings or peaks would survive the thermodynamic limit.

Chapter 20 · Design a thermometer with a countable energy spectrum

How precisely can a single equilibrated degree of freedom measure temperature? Derive the reading distribution from the canonical ensemble, find the gap that makes it most informative, and analyse a detector that records no excitations at all.

Start with how the instrument is read

Bring a two-level probe into weak thermal contact with a large bath at an unknown positive temperature T. Its lower energy is set to zero and its upper energy is $\Delta > 0$; both levels are nondegenerate. Wait for equilibration, then measure which energy level the probe occupies. Repeat with independent equilibrium preparations, either by waiting for full rethermalisation between measurements or by using many identical probes. The data consist of N binary outcomes. They are not a direct measurement of the average energy of an infinitely large sample.

Weak coupling lets us use a fixed calibrated Hamiltonian for the probe while treating the bath as a reservoir. In a canonical ensemble the ratio of excited to ground weights is $e^{-\Delta/(k_B T)}$. Write $x = \Delta/(k_B T)$. Normalising the two weights gives excited probability $p = 1/(1 + e^{x})$, partition function $Z = 1 + e^{-x}$, and mean energy $U = \Delta p$. The arbitrary zero of energy drops out of all three measurable statements about populations, energy differences and temperature sensitivity.

One low-energy observation can occur both in a very cold bath and in a relatively warm one. One upper-state observation likewise gives only partial information. Counting k excitations across N preparations produces a binomial distribution. Its dependence on T enters entirely through $p(T)$. A histogram of measured energies has only two bars, yet the relative heights carry a continuous parameter. Keeping track of the total number of shots is essential: the same 10 per cent excited fraction is a different measurement when it came from ten shots and when it came from ten thousand.

Invert the count and treat the boundaries as part of the problem

The likelihood of a hypothesised temperature is proportional to $p^{k}(1-p)^{N-k}$. At an interior maximum the derivative of its logarithm with respect to p vanishes, giving $\hat{p} = k/N$. For $0 < k < N/2$, inversion yields

$$\hat{T} = \frac{\Delta}{k_B \ln[(N-k)/k]}.$$

This inverse is nonlinear. Doubling the excited fraction does not simply double the temperature, and errors on the two sides of a temperature estimate need not be equal. At $k = 0$ the maximum lies at the limiting temperature $T \to 0^{+}$; it does not prove that the bath has reached absolute zero. At $k/N \geq 1/2$, the likelihood on the positive-temperature model is maximised as $T \to \infty$. A finite sample can cross one half by chance even when the equilibrium bath temperature is positive. Before inferring population inversion check statistical uncertainty, uncorrected readout bias and the assumptions about degeneracy.

For numerical work use $\ln L = -kx - N\ln(1 + e^{-x}) + \text{const}$. Working with log probabilities keeps a sequence of unlikely but possible observations from being rounded to zero. Plotting the likelihood normalised to its largest value on a temperature grid makes its shape visible. This curve is not automatically a probability density over T; a Bayesian posterior would additionally require a prior and normalisation in the chosen temperature coordinate.

Derive the temperature resolution twice

Differentiate the occupation with respect to temperature. Since $dp/dx = -p(1-p)$ and $dx/dT = -x/T$, the thermal response is

$$dp/dT = xp(1-p)/T.$$

The observed fraction has variance $p(1-p)/N$. When its statistical width is small enough that the inverse calibration is locally linear, propagate a small fluctuation through the inverse function:

$$\frac{\sigma_T}{T} \simeq \frac{1}{x\sqrt{Np(1-p)}}.$$

A second derivation identifies an information limit. For one observation $X \in \{0, 1\}$, the logarithmic likelihood derivative is $(X-p)p'/[p(1-p)]$. Squaring and averaging uses $\operatorname{var} X = p(1-p)$ and gives Fisher information $F_T = (p')^2/[p(1-p)]$. Information adds across independent trials. For a locally unbiased regular estimator the Cramér–Rao bound is $\operatorname{var}\hat{T} \geq 1/(NF_T)$. The preceding propagated width equals that bound asymptotically for an interior occupation probability. It does not guarantee the performance of the finite-sample inverse at a boundary.

The same expression has a thermodynamic interpretation. The energy variance is $\Delta^2 p(1-p)$, while $C = \partial U/\partial T = k_B x^2 p(1-p)$. Therefore $F_T = C/(k_B T^2)$. A large equilibrium heat capacity means that the energy histogram moves appreciably for a small temperature change compared with its own thermal width. This identity extends to an equilibrium system with a temperature-independent Hamiltonian when its full energy is measured. Coarsening energy into fewer detector bins, changing the Hamiltonian with T, or using correlations across repeated trials changes what can be inferred from the data.

Open editable notebook Download .ipynb

Where should a designer place the gap?

For fixed bath temperature and number of independent preparations, maximise $x^2 p(1-p) = x^2/[4\cosh^2(x/2)]$. Its logarithmic derivative is $2/x - \tanh(x/2)$. The optimum positive root solves

$$x\tanh(x/2) = 2, \qquad x_{\mathrm{opt}} \simeq 2.39936.$$

The product grows monotonically on the positive axis, so this root is unique. At small x the two measured outcomes occur almost equally often but very weakly depend on temperature. At very large x a rare excitation carries substantial energy yet nearly every trial gives the same result. The maximum balances the thermal slope against the fluctuations and availability of both outcomes. Its location assumes two nondegenerate levels; increasing a level's degeneracy changes both its equilibrium weight and the optimum.

For a probe designed to work near 4.0 K the optimal gap divided by k_B is about 9.60 K. The expected excited fraction is only 0.0832. With 400 independent shots the local bound on relative standard error is about 7.54 per cent, or 0.30 K. The default notebook seed happens to yield 24 excitations and a maximum-likelihood estimate near 3.49 K. A known simulation input and a fitted parameter need not coincide in one finite record.

Change the seed before changing the temperature to observe which scatter is expected from the instrument itself.

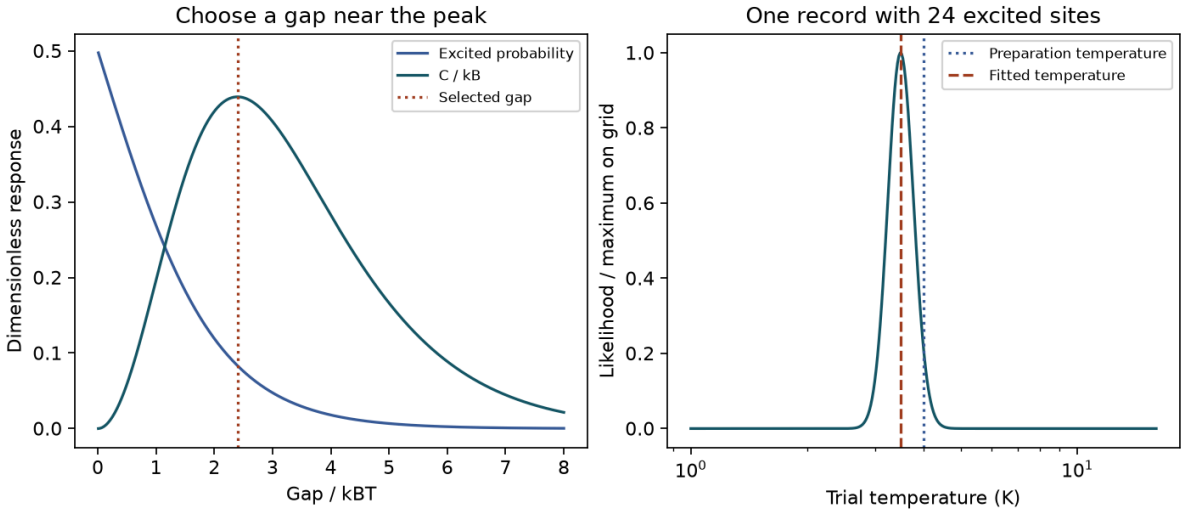

An energy gap near 2.4 times thermal energy maximises local information for a nondegenerate two-level probe. The likelihood on the right comes from one reproducible sample. Its maximum can move even when the bath temperature stays fixed.

Try it · Thermometer

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

Gap / kB (K) **400**

Observed excited count out of 200 **40**

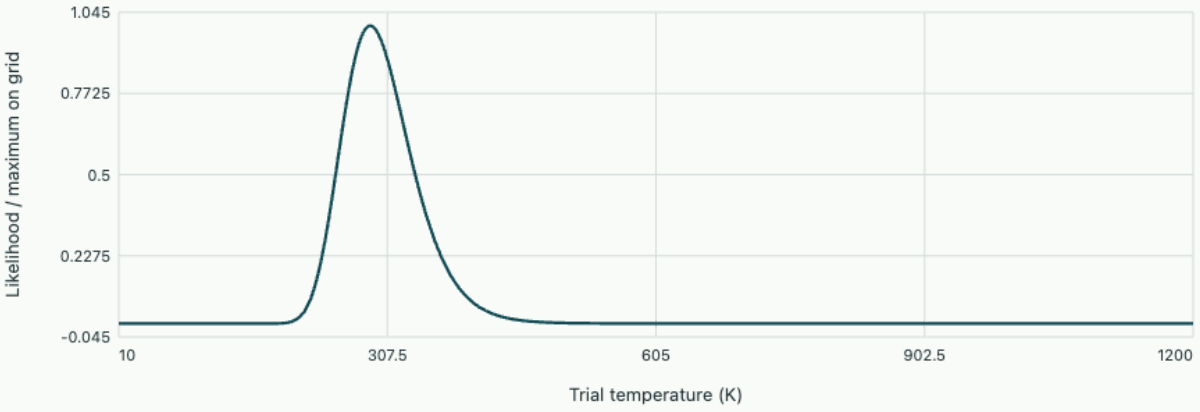

Likelihood for selected record

Positive-temperature maximum 288.54 K. The observed count is adjustable; no fresh random record is silently generated.

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

Use a gap of 0.4 K at the same 4 K bath temperature, then a gap of 40 K. In the first experiment the hot side of a likelihood plot can be extremely broad. In the second, most records are entirely in the ground state. Optimising around one known operating point also has a cost: if temperature may lie anywhere across several orders of magnitude, the best single design locally can be an insensitive probe elsewhere. One strategy is to make a coarse first measurement with one gap, then tune the second gap using that estimate and analyse the two independent batches with their own known calibration settings.

Recover an interval when zero events were observed

For $k = 0$, the exact probability of such a record at a proposed p is $(1 - p)^N$. A one-sided 95 per cent frequentist upper limit p_u solves $(1 - p_u)^N = 0.05$. The transformation to T is monotone on the positive-temperature interval, so invert p_u using the same level gap whenever $p_u < 1/2$. This confidence construction

concerns repeated experiments with the stated coverage, rather than the probability that a particular fixed temperature belongs to an already observed interval.

With 400 shots and zero excitations one finds $p_u \simeq 0.00746$. With $\Delta/k_B = 9.60$ K this implies an upper temperature near 1.96 K. The comparison with the typical record at 4 K shows how a seemingly uninformative blank histogram can rule out part of parameter space. If the detector sometimes misses excitations, however, the same blank record need not imply as cold a bath. Detection efficiency changes the observation law and must be calibrated separately or estimated using more than one preparation.

Tutorial on uncertainty, design and a defective detector

A probe with $\Delta/k_B = 6.0$ K gives 30 excitations in 300 independent preparations. Estimate the temperature and the local relative standard error.

▼ **Work through the solution**

Here $\hat{p} = 0.10$, so $x = \ln 9 = 2.197$ and $\hat{T} \simeq 2.73$ K. Substituting this interior estimate gives $1/[\ln 9\sqrt{300(0.1)(0.9)}] \simeq 0.0876$. The approximate standard error is 0.24 K. The finite-count likelihood is available if asymmetric uncertainty matters for a particular decision.

At the optimal gap how many independent observations are needed for a local relative standard deviation of one per cent? What happens if every recorded observation is simply copied twice?

▼ **Work through the solution**

At the optimum $x^2 p(1-p) \simeq 0.4392$. The design condition is $N \gtrsim 1/[0.01^2(0.4392)]$, approximately 22,800 independent observations. Duplicating each row leaves the experiment and its true uncertainty unchanged. Treating the copied rows as fresh preparations underestimates the standard error by a factor of $\sqrt{2}$.

A readout labels an excited site correctly with probability 0.9 and incorrectly labels a ground-state site as excited with probability 0.02. What quantity must be inverted if 10 per cent of readings say excited?

▼ **Work through the solution**

The recorded probability is

$$q = 0.9p + 0.02(1-p) = 0.02 + 0.88p.$$

From $q = 0.10$ infer $p = 0.0909$ and $x = \ln 10 \simeq 2.303$. Inverting the raw 0.10 instead would give $x = \ln 9$ and a higher inferred temperature for the same positive gap. The derivative of the calibrated observation probability also carries the factor 0.88 when propagating uncertainty.

Further reading · L. A. Correa and collaborators · Individual quantum probes for optimal thermometry

Problem sheet 7 · Can a fluctuation reveal attraction on a surface?

Two surfaces have identical mean coverage at a chosen gas pressure. Repeated microscopy images nevertheless show different event-to-event scatter in the number of occupied sites. Determine whether interactions can account for the excess and design a measurement that distinguishes clustering from instrumental drift.

Predict the independent-site baseline

Each image samples 100 equivalent sites in thermal equilibrium. At half coverage, predict the mean occupied population, its standard deviation and the slope of the mean population with respect to $\beta\mu$. What changes if independent sites have binding energies larger by $k_B T \ln 3$ at fixed μ?

▼ **Work through the solution**

At $\theta = 1/2$, mean population is 50 and standard deviation $\sqrt{100/4} = 5$. The number fluctuation identity gives slope 25. Multiplying the original odds by three gives coverage $3/4$, mean 75, variance $100(3/4)(1/4) = 18.75$ and standard deviation 4.33. A change of mean binding at fixed coverage, compensated by the chemical potential, cannot by itself enhance the independent-site half-coverage variance.

Check the half-coverage variance and the response to a chemical-potential shift

Map an interacting row to spins

Model the surface as a ring of sites with occupancies $n_j = 0, 1$, attraction $-w n_j n_{j+1}$ per occupied pair, and grand-energy contribution $-(\epsilon + \mu) n_j$ per site. Put $n_j = (1 + s_j)/2$. Determine the Ising exchange J and the chemical potential that gives zero Ising field.

▼ **Work through the solution**

The pair product becomes $[1 + s_j + s_{j+1} + s_j s_{j+1}]/4$. On a periodic row each site enters two pair terms. The exchange is $J = w/4$ and the field coefficient in $-H \sum_j s_j$ is $H = (w + \epsilon + \mu)/2$. Choosing $\mu = -\epsilon - w$ gives $H = 0$, and spin reversal then enforces mean occupation $1/2$. Pair attraction can therefore be hidden by a shift of the half-coverage chemical potential.

Use the correlation length to predict number noise

Let $w = 4 k_B T$ at the adjusted half-coverage point. Calculate the correlation at one and ten bonds and the predicted population deviation for 100 sites. Compare to the independent-site prediction.

▼ **Work through the solution**

Here $J/(k_B T) = 1$, $q = \tanh 1 = 0.76159$, and $\langle s_0 s_r \rangle \simeq q^r$, giving 0.7616 and 0.0657. Since $n = (1 + s)/2$, occupancy covariance is $q^r/4$. Summing the covariance over all site pairs for a ring much

longer than $-1/\ln q = 3.67$ gives

$$\operatorname{var} N \simeq (100/4)[1 + 2\sum_{r=1}^{\infty} q^r] = 25(1+q)/(1-q) = 184.7.$$

The deviation is 13.59 sites. Positive spatial covariances account for the extra noise.

Compare a 100-site ring with a shorter ring at the same temperature

Check whether the observation is spatial or temporal

Repeat the experiment with eight sites on a ring. Derive the variance from the finite-ring correlation $C_r = (q^r + q^{8-r})/(1+q^8)$. What measurements would test the interpretation of the excess variance on the larger sample?

▼ **Work through the solution**

Translational invariance gives $\operatorname{var} N = (8/4)\sum_{r=0}^{7} C_r = 2(1+q)(1-q^8)/[(1-q)(1+q^8)]$, approximately 11.77 at $q = \tanh 1$. Report a standard deviation of about 3.43 sites. Measure occupancy covariance against spatial separation and repeat after changing the delay between images. Surface correlations predict a distance-dependent decay within individual images. Temporal drift common to all sites can instead give a substantial long-range component and serially correlated frame means.

Problem sheet 8 · Count atoms in a trap and predict the onset of condensation

A dilute cloud of bosons is confined by three harmonic restoring forces. Cooling moves particles between one-particle states while the total atom number stays constant. Start with counted lattice points in quantum-number space and predict the temperature at which the excited states can no longer hold the whole cloud.

Count levels rather than ordering them one by one

Measure energy from the single-particle ground state. Suppose the angular frequencies are $\omega_x, \omega_y, \omega_z$. For energies far above one quantum, recover the density of levels $g(E)$ by finding the volume of the allowed region in (n_x, n_y, n_z). Check the dimensions. Should one divide by another factorial for indistinguishable atoms?

▼ **Work through the solution**

Each nonnegative integer triple is a distinct one-particle state. Its excitation energy equals $\hbar(\omega_x n_x + \omega_y n_y + \omega_z n_z)$. States below energy E fill a tetrahedron with three intercepts $E/(\hbar\omega_i)$. Its volume is one sixth of their product. Differentiating the integrated count supplies a density per unit energy.

$$G(E) \simeq \frac{E^3}{6\hbar^3 \omega_x \omega_y \omega_z},$$

$$g(E) \simeq \frac{E^2}{2(\hbar\bar{\omega})^3},$$

$$\bar{\omega} = (\omega_x \omega_y \omega_z)^{1/3}.$$

The factor of six came from geometry. Indistinguishability enters later in the count of occupations, through the Bose distribution. Dividing the one-particle density by a factorial of atom number would change the number of available orbitals when one added a second particle. The continuous tetrahedron loses information on its faces and edges, so the estimate is best when $k_B T$ spans many quanta in every direction.

Find how many atoms the excited cloud can hold

The Bose occupation is $1/[\exp((E-\mu)/k_B T) - 1]$. The chemical potential measured from the lowest state approaches zero from below as the population grows. Evaluate the excited-state population at this limit and determine the condensed fraction as a function of temperature. Explicitly account for the state at zero.

▼ **Work through the solution**

The smooth density vanishes at zero and loses the one discrete ground orbital. Retain its population separately, $N = N_0 + N_{\mathrm{ex}}$. Insert the smooth density only for the excited contribution. Put $x = E/(k_B T)$. The integrand contains $x^2/(e^x - 1)$. Expanding the reciprocal as $\sum_{\ell \geq 1} e^{-\ell x}$, integrate each nonnegative term. Its integral is $2/\ell^3$. The result is

$$N_{\mathrm{ex}}^{\max} = \zeta(3)\left(\frac{k_B T}{\hbar\bar{\omega}}\right)^3,$$

$$T_c = \frac{\hbar\bar{\omega}}{k_B}\left(\frac{N}{\zeta(3)}\right)^{1/3}.$$

Below this temperature the excited cloud has the stated capacity and the balance occupies the ground state. Hence

$$N_0/N = 1 - (T/T_c)^3.$$

Above T_c the chemical potential is negative and one must solve a number constraint. Extending the cubic formula above the transition would give a negative condensate fraction. This is an especially clear example of a constraint selecting which branch of a calculation has physical meaning.

A trap calibration as a numerical example

Use $N = 100000$ and cyclic frequencies $f_x = f_y = 80\,\mathrm{Hz}$, $f_z = 160\,\mathrm{Hz}$. Compute T_c, the ideal fraction at 140 nK and the result of halving only the axial frequency. Distinguish the quantities kept fixed in these two comparisons.

▼ **Work through the solution**

Since $\hbar\omega = hf$, the geometric-mean cyclic frequency is $100.79\,\mathrm{Hz}$ and $h\bar{f}/k_B = 4.84\,\mathrm{nK}$. With $\zeta(3) = 1.20206$ the dimensionless temperature multiplier is about 43.65. Thus T_c is about 211 nK and the fraction at 140 nK is about 0.71. Halving one frequency at fixed particle number lowers T_c by $2^{-1/3}$ to roughly 168 nK. At the same absolute temperature the ideal fraction then falls to about 0.42. At a fixed ratio T/T_c both clouds would have the same fraction.

The same equations predict the energy below T_c. Multiplying the occupation integrand by E changes the sum to $6\zeta(4)$. Thus $U = 3\zeta(4)(k_B T)^4/(\hbar\bar{\omega})^3$, measured relative to all atoms in the ground state. Differentiating at fixed trap gives

$$C = 12 k_B \zeta(4)(k_B T/\hbar\bar{\omega})^3.$$

The population and energy probe different moments of the same density of states. A wrong power of energy in the state count changes both predictions.

Test what a continuum approximation misses

Use the isotropic trap experiment below. For frequency f, level n has energy nhf. Obtain its exact degeneracy before comparing the discrete excited-state sum with the continuum. Vary N and the frequency separately. Interpret the crossing of each capacity curve with one.

Open this experiment online

Try the calculation

Atom number **20000**

Isotropic trap frequency (Hz) **100**

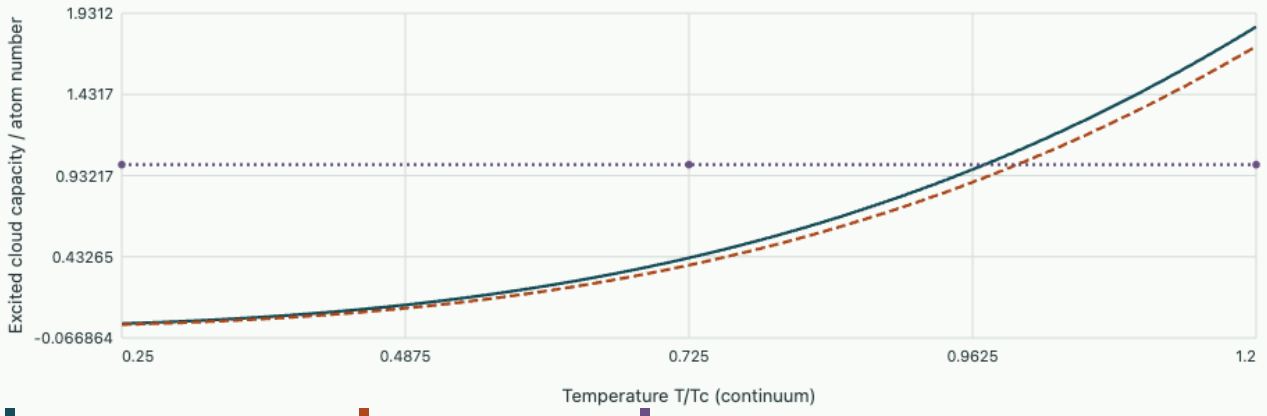

Continuum Tc 122.52 nK. The discrete excited levels reach N at 119.17 nK.

▼ Work through the solution

For each $n_x = 0, \ldots, n$, there are $n - n_x + 1$ choices of n_y. Summing gives $(n+1)(n+2)/2$. The total capacity sums this degeneracy times $1/[\exp(nhf/k_BT) - 1]$ for $n \geq 1$. At a given temperature the faces and edges of the lattice supply positive state-count corrections to the leading volume estimate. Consequently the discrete capacity reaches N at a lower temperature. Increasing N makes the fractional correction smaller. Changing f rescales the temperature axis in kelvin while leaving the discrete curve versus T/T_c unchanged at fixed N. The reported discrete crossing is a saturation convention for a finite ideal sample. A finite set of atoms has a smooth occupation change and interactions introduce further shifts.

Compare the trap with a gas in a box by reading the quantum-gas chapter in David Tong's statistical physics course. Repeating the level count for a box changes the power of energy in the density of states.

Notebook practice · Design a thermometer with a countable energy spectrum

Statistical mechanics from microscopic counting to material response

Count states, construct ensembles and derive macroscopic predictions. Use fluctuations, time dependence and finite sample size to test how well a thermodynamic limit describes a measurement.

1. ## Counting, absolute entropy and the entropy of mixing

   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.

2. ## Ensembles, thermodynamic potentials and measurable curvature

   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.

3. ## Pair correlations, the virial gas and a pressure measurement

   A dilute nonideal gas can display almost ideal pressure while its particles strongly avoid one another at short distance. Connect a two-particle interaction to pressure, internal energy and a measurable pair correlation. Then use the resulting integral equations to ask which features of a potential can actually be inferred from a limited set of pressure readings.

4. ## Rotational and vibrational populations in a molecular gas

   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.

5. ## Bose condensation, dimensionality and a finite trap

   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.

6. ## Degenerate fermions, heat and magnetic 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.

7. Master equations, detailed balance and probability currents

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.

8. From binary collisions to the Boltzmann equation

An initially anisotropic velocity distribution can relax even though every individual collision conserves energy and momentum. Describe a dilute gas with a phase-space density, derive the structure of its collision term and prove which distribution can stop producing entropy. Then examine the assumptions behind that proof and the measurements most likely to reveal their failure.

9. Transport coefficients from equilibrium fluctuations

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.

10. Finite-size criticality and the evidence for a phase transition

A finite magnet in zero field can reverse between two equally magnetised directions. Its exact mean magnetisation is then zero even below the temperature where a bulk material orders. Learn how to extract order, correlation length and critical exponents from finite measurements, and which apparent crossings or peaks would survive the thermodynamic limit.

11. Phase nucleation, metastability and waiting-time statistics

A cooled liquid can remain liquid below its equilibrium freezing temperature and then transform abruptly. How can a bulk phase be thermodynamically favoured while a small embryo of it tends to shrink? Balance interfacial cost against bulk gain, convert a barrier into a probabilistic waiting-time prediction and test which measurements can distinguish surface nucleation from a bulk mechanism.

12. Work fluctuations and equilibrium free-energy differences

Repeated rapid changes of a laboratory control can give a broad distribution of work even when every run follows the same prescribed protocol. Derive why a particular exponential average retrieves an equilibrium free-energy difference, why the average work is generally larger, and why a finite experiment can miss the rare trajectories that make the equality true.

13. Put a thermometer on a bounded spin spectrum

A collection of identical two-level degrees of freedom has both a lowest and a highest energy. Count the accessible states at each energy and determine from adjacent entropy differences which way energy flows when two systems are brought into weak thermal contact.

14. Follow a refrigerant through four devices and a gas through a liquefier

A thermometer on a valve tells us little about the cooling capacity of a refrigerator. We need to follow one kilogram through its entire cycle and distinguish flash evaporation inside the valve from heat absorbed from the cold load.

15. Eutectic mixtures, freezing-point depression and boiling-point elevation

Adding a nonvolatile solute can lower the temperature at which a solvent begins to freeze. In a binary eutectic both species can solidify. The two observations follow from the same equality of chemical potentials, with different phases available for equilibrium.

16. From reaction affinity to a reversible cell voltage

An electrochemical cell transfers reaction free energy into electrical work through a controlled charge transfer. A measured voltage can therefore tell us about a reaction quotient, provided the electrodes are in reversible equilibrium.

Statistical mechanics from microscopic counting to material response · 1. Counting, absolute entropy and the entropy of mixing

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 $2.53k_B$. 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 $\exp[-NI(x)]/\sqrt{2\pi N x(1-x)}$, 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**

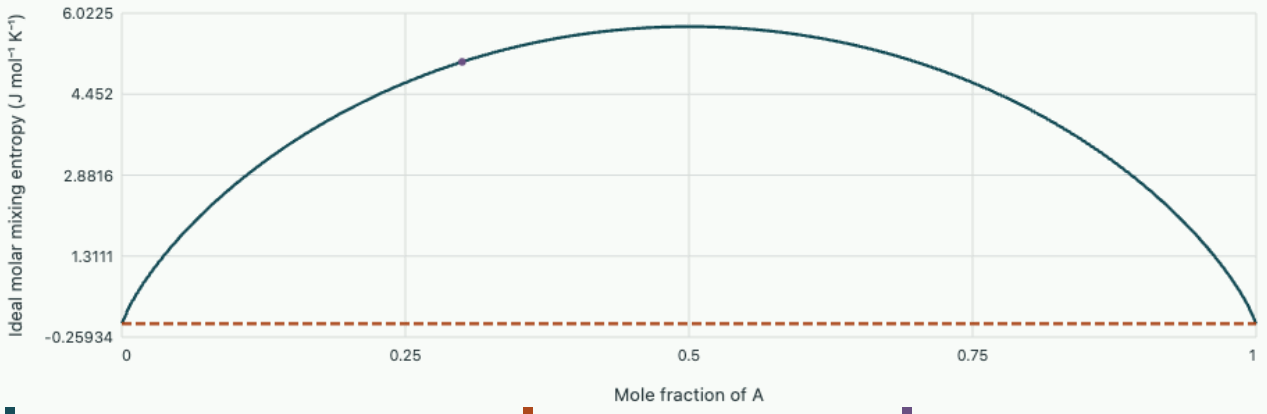

Two distinguishable ideal components | Identical gas on both sides | Selected composition

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

Statistical mechanics from microscopic counting to material response · 2. Ensembles, thermodynamic potentials and measurable curvature

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, $\Omega = -pV$.

$$\Xi = \sum_N e^{\beta\mu N} Z_N,$$

$$\langle N \rangle = \partial_{\beta\mu} \log \Xi,$$

$$\operatorname{var}(N) = k_B T \left(\frac{\partial \langle N \rangle}{\partial \mu} \right)_{T,V}.$$

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 $\operatorname{var} N / \langle N \rangle = n k_B T \kappa_T$. 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

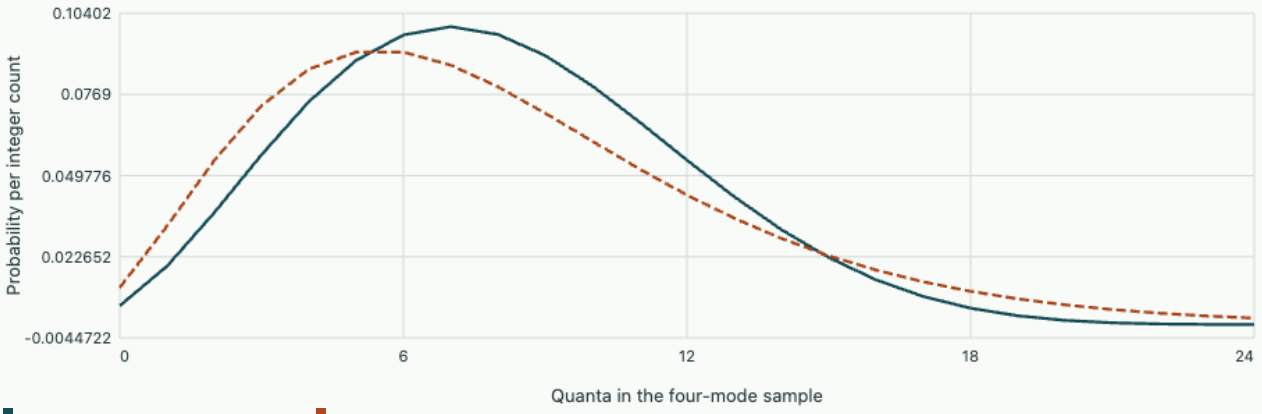

Isolated pair of solids Canonical sample

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

Statistical mechanics from microscopic counting to material response · 3. Pair correlations, the virial gas and a pressure measurement

A dilute nonideal gas can display almost ideal pressure while its particles strongly avoid one another at short distance. Connect a two-particle interaction to pressure, internal energy and a measurable pair correlation. Then use the resulting integral equations to ask which features of a potential can actually be inferred from a limited set of pressure readings.

Separate kinetic motion from the configurational count

For a classical monatomic gas at fixed temperature take pair potential $u(r)$ and total interaction $\sum_{i<j} u(r_{ij})$. The canonical partition function factors into a Gaussian momentum part and an integral over all allowed particle positions. Divide by $N!\lambda_T^{3N}$ for indistinguishable particles. In the absence of interactions the configurational integral is V^N. A hard core removes configurations with overlapping particles, while an attractive well increases their equilibrium weights instead of adding geometric space.

Use the Mayer function $f(r) = e^{-\beta u(r)} - 1$. The position-dependent Boltzmann weight is exactly $\prod_{i<j}[1 + f(r_{ij})]$. Expanding the product separates independent particles, terms with one pair bond, and terms involving several bonds. Translation invariance allows each isolated pair integral to be written as $V \int f(\mathbf{r}) d^3r$ when boundary corrections are negligible. The logarithm of the configurational partition function retains connected contributions, a crucial step before taking the large-volume limit.

$$\frac{F}{N k_B T} = \log(n\lambda_T^3) - 1 + n B_2(T) + O(n^2),$$

$$B_2 = -\frac{1}{2}\int f(\mathbf{r})\, d^3r.$$

A naive truncation of the full partition function after its first pair term can fail even for small density because there are of order N^2 pairs. The connected expansion organises the intensive free energy, not the sum of all events in a macroscopic gas. Differentiating F with respect to volume at fixed T, N gives

$$p/(n k_B T) = 1 + n B_2 + \cdots.$$

It is the smallness of successive density terms that controls the virial prediction. A small accidental value of B_2 does not bound the next coefficient.

Compute excluded volume and attraction as separate contributions

For a spherical hard core of diameter σ, $f = -1$ for $r < \sigma$, zero otherwise. The second virial coefficient is $2\pi\sigma^3/3$, four times the physical volume of a sphere of diameter σ. The forbidden region for the centre of a second particle has radius σ, twice a particle's radius, and each unordered pair is counted once. Treating the excluded volume per particle as merely its physical volume loses both geometric factors.

A square well of depth $\epsilon > 0$ extends from σ to $\lambda\sigma$. Outside the well the interaction vanishes. Evaluate the shell integrals directly:

$$B_2 = b[1 - (\lambda^3 - 1)(e^{\epsilon/(k_B T)} - 1)],$$

$$b = 2\pi\sigma^3/3.$$

The Boyle temperature is the solution $B_2 = 0$, namely $k_B T_B/\epsilon = 1/\log[1 + 1/(\lambda^3 - 1)]$. At T_B hard-core and attractive pair contributions cancel in pressure through first order in density. Particles still have nontrivial correlations and excess internal energy. The Boyle point is a pressure cancellation, not a noninteracting gas state. At fixed T and very low density the absolute pressure correction is small even when $|B_2|/b$ is large.

For $\lambda = 1.50$ and $\epsilon/k_B = 100$ K, $T_B \simeq 285$ K. At 200 K the ratio $B_2/b = 1 - 2.375(e^{1/2} - 1) = -0.541$. If $nb = 0.030$, the compressibility factor $Z_p = p/(nk_BT)$ is approximately 0.9838. At $T = 400$ K the ratio becomes 0.325 and the pressure exceeds ideal by about 0.98 percent for the same density. A precise comparison should also estimate $n^2 B_3$, especially when the objective is to infer model parameters from percent-level deviations.

Sweep well range and temperature through the pressure cancellation

An interaction can exchange heat even at the Boyle point

The canonical identity $U = \partial(\beta F)/\partial\beta$ at fixed volume and particle number gives the leading excess internal energy per particle as $-k_B T^2 n\, dB_2/dT$. For the square well,

$$\frac{U_{\mathrm{ex}}}{N} = -nb(\lambda^3 - 1)\epsilon e^{\epsilon/(k_B T)} + O(n^2).$$

This result is negative for an attractive well. Repulsive hard-sphere exclusion instead changes pressure but has no potential energy on any permitted configuration. These contrasting examples show that a pressure deviation alone does not determine an energy deviation without its temperature dependence. Different virial models may pass through the same pressure at one temperature while predicting very different heat capacities.

In the numerical well at 200 K the excess energy per particle equals $-0.1175\,\epsilon$, or approximately minus 1.62×10^{-22} J. Compare with the positive translational mean $3k_BT/2$; the correction is about minus four percent of that kinetic energy. A configuration-based simulation offers a direct independent check: count pairs in the well, assign minus ϵ per unordered pair, and average. Counting both neighbours of every pair without halving would double the inferred mean energy.

Thermodynamic consistency supplies another useful cross-check. The identity $(\partial U/\partial V)_T = T(\partial p/\partial T)_V - p$ requires the pressure and energy virial calculations to share the same derivative of B_2. For a hard-sphere gas with temperature-independent core the two sides vanish through this order. For an attraction they do not vanish. If a fitted empirical repulsion diameter depends on temperature, differentiate that dependence as part of the fit and interpret the resulting effective energy with care.

What a pair histogram measures

Choose one particle as origin and count the expected number of other centres in a spherical shell $4\pi r^2 dr$. Define $g(r)$ by setting this count equal to $ng(r)4\pi r^2 dr$. In the dilute limit a single neighbour at separation r

sees the relative Boltzmann weight $e^{-\beta u(r)}$, so $g(r) \to e^{-\beta u(r)}$. A hard core produces an empty inner region, an attractive well a finite elevated shell and large separations approach one. Histograms per unit radius without division by the shell volume grow as r^2 even for an ideal gas.

The leading configurational energy follows by integrating the pair energy weighted by the expected number of neighbours and dividing by two for pair sharing. Similarly, expanding a system uniformly by changing scaled coordinates in the configurational integral gives the virial pressure for a differentiable finite-range pair potential:

$$\frac{U_{\mathrm{ex}}}{N} = \frac{n}{2} \int u(r) g(r)\, d^3 r,$$

$$p = n k_B T - \frac{2\pi n^2}{3} \int_0^\infty r^3 u'(r) g(r)\, dr.$$

The pressure formula includes the radial force, whereas the energy integral includes the potential. Taking the dilute Boltzmann form and integrating by parts returns the second-virial result. For a hard discontinuity ordinary pointwise differentiation of $u(r)$ is inappropriate. Treat it through contact jump conditions or a limit of a smooth steep interaction. A code that samples a hard-core derivative as zero everywhere it is finite will miss exactly the repulsive momentum transfer it seeks to measure.

An optical or simulated correlation measurement with a finite point spread blurs the exclusion edge and can populate apparent distances below the true diameter. Calibrate spatial resolution and two-particle tracking efficiency, particularly where particle images overlap. In a periodic cube, shells whose radius approaches half the box edge need an appropriate geometric normalisation. Removing the outer shell deficit by an arbitrary normalisation may preserve a pretty asymptote while biasing an integral used for compressibility.

Connect density fluctuations to the integral of the pair structure

For a large observation window in an open homogeneous fluid, total particle-number variance can be calculated in two ways. The diagonal single-particle terms give the mean number. Distinct pairs add the integral of the connected density correlation. Provided correlations decay rapidly enough compared with the window size, one obtains a zero-wavevector structure factor

$$S(0) = \frac{\operatorname{var} N}{\langle N \rangle} = 1 + n \int [g(r) - 1]\, d^3 r = n k_B T \kappa_T.$$

To leading order $g - 1 = f$, so $S(0) = 1 - 2nB_2$. The compressibility derived by differentiating $p = n k_B T (1 + n B_2)$ is consistent to this order. The factor two relative to the pressure correction arises because pressure contains $n^2 B_2$ before differentiation. A mean pressure near the ideal value therefore need not imply that the full spatial correlation function is nearly flat.

One must specify the counting window. The exactly fixed total number in a sealed simulation box gives zero variance of that total. Measuring number fluctuations in subvolumes or the small-nonzero-wavenumber structure factor can estimate the thermodynamic value, with finite-size corrections. If the correlation length approaches the box size near a transition, extrapolation to zero wavevector using only one or two available modes is poorly conditioned.

An exactly soluble interacting gas as a test of truncation

A one-dimensional gas of hard rods provides a rare opportunity to compare an entire exact equation of state with its virial truncation. Take N rods of length a in an interval of length L with hard end walls. Order them by their left endpoints $x_1 < \cdots < x_N$. The no-overlap constraints become $x_{j+1} \geq x_j + a$, $0 \leq x_1$ and $x_N + a \leq L$. Introduce free coordinates $y_j = x_j - (j-1)a$. They obey $0 \leq y_1 \leq \cdots \leq y_N \leq L - Na$. Their configuration volume is consequently $(L - Na)^N / N!$. We have already chosen one ordering sector for indistinguishable rods, so another division by $N!$ would double-count the symmetry correction.

The momentum integral contributes one thermal wavelength per rod. Differentiating log partition function with respect to length gives the one-dimensional pressure, which has units of force:

$$p = \frac{N k_B T}{L - Na},$$

$$\frac{p}{n k_B T} = \frac{1}{1 - na} = 1 + na + (na)^2 + \cdots,$$

$$n = N/L.$$

In this geometry the second coefficient is a and successive virial terms are successive powers of the occupied length fraction. At $na = 0.10$ the first two terms give 1.10 rather than the exact 1.111, a one-percent relative pressure error. At $na = 0.40$ they give 1.40 instead of 1.667, a sixteen-percent error. The pressure becomes singular as free length vanishes even though the potential energy of every permitted configuration is zero.

A measurement of the mean internal energy alone would therefore find the same kinetic value as a one-dimensional ideal gas at each temperature and miss the entire excluded-volume pressure. Conversely a pressure fit over a narrow high-density interval can return an effective second coefficient quite unlike its true low-density limit. The exact hard-rod solution makes these two warnings quantitative without relying on a simulation or on a vaguely defined packing parameter.

There is also a contact-force interpretation. An arbitrary ordering gap between neighbouring rods has a broad distribution even if the total free length is strictly fixed. In the large-system limit the probability density for a free gap is exponential with mean $(1 - na)/n$. The pressure fixes its inverse scale $p/(k_B T)$, a one-dimensional illustration of how contact probabilities encode force on a wall. For a finite chain the fixed sum of all gaps correlates them, and their exact marginal distribution has a finite upper bound. A direct histogram of gap sizes is thus an independent test of whether a pressure-inferred excluded length really represents the particle geometry.

Infer a well range and depth from a two-temperature experiment

Suppose the core diameter has been determined from a high-temperature data set, and low-density pressure intercepts determine $B_2/b = -0.541$ at 200 K and $B_2/b = 0.325$ at 400 K. In the square-well model define $A = \lambda^3 - 1$ and $x = e^{\epsilon/(400 k_B)}$. The two observations imply $A(x^2 - 1) = 1.541$ and $A(x - 1) = 0.675$. Divide to obtain

$$x + 1 = 1.541 / 0.675 = 2.283.$$

Therefore $\epsilon/k_B = 400 \log 1.283 \simeq 99.8$ K, $A \simeq 2.38$, and $\lambda \simeq 1.50$. The algebra retrieves the earlier microscopic model using only two temperatures and a separately known core size.

The ratio calculation also reveals an ill-conditioned regime. At temperatures both high relative to the well depth, x differs from one by a small amount. A modest relative uncertainty in the measured excess over the hard-core baseline can then dominate $x - 1$, and range and depth will move in opposite directions in a fit. Put the estimated model back into both original pressure values, then predict a third intermediate temperature without refitting. A correlated misestimate of the core diameter shifts both ratios and can survive the first two consistency tests.

Consider an alternative attraction with a smooth tail but the same two measured virial coefficients. Since each coefficient is a weighted integral over all separations, agreement at two temperatures imposes only two scalar constraints on an entire function of distance. Measuring the dilute $g(r)$ supplies local information and can distinguish a broad shallow well from a narrow deep one. A third highly precise pressure reading cannot guarantee a unique functional form in the absence of a justified parametrisation, though it can falsify a particular two-parameter guess.

When estimating the pressure slope from experiment, do not fit Z_p at a single density as if the ideal intercept were measured there. Convert density units consistently, include the uncertainty in n entering both axes, fit several low-density points and test sensitivity to a quadratic term. If the inferred low-density intercept differs from one, check absolute pressure and temperature calibration before interpreting an empirical nonzero intercept as an interaction. No finite pair potential with a conventional dilute classical limit should create a constant deviation from ideality at zero density.

Use reciprocal-space scattering to test density fluctuations

A weak coherent probe scattered elastically by identical particles measures a momentum-transfer dependent structure factor. For positions $\mathbf{r}_j$, define $\rho_{\mathbf{q}} = \sum_j e^{-i\mathbf{q}\cdot\mathbf{r}_j}$. For nonzero $\mathbf{q}$ in a uniform sample the dimensionless static factor is $S(q) = \langle |\rho_{\mathbf{q}}|^2 \rangle / N$, after subtracting coherent mean components where required. Expanding the absolute square gives N self terms and pair terms. Isotropy then connects the measured angular spectrum to the radial correlation via

$$S(q) = 1 + 4\pi n \int_0^\infty r^2 [g(r) - 1] \frac{\sin(qr)}{qr}\, dr.$$

At large q the oscillatory factor tends to cancel smooth pair contributions, approaching the self term for suitable fluids. At very small q the factor tends to one, giving the compressibility sum rule stated earlier. An experiment rarely measures exactly $q = 0$ because the incident beam itself occupies that direction. Extrapolation from small but finite angles is required. A correlation extending farther in real space produces structure at smaller q, so insufficient low-angle coverage is equivalent to limited knowledge of long separations.

For hard spheres at leading density order, $g - 1 = -1$ inside the core and zero outside. Integrate the sine kernel over that sphere to obtain a pair suppression proportional to $[\sin(q\sigma) - q\sigma\cos(q\sigma)]/(q\sigma)^3$. Its limit is one third before multiplication by the geometric coefficient. At $q = 0$ the correction is $-4\pi n\sigma^3/3 = -2nB_2$, exactly the pressure derivative prediction. At finite angle the correction oscillates even though the pair potential contains no oscillatory part. Interpreting every peak in an angular spectrum as a separate preferred particle spacing would overinterpret a Fourier transform.

Actual measured intensity also contains a one-particle form factor, detector background, finite beam shape and instrumental angular response. Determine the form factor in an independently diluted reference, then divide only in bands where it is sufficiently far from zero. Corrected $S(q)$ should approach one at sufficiently large accessible q for a simple fluid model. A pressure-based low-q estimate and a independently estimated real-space core radius give two anchors on the scattering normalisation. Agreement across all three measurements is more discriminating than an arbitrary smooth normalisation of a single spectrum.

Separating exchange from interaction in a dilute virial measurement

An observed second virial correction is not always purely the configuration-space Mayer integral. Even an ideal quantum gas has exchange corrections when the thermal wavelength is appreciable. Expand the fugacity series of an ideal three-dimensional gas at small fugacity. With internal degeneracy one,

$$n\lambda_T^3 = z \pm z^2/2^{3/2} + \cdots,$$
$$p\lambda_T^3/(k_B T) = z \pm z^2/2^{5/2} + \cdots,$$

where the upper signs are for bosons. Eliminate z to second order to find

$$B_2^{\mathrm{exchange}} = \mp \lambda_T^3/2^{5/2}.$$

A Bose gas has a negative pressure correction and a Fermi gas positive at fixed density even without a material pair potential.

The relative importance scales as $n\lambda_T^3$, growing on cooling and for lighter particles. This can be distinguished from a hard-sphere core whose classical second coefficient is temperature independent. A real weakly interacting quantum gas can have both effects, and a full quantum virial coefficient may require scattering phase shifts or bound-state information. A classical configuration integral with an added exchange term is at best an explicitly justified limit, not a universally accurate combination.

Predict a Joule expansion from the measured virial curve

An independently calibrated expansion gives an energy-based test of the low-density pressure law. Suppose an insulated vessel of gas expands freely into a second evacuated chamber. No external work or heat is exchanged, so the total internal energy is constant. For a general single-component simple fluid write $dU = C_V dT + [T(\partial p/\partial T)_V - p]dV$. This identity follows from the Helmholtz mixed derivative and is testable even if detailed pair potentials are inaccessible. Insert the leading virial pressure at fixed particle number to connect temperature change directly with $B_2'(T)$.

The excess energy per particle in this approximation is $-k_B T^2 n B_2'(T)$. For weak nonideality use the ideal monatomic per-particle heat capacity $3k_B/2$ in the leading temperature correction and evaluate B_2' at the initial temperature. At fixed energy this gives

$$\Delta T \simeq (2/3) T^2 B_2'(T)(n_f - n_i).$$

An attractive well with $B_2' > 0$ cools on expansion: fewer attractive neighbours raise potential energy, reducing the kinetic contribution. Pure hard-core exclusion has $B_2' = 0$ and does not change internal energy

at fixed temperature in this model.

In the earlier square-well example the initial excess energy at 200 K and $n_i b = 0.030$ is approximately minus $11.75 k_B$ kelvin per particle. Doubling the available volume at nearly unchanged temperature halves that excess magnitude. The approximate increase in configurational energy is $5.88 k_B$ kelvin per particle, implying a translational temperature drop near 3.92 K. A fuller calculation solves the constant-energy equation with the final temperature inside the well Boltzmann factor; the difference between that and the linear estimate is a measure of the size of the correction itself.

Distinguish this experiment from throttling through a steady porous restriction, where enthalpy rather than internal energy is approximately conserved under idealised insulated flow conditions. A real-gas Joule–Thomson coefficient depends also on thermal expansion and molar volume. Pressure and density units can look almost identical in the two calculations, but the device boundary and the work term differ. Logging the direction of mass transfer and whether the total confined particle number remains fixed prevents a common misidentification.

Repeat the free expansion at several initial densities while controlling chamber heat capacity. A temperature change proportional to the density difference supports the leading pair correction; a constant offset suggests a thermometer or wall heat effect. An accidentally zero pressure deviation at the Boyle temperature need not make the expansion temperature change vanish, because the temperature derivative of the virial coefficient remains finite there. Jointly observing near-ideal pressure and measurable cooling is exactly the distinction between a function crossing zero and its derivative crossing zero.

A useful uncertainty estimate comes directly from the linear expansion thermometer. In the illustrated twofold expansion the roughly four-kelvin temperature shift scales with the initial density and with the temperature derivative of the virial coefficient. If both wall thermometers share a common zero offset, an initial-minus-final temperature difference largely cancels it, but a drifting temperature scale can remain. Use a control expansion of a gas well represented as ideal under the same laboratory conditions; an apparent cooling there bounds residual wall and valve effects. Independently reversing the gas transfer between two matched bulbs changes wall-specific artefacts while retaining the expected sign for true free expansion. Quantitative agreement should involve the gas energy change plus heat transferred to thermometer mounts and vessel walls, not merely two indicated peak temperatures.

Recover a virial coefficient from a pair interaction

Vary the range and temperature of an attractive interaction. Identify where its attractive and excluded-volume contributions cancel.

Open this experiment online

Try the calculation

Outer radius / core diameter 1.5

Reduced temperature $k_\beta T/\varepsilon$ 2

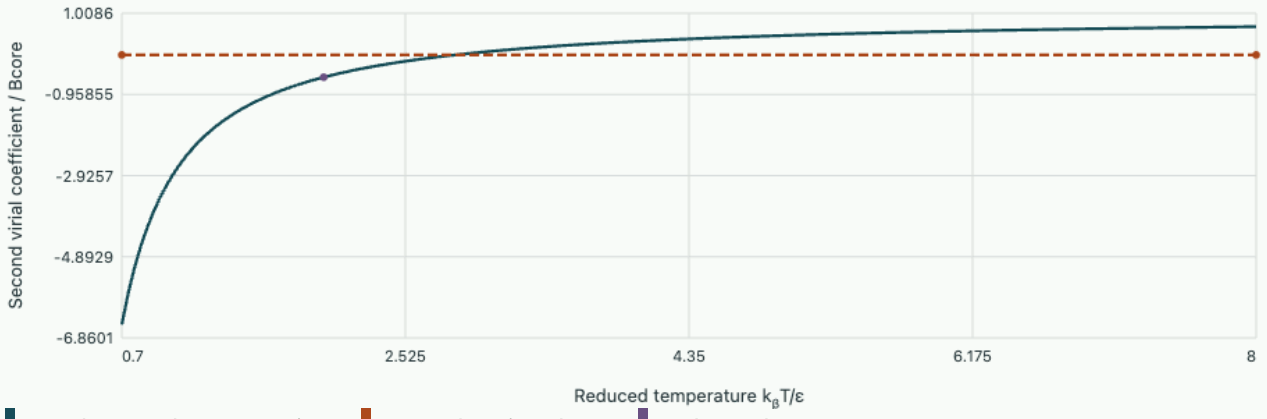

Hard core plus attraction | Zero density slope | Selected temperature

B_2/Bcore = -0.5407. The Boyle temperature is 2.846 ε/k_β.

Tutorial · Determine an effective diameter from high-temperature pressure

A gas at high enough temperature for its attraction to be negligible is dilute. A fit of $Z_p - 1$ versus particle density gives a slope 1.00×10^{-28} cubic metres. Infer the hard-sphere diameter. A student instead interprets that slope as the volume of one sphere. Quantify the error.

▼ **Apply the pair geometry and inspect the fit range**

Set the slope equal to $2\pi\sigma^3/3$, giving $\sigma = 3.63 \times 10^{-10}$ m. Using the physical sphere volume $\pi\sigma^3/6$ instead makes the diameter too large by $4^{1/3} = 1.587$, giving 0.576 nm rather than 0.363 nm. Check the fit extrapolated toward zero density and examine curvature as the maximum included density is reduced. A small slope uncertainty based on a linear regression does not include a systematic neglected third virial term.

Tutorial · Use energy to break an attractive-potential degeneracy

At one temperature two square wells with the same core yield the same negative B_2, one deep and narrow, one shallow and broad. Show how an independent calorimetric energy measurement can distinguish them and specify a limiting regime in which the additional measurement becomes weak.

▼ **Identify two different combinations of range and depth**

At fixed temperature the pressure supplies $(\lambda^3 - 1)(e^x - 1)$, with $x = \beta\epsilon$. The excess energy divided by $-nbk_BT$ instead supplies $(\lambda^3 - 1)xe^x$. Their ratio is $x/(1 - e^{-x})$, a monotone function of positive x. With a known core and sufficiently accurate data it determines x, then the original pressure combination determines

range. At $x \ll 1$ the ratio approaches $1 + x/2$; its excess above one becomes small compared with calibration noise. Several temperatures can then be more informative than one highly uncertain difference.

Tutorial · Design an independent histogram check of pressure and energy

A dilute well fluid has measured $g = 0$ in the core, $g = 1.65$ across its attractive shell and $g = 1$ outside. Take $\lambda = 1.5$, $nb = 0.03$ and $T = 200$ K. Infer approximate well depth and mean neighbours in the shell. Compare the resulting potential energy with the thermodynamic prediction.

▼ **Integrate the actual shell volume and count unordered pairs**

The dilute relation gives

$$\epsilon/k_B = T \log 1.65 = 100.2\,\mathrm{K}.$$

The shell volume is $2b(\lambda^3 - 1)$, giving mean shell neighbours per particle $0.03(2)(2.375)(1.65) = 0.235$. Dividing by two gives mean potential energy -0.1176ϵ per particle, matching the derivative calculation. Using an unnormalised radial pair histogram or counting the same bond twice would spoil this agreement.

Now repeat at higher density. A measured effective potential $-k_B T \log g$ is then a potential of mean force containing the effect of other neighbours. Inserting it into the original low-density two-body virial integral as if it were the bare pair potential double-counts many-body correlations. The direct energy integral using a known actual pair potential and the measured full g remains appropriate for a pairwise-additive fluid.

Cambridge Statistical Physics · interacting classical gases and cluster expansion

Statistical mechanics from microscopic counting to material response · 4. Rotational and vibrational populations in a molecular gas

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 Nk_BT/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 $Nk_B \log g$ 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 $\theta_r = \hbar^2/(2Ik_B)$. 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**

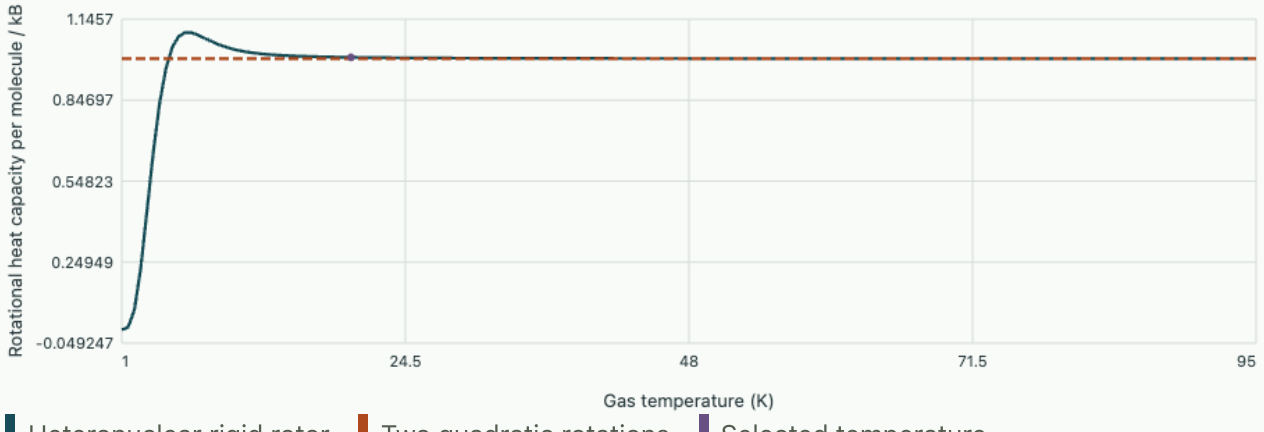

First excited rotational level has E_1/kB = 16 K. At the selected temperature Crot/kB = 1.0051. Vibrational and translational contributions are excluded.

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

Statistical mechanics from microscopic counting to material response · 5. Bose condensation, dimensionality and a finite trap

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

$$\frac{N_{\mathrm{ex}}}{V} = \lambda_T^{-3} g_{3/2}(z),$$

$$p = k_B T \lambda_T^{-3} g_{5/2}(z),$$

$$U_{\mathrm{ex}} = \frac{3}{2} pV.$$

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

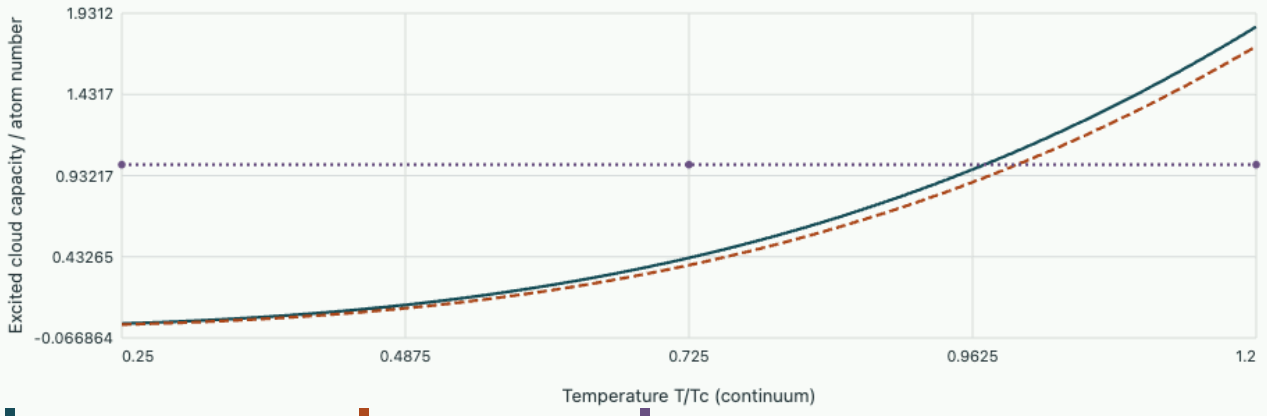

Counted harmonic levels | Continuum density | All atoms

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

Statistical mechanics from microscopic counting to material response · 6. Degenerate fermions, heat and magnetic 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

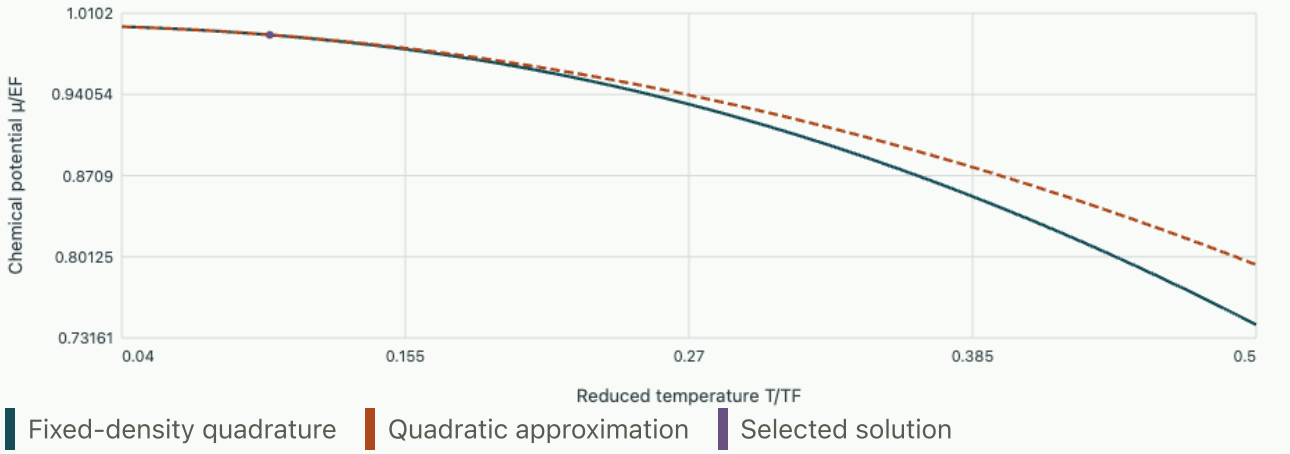

Fixed-number μ/EF = 0.9916, quadratic prediction 0.9918. Leading heat capacity $C_v/(Nk_\beta) = 0.4935$.

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

Statistical mechanics from microscopic counting to material response · 7. Master equations, detailed balance and probability currents

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 $2k_B T$, and put a transition configuration $10k_B T$ 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 $2k_B T$ 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

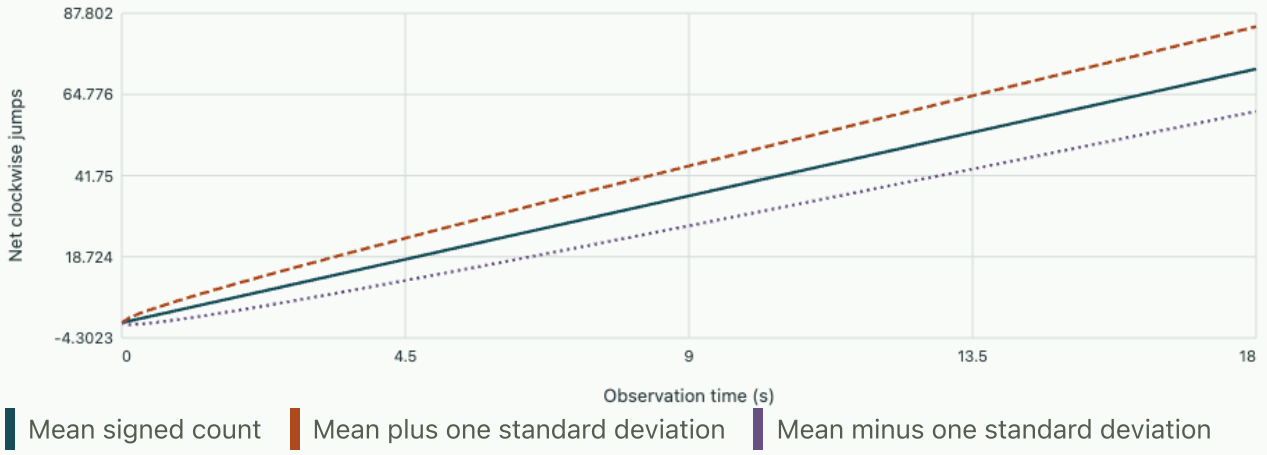

Every stationary population is 1/3. Net clockwise edge current 1.3333 s⁻¹. Entropy production 4.3944 kB/s.

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

Statistical mechanics from microscopic counting to material response · 8. From binary collisions to the Boltzmann equation

An initially anisotropic velocity distribution can relax even though every individual collision conserves energy and momentum. Describe a dilute gas with a phase-space density, derive the structure of its collision term and prove which distribution can stop producing entropy. Then examine the assumptions behind that proof and the measurements most likely to reveal their failure.

What the one-particle function counts

Let $f(\mathbf{r}, \mathbf{v}, t) d^3r\, d^3v$ be the expected number of particles in a small phase-space cell. Integrating over velocities gives number density n, and the mean velocity is $\mathbf{u} = n^{-1} \int \mathbf{v} f\, d^3v$. Define the peculiar velocity $\mathbf{c} = \mathbf{v} - \mathbf{u}$; its second moment sets local kinetic energy after removing the organised flow. A gas can have large bulk speed but low temperature.

Follow a particle ballistically over a short interval under acceleration $\mathbf{a} = \mathbf{F}/m$. With no collisions, phase-space probability is advected along its trajectories. For a standard position-dependent force, phase-space flow has zero divergence, giving the streaming operator $\partial_t f + \mathbf{v} \cdot \nabla_r f + \mathbf{a} \cdot \nabla_v f$. Magnetic Lorentz acceleration also has zero velocity-space divergence. A dissipative velocity-dependent drag generally requires the divergence form $\nabla_v \cdot (\mathbf{a} f)$ instead. It should not silently be inserted into the Hamiltonian streaming expression.

We now assume a dilute classical gas of identical particles with short-ranged binary elastic collisions, approximated as local on the spatial scale resolved by f. Different particles approaching collision are assumed uncorrelated apart from their individual one-particle distributions at that position. This molecular-chaos approximation is an assumption about incoming pairs, not a claim that the outgoing particles remain uncorrelated. It is the step that allows the evolution to close on a one-particle function.

Count pairs crossing a scattering area

Consider a tagged incident velocity $\mathbf{v}$ and a possible partner $\mathbf{v}_1$. Their relative speed $g = |\mathbf{v} - \mathbf{v}_1|$ converts a differential scattering cross section $d\sigma/d\Omega$ into the swept interaction volume per unit time. A binary collision rotates the relative velocity in the centre-of-mass frame while preserving its magnitude. With equal masses, set $\mathbf{V} = (\mathbf{v} + \mathbf{v}_1)/2$ and $\mathbf{g} = \mathbf{v} - \mathbf{v}_1$. The partner and tagged postcollision velocities are $\mathbf{V} \pm \mathbf{g}'/2$, where $|\mathbf{g}'| = |\mathbf{g}|$. This parametrisation enforces momentum and kinetic energy conservation before choosing a scattering direction.

The loss term counts incoming pairs with velocities $\mathbf{v}, \mathbf{v}_1$. The gain counts inverse collisions that produce the specified velocity. Microreversibility of elastic scattering and invariance of the integration measure allow the compact expression

$$C[f](\mathbf{v}) = \int d^3v_1\, d\Omega\, g \frac{d\sigma}{d\Omega} \left[f(\mathbf{v}') f(\mathbf{v}_1') - f(\mathbf{v}) f(\mathbf{v}_1) \right].$$

Both products refer to the same position and time. A change in which pair is called primed merely swaps gain and loss if one reverses the parametrisation consistently. This formula uses a consistent distinct-partner convention for the full scattering measure. When counting the total number of collision events per volume

rather than the loss of tagged particles, divide the sum over all ordered pairs by two. Confusing those two observables gives an erroneous collision frequency.

For hard spheres of molecular diameter d, the total geometric cross section is πd^2. In an equilibrium Maxwellian the mean relative speed is $\sqrt{2}$ times the mean single speed for identical masses. The mean free path under the usual hard-sphere convention is $\ell = 1/(\sqrt{2} n \pi d^2)$. For $d = 0.36$ nm, $n = 2.4 \times 10^{25}$ per cubic metre and molecular mass 28 unified atomic units at 300 K, $\ell = 72.4$ nm and $\bar{v} = 476$ m/s. Their ratio gives a characteristic flight time 0.152 ns.

Take a weak moment of the collision operator

For a test function $\psi(\mathbf{v})$, multiply the collision term by ψ and integrate over tagged velocity. Exchange the identities of both particles and then exchange incoming and outgoing velocities in the gain integral. The result expresses the change of the pair's total test-function value:

$$\int \psi C[f] \, d^3v = \frac{1}{2} \int d^3v \, d^3v_1 \, d\Omega \, g \frac{d\sigma}{d\Omega} f f_1 [\psi' + \psi_1' - \psi - \psi_1].$$

For $\psi = 1$, $m\mathbf{v}$ and $mv^2/2$ the square bracket vanishes collision by collision. Integrating the kinetic equation gives exact local conservation of mass, momentum and energy with appropriate streaming fluxes, regardless of how close the distribution is to Maxwellian. Conservation alone leaves a hierarchy: the momentum flux depends on the pressure tensor and the energy flux on a heat flux, both higher velocity moments.

A numerical collision solver should check all these invariants separately. Equal masses allow a direct test with velocities $(300, 100, 0)$ and $(-100, 100, 0)$ m/s. Their centre velocity is $(100, 100, 0)$ and relative velocity has magnitude 400 m/s. Rotating it from x to y gives outgoing velocities $(100, 300, 0)$ and $(100, -100, 0)$. Total velocity and total squared speed agree before and after, but the directional kinetic energy has changed. A field of such collisions can isotropise the pressure tensor while leaving total energy fixed.

Derive a sign-definite entropy production

Define the local kinetic H density as $\int f \log(f/f_\star) d^3v$ with a constant reference of matching units. The reference only changes a conserved number contribution. For a homogeneous isolated gas the collision contribution to its time derivative follows by applying the weak form with $\psi = \log f$. Average the result with its incoming–outgoing interchange. Setting $A = f f_1$, $B = f' f_1'$, one obtains

$$\frac{dH}{dt} = -\frac{1}{4} \int d^3v \, d^3v_1 \, d\Omega \, g \frac{d\sigma}{d\Omega} (A - B) \log(A/B) \le 0.$$

For positive A, B, logarithm and difference have the same sign; appropriate limits handle zero occupation. Thus the Boltzmann entropy $-k_B H$ does not decrease from these local binary collisions. In an inhomogeneous region, advection additionally transports entropy through the boundaries. A local cell's kinetic entropy can drop because entropy leaves the cell, while collision production remains nonnegative.

Equality requires $f f_1 = f' f_1'$ for collisions connected by the scattering kernel. Equivalently $\log f + \log f_1$ is a collision invariant. For sufficiently rich elastic scattering in three dimensions the smooth additive collision invariants are linear combinations of particle number, momentum and energy. Therefore

$\log f = a + \mathbf{b}\cdot\mathbf{v} + cv^2$, with $c < 0$ for normalisation. Completing the square and matching moments yields the local Maxwellian:

$$f_M = n\left(\frac{m}{2\pi k_B T}\right)^{3/2} \exp[-m|\mathbf{v}-\mathbf{u}|^2/(2k_B T)].$$

A Maxwellian at each location makes the collision term zero locally, but spatial gradients usually leave a nonzero streaming term. Such a distribution can evolve because neighbouring cells exchange particles. Conversely, a truly global equilibrium in an external conservative potential includes the appropriate spatial Boltzmann factor. In one-dimensional equal-mass elastic collisions particles merely exchange velocities, so the assumed richness of three-dimensional scattering fails; an arbitrary velocity histogram need not thermalise through those collisions.

Assess the kinetic and hydrodynamic descriptions separately

Define Knudsen number $Kn = \ell/L$ using the length over which the distribution or boundary condition changes. In the earlier gas $Kn = 7.2 \times 10^{-5}$ for a 1 mm chamber feature but 0.072 for a 1 micrometre channel gap. The same binary collision model can remain dilute in both, yet a no-slip bulk hydrodynamic closure may be much less accurate in the microchannel. A mean free path small relative to the overall device length is insufficient if a sharp boundary layer develops on a much shorter scale.

A useful relaxation-time collision approximation is $C_{\mathrm{BGK}} = -(f - f_M)/\tau$, with f_M computed from the local conserved moments of f. This enforces their conservation and models return to local Maxwellian without evaluating an angular integral. If an externally fixed Maxwellian is used irrespective of actual local density, momentum and energy, the approximation can remove these quantities spuriously. A single fitted relaxation time also need not reproduce both stress and heat transport for a real gas.

A particle simulation resolving binary encounters must choose spatial cells small enough to represent local variation, and event or collision timesteps short against relevant flight or collision times. Statistical noise decreases with independent sampled pairs; refining cells while keeping total simulated particle count fixed can make local collision estimates noisier. Compare conservation, relaxation of anisotropic moments and a transport observable under separate refinements of geometry and sample count. A correct mean pressure alone is a weak test of a collision operator.

Compare the flux-weighted velocities of escaping particles with the bulk velocity distribution

Take the next moment and exhibit the source of viscosity

Write momentum density as $\rho u_i = m\int v_i f\, d^3v$. The flux of i-momentum through a surface with normal in direction j is $m\int v_i v_j f\, d^3v = \rho u_i u_j + P_{ij}$, where $P_{ij} = m\int c_i c_j f\, d^3v$. The scalar pressure is one third of its trace in three dimensions. Multiplying the kinetic equation by mv_i and using collision invariance yields the momentum balance with $\partial_j P_{ij}$ explicitly. No constitutive assumption has yet set the off-diagonal components to zero or made them proportional to strain rate.

For a local Maxwellian, the odd moments of peculiar velocity vanish and $P_{ij} = nk_B T\delta_{ij}$. Now consider slow shear with mean velocity $u_x(y)$, and approximate collisions by a relaxation time τ. In a local coordinate frame

moving with the fluid write $f \simeq f_M - \tau v_y \partial_y f_M$ to first order in gradient. At constant density and temperature,

$$\partial_y f_M = (m c_x / k_B T) f_M \partial_y u_x.$$

The off-diagonal momentum flux correction is therefore a fourth Gaussian velocity moment:

$$P_{xy} = -\tau \frac{m^2}{k_B T} \partial_y u_x \int c_x^2 c_y^2 f_M d^3 v = -n k_B T \tau \, \partial_y u_x.$$

The integral factorises into variances $k_B T / m$ in the two directions. Under the kinetic pressure-tensor convention the result has a minus sign, because random motion transports high x-momentum toward lower mean velocity. The viscous part of the Cauchy stress is commonly written with the opposite sign convention. State which tensor enters the balance before declaring a predicted shear stress negative. The dynamic viscosity of this simple model is $\eta = p\tau$.

If $p = 1.0 \times 10^5$ Pa and the fitted stress relaxation time is 1.8×10^{-10} s, this gives 1.8×10^{-5} Pa s. At a shear rate of 10^3 per second the dissipative stress magnitude is 0.018 Pa, tiny compared with equilibrium pressure and with dimensionless shear rate $\tau |\partial_y u_x| = 1.8 \times 10^{-7}$. A stress measurement must resolve a small anisotropic contribution on top of a large isotropic pressure. A pressure gauge alone cannot test this kinetic viscosity relation.

In a dilute hard-sphere gas the mean collision time scales inversely with density. The simultaneous pressure factor then makes dynamic viscosity approximately independent of density at fixed temperature in the dilute continuum regime. The kinematic viscosity, η / ρ, still changes strongly. A gas made sufficiently rarefied no longer supports ordinary no-slip continuum flow through a fixed channel, so an apparent pressure-dependent flow resistance can coexist with an essentially density-independent bulk dynamic viscosity.

Infer accommodation from a wall momentum measurement

The interior collision operator does not prescribe the boundary condition at a solid wall. Specular reflection reverses normal velocity while preserving tangential velocity in the wall frame. Diffuse thermal re-emission from a fully accommodating stationary wall instead draws outgoing velocities from a wall-temperature distribution appropriate to the outward flux and erases incident tangential drift. A real clean or rough surface may lie between these limits. The boundary scattering model matters particularly when particles encounter the wall as often as other particles.

For a simple moving plate in a dilute gas at the same temperature, suppose a fraction α of incident particles accommodates tangentially to plate speed U, and the rest reflects specularly. To leading order in small tangential drift, accommodated particles carry away an average tangential momentum increment mU if incident gas has zero mean tangential speed. The incident number flux on one plate face is $n\bar{v}/4$. Thus the leading tangential drag magnitude on the plate is $\alpha m U n \bar{v} / 4$ per area in this ideal free-molecular geometry.

Take $n = 10^{21}$ per cubic metre, molecular mass 4.65×10^{-26} kg, $\bar{v} = 480$ m/s, $U = 2.0$ m/s and $\alpha = 0.80$. The predicted shear is about 8.93×10^{-3} Pa per exposed face. The effect scales linearly with density here, unlike the dilute continuum viscosity at a fixed macroscopic velocity gradient. Both results derive from momentum transport, but the first describes particles crossing a wall-bounded collisionless region and the

second local shear sustained by many intermolecular collisions. Compare them only after specifying the Knudsen number and wall geometry.

To test the accommodation model, reverse the plate motion to remove static force offsets, vary density across a regime that remains free molecular and change the exposed surface while keeping area fixed. A single measurement gives the combination of accommodation fraction, incident number flux and velocity difference. It cannot separate a changing incident gas drift from a changing material accommodation coefficient. Measuring the upstream velocity distribution or using opposed plates with known relative motion supplies additional constraints.

A discrete collision quartet demonstrates the H theorem numerically

A single allowed reversible collision connects two incoming velocity cells to two outgoing cells. Suppose cell-volume and scattering factors have been normalised so that its net flux is $K(A - B)$, with $A = f_1 f_2$, $B = f_3 f_4$. It removes particles from cells one and two and adds them to three and four when $A > B$. Conservation requires the representative cell velocities to satisfy the same pair momentum and energy constraints used for a physical collision. Assigning arbitrary four velocities may conserve the number of particles but will generally create energy.

Choose the velocities from the earlier rotated-pair example and occupancies $f_1 = 4$, $f_2 = 1$, $f_3 = 2$, $f_4 = 1$ in a common dimensionless reference. Then $A = 4$, $B = 2$ and the forward flux is $2K$. Its contribution to dH/dt is $-2K \log 2$, negative. The exact continuous rate tends to zero as the products equalise. A forward Euler step larger than $f_2/(2K)$ would make the second occupancy negative and destroy the logarithm's meaning. Nonnegative entropy production for the continuum collision model does not guarantee it for an arbitrary numerical time step.

An implicit or positivity-preserving update can avoid this failure, but still requires independent tests of conserved moments and of the correct equilibrium. For this isolated quartet the condition $f_1 f_2 = f_3 f_4$ alone leaves many possible stationary cell distributions. A network of collision quartets sufficiently connecting the velocity space is needed to enforce a Maxwellian-like equilibrium. On a coarse discrete velocity lattice too few permitted momentum- and energy-conserving quartets may leave artificial invariants. A code can display exactly conserved energy and monotonically decreasing H while relaxing to the wrong distribution because its collision graph is incomplete.

A velocity-selective aperture tests collision equilibration

A thin aperture feeding a vacuum produces a flux-weighted distribution from a local Maxwellian reservoir. A detector placed downstream generally selects a finite angular cone and may respond differently to fast and slow particles. Consider two narrow speed bins with the same width in speed and the same angular acceptance. For an effusive Maxwellian their flux ratio contains a cube of speed multiplied by a Boltzmann exponential. A bulk sampler instead has a square of speed. Confusing these prefactors shifts a temperature inferred from ratios even if the gas truly is locally equilibrated.

For bins at speeds v and $2v$, and choosing $mv^2/(2k_B T) = 1$, the effusive high-to-low ratio is

$$8e^{-3} = 0.398.$$

The bulk number-density ratio at those speeds would be $4e^{-3} = 0.199$. A downstream detector measuring particles present in a fixed optical observation volume may weight by travel time, introducing an additional approximate inverse speed that moves the response closer to density weighting. Distinguish arrivals counted per time from particles counted during an exposure.

Changing aperture diameter is an independent control. A broad opening or high upstream density can produce collisional expansion and a directed jet, altering both angular and speed distributions. Measure pressure and beam shape across aperture sizes and compare the small-aperture limit before attributing a non-Maxwellian downstream spectrum to a failure of local thermal equilibrium throughout the reservoir.

A flight-time measurement tests free paths and speed selection separately

Imagine a dilute molecular tracer introduced into a uniform bath. For a deliberately simplified constant collision-rate model with rate r, probability of no collision by time t is e^{-rt}. For tracers conditioned to have one speed v, the distribution of collision-free distance is then exponential with mean v/r. If the collision rate itself depends on speed through the average relative velocity against the bath, the conditional mean free path is $v/r(v)$. Mixing all initial speeds need not produce a single exponential in distance, even if every conditional waiting time is exponential.

Place a detector a distance L downstream that accepts only trajectories maintaining their initial direction. For a narrow injected speed class, the unscattered fraction is approximately $e^{-r(v)L/v}$ if any collision removes it from the acceptance. For a class at 500 m/s and mean free path 0.20 m, a detector at 0.10 m receives an unscattered fraction 0.607; at 0.30 m it receives 0.223. A ratio of the two readings removes the unknown injection strength and recovers the attenuation coefficient if detector acceptances and geometric factors match.

If forward-scattered trajectories remain inside the acceptance, the apparent attenuation is smaller than the total collision rate. Angularly narrowing the receiver should approach the total-removal model for suitable scattering. A finite source pulse and finite time-of-flight bin also select velocities even if the source initially emitted a Maxwellian. Fit survival separately in narrow speed bins first. A nonexponential whole-beam attenuation then need not imply that the microscopic collision process has memory; it can arise purely from an unresolved distribution of rates.

A simulation can distinguish these alternatives by logging every collision event and comparing the full collision count with the detector-defined removed count. First test conditional survival for a prepared monochromatic beam, then add a known distribution of incident speeds and calculate the mixture prediction explicitly. Finally include a finite angular acceptance and integrate the differential cross section over the appropriate retained directions. These steps separate the kinetic collision operator, the initial distribution and the observation rule.

The same logic applies to a relaxation experiment reporting a single characteristic mean free time. The decay time of a directional flux depends on how much momentum is deflected per collision, not simply on how often particles meet. Strongly forward-peaked scattering can have a large total cross section and a much smaller momentum-transfer cross section weighted by $1 - \cos\theta$. A measured attenuation and a measured viscosity may consequently constrain different moments of precisely the same differential scattering process.

For example, if every collision deflects a test particle through a small fixed angle of ten degrees in the relevant frame, the momentum-transfer weighting is

$$1 - \cos 10^\circ = 0.0152.$$

The corresponding transport cross section can then be far smaller than the total cross section even when collisions are frequent. A velocity-randomising approximation that uses the total collision time as its stress-relaxation time would predict much faster momentum relaxation than this angular picture supports. Actual stress and energy transport involve additional angular weights and velocity averages, so this numerical comparison is a scale diagnostic rather than an exact viscosity formula. Measure several angular moments of the differential scattering kernel or compare at least two independently sensitive transport observables before assigning a universal relaxation time to a real collisional gas.

Follow an H functional during a reversible elementary collision

Prepare unequal populations in two velocity pairs. Count forward and reverse collision rates and compare the sign of entropy change with the population imbalance.

Open this experiment online

Try the calculation

Initial fraction in velocity pair A **0.83**

Collision rate coefficient (s⁻¹) **0.8**

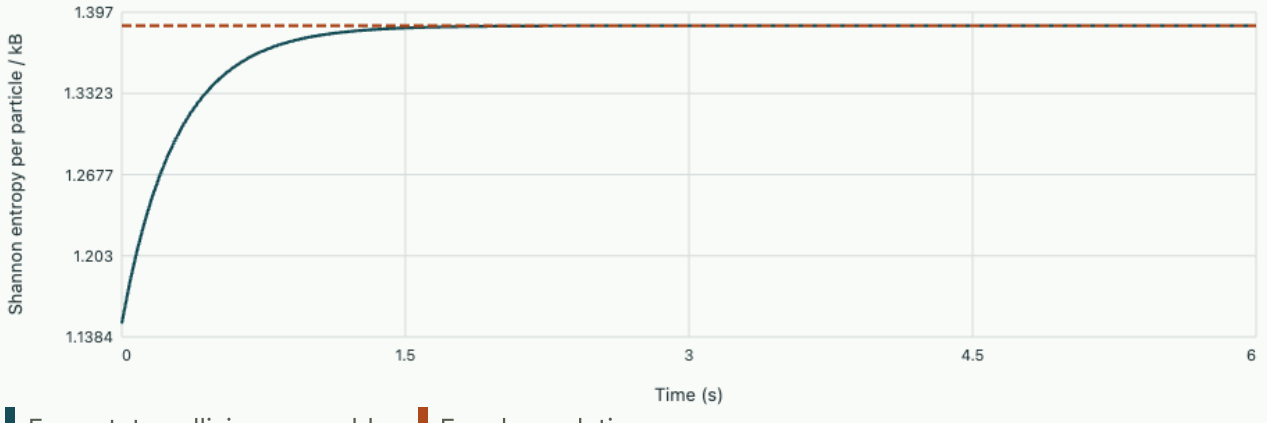

Initial clockwise-minus-reverse pair imbalance x²−(1−x)² = 0.66. Initial d(S/kB)/dt = 0.83721 s⁻¹.

Tutorial · Find the equilibrium temperature of an anisotropic preparation

A homogeneous gas has zero mean flow and component kinetic temperatures 600, 300 and 150 K, defined by $m\langle v_i^2\rangle = k_B T_i$. In an isolated elastic gas, predict its equilibrium temperature and pressure tensor change.

Could a thermostat holding the velocity variance in one direction at its original value permit the same inference?

▼ **Retain all three conserved energy contributions**

Conservation of total kinetic energy implies $T_f = (600 + 300 + 150)/3 = 350$ K. Initially the diagonal kinetic pressure components are nk_BT_i; ultimately each becomes $nk_B(350\,\mathrm{K})$. Their sum is unchanged. A component singled out by a thermostat can exchange work or heat with its environment. Its final common temperature, if one exists, must be obtained from the full energy balance including that exchange. Isotropisation cannot simply be assumed to retain the initial isolated mean.

Tutorial · Estimate the error in a collision-frequency inference

A detector reports a mean speed of 480 m/s and a geometric collision cross section 4.0×10^{-19} square metres at $n = 2.0 \times 10^{25}$ per cubic metre. Estimate the collision rate per tagged particle and total events per unit volume per second for an equilibrium identical-particle gas. What changes in a molecular beam with almost parallel velocities?

▼ **Use the relative rather than the single-particle speed**

The equilibrium mean relative speed is $\sqrt{2}(480) = 679$ m/s. The rate per tagged particle is $ng\sigma \simeq 5.43 \times 10^9$ per second. Multiplying by $n/2$ gives event rate 5.43×10^{34} per cubic metre per second. In a nearly monoenergetic collimated beam the relative speed can be much less than the mean laboratory speed. Evaluate its pair-distribution average explicitly; translating every velocity by a common vector leaves all collision rates unchanged at fixed distribution of relative velocities.

Tutorial · Does reversing all velocities contradict the H theorem?

In a dilute-gas simulation, reverse all particle velocities at a chosen instant with exact precision. An ideal reversible microscopic integrator would retrace the prior trajectory. Explain how this outcome can coexist with a kinetic equation predicting nondecreasing entropy, and propose a comparison that isolates the missing information.

▼ **State which distributions are controlled by the reversal**

The reversed incoming pairs carry precise correlations inherited from earlier outgoing collisions. Their two-particle distribution generally does not factor into the product of their one-particle distributions just before future collisions. The molecular-chaos step used to derive the collision operator therefore fails for this deliberately prepared history. Microscopic Hamiltonian evolution and the coarse one-particle closure make different state specifications.

Prepare a second gas with the same reversed one-particle velocity histogram but with spatial and partner correlations randomised while conserving total momentum and energy. Compare the two subsequent anisotropy histories. The kinetic equation has the same initial one-particle input for both, while a microscopic simulation can distinguish them through higher-order correlations. Imperfect numerical reversal introduces a

separate error; test exactness and time-step dependence before interpreting a failure to retrace as a fundamental thermodynamic observation.

Cambridge Kinetic Theory · the Boltzmann equation, collision invariants and hydrodynamics

Statistical mechanics from microscopic counting to material response · 9. Transport coefficients from equilibrium fluctuations

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

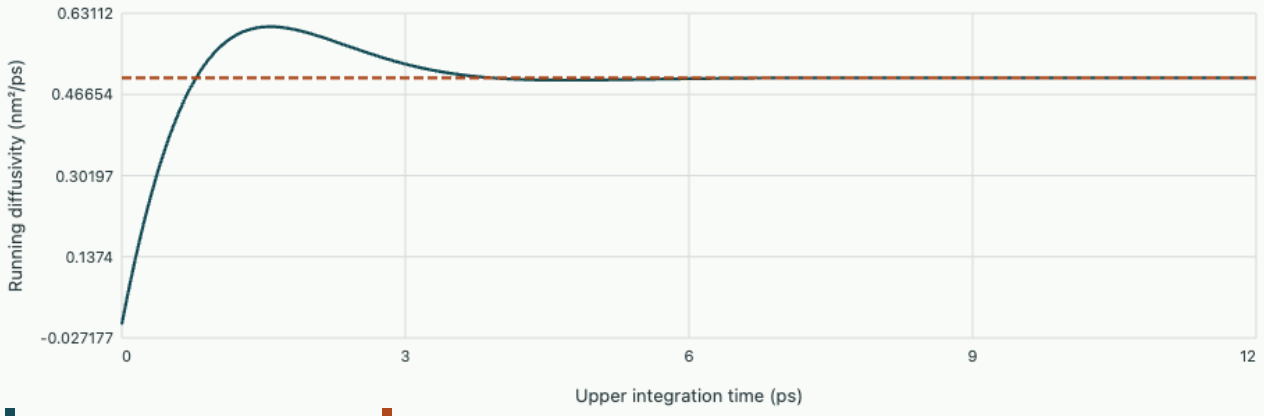

With C(0) = 1 nm²/ps² the exact 1D diffusion coefficient is 0.5 nm²/ps. At 2τ the integral gives 0.58969 nm²/ps.

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

Statistical mechanics from microscopic counting to material response · 10. Finite-size criticality and the evidence for a phase transition

A finite magnet in zero field can reverse between two equally magnetised directions. Its exact mean magnetisation is then zero even below the temperature where a bulk material orders. Learn how to extract order, correlation length and critical exponents from finite measurements, and which apparent crossings or peaks would survive the thermodynamic limit.

What a finite magnet can and cannot spontaneously break

Consider a ferromagnetic Ising model with spins $s_i = \pm 1$, $N = L^d$ sites, short-range coupling $J > 0$, and uniform field h in energy units. Its Hamiltonian is $-J \sum_{\langle ij \rangle} s_i s_j - h \sum_i s_i$. A finite partition sum contains finitely many positive analytic terms at positive temperature. For finite N it cannot develop a true nonanalyticity on the real positive-temperature axis. Yet derivatives can become large and measured distributions strongly bimodal as size grows.

Define total magnetisation $M = \sum_i s_i$ and intensive order $m = M/N$. At zero field, spin inversion pairs configurations of M and $-M$, making the exact finite equilibrium mean zero. To define spontaneous order in the positive direction, take the thermodynamic limit at positive h, then let $h \to 0^+$. Reversing that order gives zero. A finite record that stays on one side because reversals are rare has not performed the full finite-system equilibrium average. It can approximate a symmetry-selected state only with that interpretation stated.

The absolute magnetisation or its root-mean-square value remains positive in finite zero-field calculations. Even independent disordered spins produce $\langle |m| \rangle$ of order $N^{-1/2}$. Therefore a nonzero magnitude in one simulation cannot establish bulk order. Repeat over sizes and inspect its scaling. A weak applied field, surfaces or slow field drift can also select one sign in a laboratory, modifying the finite-size comparison.

Derive a response from connected fluctuations

Differentiate the finite partition function at fixed temperature. The susceptibility per site to energy field h is

$$\chi_L = \frac{\partial \langle m \rangle}{\partial h} = \frac{\beta}{N} \left(\langle M^2 \rangle - \langle M \rangle^2 \right) = \beta \sum_{\mathbf{r}} G_c(\mathbf{r}),$$

where the last equality assumes a translationally invariant system and $G_c(\mathbf{r}) = \langle s_0 s_{\mathbf{r}} \rangle - \langle s_0 \rangle \langle s_{\mathbf{r}} \rangle$. Long correlations amplify the response because many sites contribute together when one responds to a field. For an experiment using magnetic induction instead of the energy field, include magnetic moments and unit-system factors.

In a symmetric finite zero-field sample below the ordering temperature, transitions between the two bulk-like peaks can give a large contribution proportional to N to the full fluctuation susceptibility. To estimate within-phase fluctuations one sometimes forms $\beta N \left(\langle m^2 \rangle - \langle |m| \rangle^2 \right)$. This is a different observable and not exactly the derivative of the finite symmetric mean with respect to h. Report which definition is plotted. Mixing those definitions in a size collapse can change both amplitudes and the apparent location of extrema.

A correlation length may be estimated from the large-distance connected correlation decay away from criticality, or from the ratio of small-wavevector structure factors. In lattice units, an often used periodic

second-moment length is

$$\xi_L = \frac{1}{2\sin(\pi/L)}\sqrt{\frac{S(0)}{S(k_{\min})} - 1},$$

$$k_{\min} = 2\pi/L.$$

This relation uses the small-k expansion of the correlation structure factor and a lattice wavevector convention. Subtract means consistently, especially if estimating connected rather than symmetric correlations. If a small noisy $S(0)/S(k_{\min}) - 1$ is negative, replacing it by zero without describing the noise can bias a length. Larger samples or more suitable observables may be preferable.

Derive the finite-size variable from competing lengths

Write reduced temperature $t = (T - T_c)/T_c$. Suppose the bulk correlation length diverges as $\xi_\infty = \xi_0|t|^{-\nu}$. A finite periodic sample cannot support correlations extending arbitrarily beyond L. The competition is governed by L/ξ_∞, equivalently $tL^{1/\nu}$ with its sign and a metric factor. If the bulk magnetisation, susceptibility and heat capacity have exponents β_m, γ, α, a standard finite-size ansatz below the appropriate upper critical dimension and without extra dangerous variables is

$$\langle|m|\rangle = L^{-\beta_m/\nu}F_m(tL^{1/\nu}),$$

$$\chi_L = L^{\gamma/\nu}F_\chi(tL^{1/\nu}) + \chi_{\text{reg}}.$$

Here β_m is an order-parameter exponent, distinct from reciprocal thermal energy β. The functional forms match the bulk powers when the scaled argument becomes large at fixed nonzero t. A rounded finite-size peak of χ occurs at an approximately fixed scaled argument, so its temperature shift and width scale roughly as $L^{-1/\nu}$, allowing boundary and subleading corrections. Its height has an additional smooth background that can be significant for modest sizes.

For a two-dimensional short-range Ising universality class the known values $\nu = 1$, $\gamma = 7/4$, $\beta_m = 1/8$ give useful benchmarks. Doubling L at criticality should raise leading susceptibility by $2^{7/4} = 3.364$ and reduce order magnitude by only $2^{-1/8} = 0.917$. The very slow decrease of order is easy to mistake for nonzero bulk order at the critical temperature. The heat-capacity singularity in this example is logarithmic rather than a simple positive power; fitting a few peak heights to a power law can give a misleading small apparent exponent.

Compare a mean-field onset and susceptibility with the finite-size scaling argument

A dimensionless crossing and its systematic drift

A convenient shape diagnostic is $U_L = 1 - \langle m^4\rangle/[3\langle m^2\rangle^2]$. A zero-mean Gaussian distribution has fourth moment three times variance squared and gives zero. Two infinitely narrow equal-weight peaks at $\pm m_0$ give $2/3$. At criticality, rescaling m by $L^{-\beta_m/\nu}$ cancels between numerator and denominator, so curves for different sizes can approach an approximately common crossing. The limiting value depends on geometry, boundary conditions and convention, despite the exponent universality.

Include a leading correction $U_L = U_* + a\,tL^{1/\nu} + bL^{-\omega} + \cdots$. Equating the expansions for L and sL shows that the pair crossing drifts as $t_\times \propto L^{-(1/\nu+\omega)}$ when the coefficients are nonzero. An exactly common crossing is therefore not guaranteed in finite data. Sample temperature resolution, uncertainty of higher moments and correlations between neighbouring reweighted temperatures can also shift the visual crossing.

For a data set with sizes 16, 32 and 64, suppose adjacent pair crossings occur at reduced-temperature-like coordinates 2.31 and 2.28 when plotted as k_BT/J. Without a calibrated correction exponent, taking their simple mean gives no justified bulk critical temperature. If a model independently predicts a crossing displacement decreasing by four per doubling, solve $T_{16,32} = T_c + A$, $T_{32,64} = T_c + A/4$, giving $T_c = 2.27J/k_B$. Vary the correction assumption and use an additional larger size to test it. Extrapolated digits considerably beyond the scatter of these alternatives are not warranted.

Separate thermodynamic rounding from sampling error

Near criticality long spatial correlations often coincide with long relaxation times. A sequential Monte Carlo update can require many sweeps to produce independent configurations. If an observable has normalised autocorrelation $\rho(k)$, its effective independent sample count for a long record is approximately $N_{\mathrm{samp}}/[1 + 2\sum_{k>0}\rho(k)]$. Blocking time records and recomputing Binder ratios from blocks is more reliable than treating individual sweeps as independent.

If a susceptibility peak is represented by ten nearby temperatures generated from the same energy histogram, their errors are correlated. Temperature reweighting uses factors $e^{-(\beta'-\beta)E}$ and is dependable only where sampled energy distributions overlap. A smooth reweighted curve extending outside sampled energies can look precise while depending on a few rare configurations. Use independent anchor simulations or directly measured temperatures to verify both tails.

Inspect distributions as well as means. A first-order transition can also show a large susceptibility peak and bimodality, with leading peak scaling proportional to volume in many finite systems. A continuous transition's order distribution may itself be broad or even possess two lobes. Distinguish using latent heat or its size dependence, interface free-energy barriers, scaling of peak width and appropriate thermal observables. One finite-size double peak in magnetisation is not a universal transition-order test.

Check how a long but finite correlation length can arise without a positive-temperature transition in one dimension

An exact finite chain tests what a rapidly growing length means

A one-dimensional nearest-neighbour Ising chain in zero field has no positive-temperature bulk ordering transition, but at low temperature its correlation length grows extremely rapidly. This makes it a valuable counterexample when interpreting a finite magnet. The transfer matrix has entries $e^{Kss'}$ with $K = J/(k_BT)$. Its two eigenvalues are

$$\lambda_+ = 2\cosh K,$$
$$\lambda_- = 2\sinh K.$$

For a periodic chain of L spins,

$$Z_L = \lambda_+^L + \lambda_-^L.$$

Keeping only the larger eigenvalue is permissible in the thermodynamic free-energy density at fixed positive temperature but need not accurately represent a short cold chain.

In the infinite chain the zero-field correlation is $\langle s_0 s_r \rangle = (\tanh K)^r$, yielding $\xi = -1/\log(\tanh K)$ in lattice spacings. At low temperature $\xi \simeq e^{2K}/2$. A chain of length 100 can therefore be strongly correlated across its entire length around $K \simeq 2.65$, while the infinite chain still has zero spontaneous order at every finite K. Setting $\xi \sim L$ gives an apparent ordering scale $k_B T_L / J \simeq 2/\log(2L)$, which moves slowly toward zero as the sample is enlarged.

For a periodic finite chain there are two paths connecting site zero and site r. The exact correlation is $[q^r + q^{L-r}]/[1 + q^L]$, $q = \tanh K$. Fitting it to a single exponential out to separations near half the box would overestimate the length because the second periodic image contributes. Use the periodic expression explicitly or restrict fits to a window that is both larger than microscopic distances and smaller than half the sample. A high-quality apparent fit over a narrow middle interval can still mask the wrong boundary geometry.

Numerically simulate chains of lengths 32, 64 and 128 at a sequence of K. Compare the measured absolute magnetisation, full connected correlation and transfer-matrix predictions. If a simulation remains stuck in one sign sector its short-distance correlation can agree with theory while the time average of signed magnetisation does not. Increasing the number of Monte Carlo sweeps and increasing the physical chain length are separate limits. Following only one of them cannot establish a nonzero thermodynamic order parameter.

A response derivative offers an exponent without partition-function normalisation

Let U_L be a dimensionless shape ratio such as the Binder cumulant. If $U_L = \mathcal{U}(tL^{1/\nu}) + \cdots$, its slope with respect to reduced temperature at the critical point scales as $L^{1/\nu}$. A logarithmic slope ratio between two sizes can estimate the correlation-length exponent independently of the magnetisation or susceptibility exponents. Corrections to scaling still matter, especially if the sampled point is not near a fixed critical argument for both sizes.

An equilibrium derivative can also be computed without a finite temperature difference. For any state observable A with no explicit temperature dependence, differentiating its canonical mean gives

$$\partial_\beta \langle A \rangle = -\operatorname{cov}(A, E).$$

Apply this to the second and fourth moments entering a Binder ratio, then use the chain rule. The resulting slope uses energy–magnetisation correlations from one run. Taking two numerically close independent temperatures instead can amplify noise by division by their small separation. A covariance estimator avoids that particular subtraction but requires accurate sampling of joint tails.

Suppose measured absolute crossing slopes on sizes 24, 48 and 96 are 5.7, 10.9 and 21.5 per unit reduced temperature. Adjacent estimates are $1/\nu = \log(10.9/5.7)/\log 2 = 0.935$ and 0.980. They approach the expected value one in this benchmark. A single fit through all sizes would give a smaller formal statistical uncertainty but might obscure a systematic drift. Report adjacent effective exponents and test sensitivity to omitting the smallest size.

The numerical derivative must use the same temperature variable for each size. A slope against inverse temperature differs by $d\beta/dT$, a known factor near a common critical temperature; plotting against

independently fitted centre temperatures can introduce an additional calibration dependence. A nonlinear temperature sensor or a slowly drifting coupling also changes the physical conversion from laboratory temperature to the simulated dimensionless control. Use an independently determined parameter scale if the goal is to compare experiment with an absolute transition temperature rather than just a scaling exponent.

Estimate an interface cost from a probability trough

At a first-order transition, a finite periodic sample may switch between two phases separated by an interfacial configuration. A histogram of an appropriate extensive observable can show peaks corresponding to nearly pure phases and a trough representing intermediate states. At equal bulk phase weight, a probability ratio gives a free-energy difference $\Delta F_L = -k_B T \log(P_{\mathrm{trough}}/P_{\mathrm{peak}})$, provided the same bin measure or a properly defined density is used. An exponentially small trough amplifies sampling errors, so the measured minimum is often only an upper bound on its probability.

In a slab geometry under periodic boundaries, an intermediate configuration may contain two approximately flat interfaces. If each interface costs γL^{d-1}, the leading barrier is $2\gamma L^{d-1}$, with subleading capillary fluctuations, geometry corrections and possible droplet-to-slab crossovers. A cube with open boundaries can instead permit interfaces of different number or shape. Counting two interfaces automatically for every order-parameter histogram would misestimate the inferred tension.

For an illustrative two-dimensional periodic square, suppose the measured barrier grows from $8k_BT$ at $L = 40a$ to $15k_BT$ at $L = 80a$, where a is lattice spacing. From the difference of the two barriers and the increase in total interface length $2(40a)$, an effective line tension is $0.0875\,k_BT/a$. Using the absolute smaller-size barrier divided by $2L$ gives $0.100\,k_BT/a$, revealing a finite-size offset. More sizes, an observed slab geometry and adequate trough sampling are needed before assigning a bulk interface tension.

This analysis should not be transplanted unchanged to a continuous critical distribution with two shoulders in finite volume. A critical finite-size order distribution can be non-Gaussian without an interfacial barrier growing as L^{d-1} at its critical point. Plot the barrier's size dependence and inspect actual spatial configurations. A peak-height ratio alone contains no direct information about how many interfaces or what type of correlations generated the intermediate states.

Estimate a ratio of moments without losing its covariance

The second and fourth magnetisation moments in a Binder cumulant usually come from the same record. Treating their errors as independent can grossly misestimate the uncertainty because the fourth moment tends to be large precisely in configurations contributing a large squared moment. Divide the record into blocks longer than its integrated autocorrelation time, recompute the ratio after deleting each block and estimate uncertainty using a block jackknife or use independent replicate simulations. Recomputing the nonlinear statistic inside each resample retains numerator–denominator covariance.

Histogram reweighting across nearby temperatures should likewise be recomputed within each block when extracting a crossing. Fitting smooth curves to the full sample and then adding independent random errors to individual temperature points ignores strong correlations induced by the shared configuration weights. Repeat the root-finding step for each independent or block-resampled pair of systems to obtain an uncertainty on the crossing temperature itself.

Check block-size convergence. If estimated errors continue growing as block duration is increased, slow modes may not be resolved or the record may have drifted. A huge number of recorded spin configurations cannot compensate for fewer than a handful of independent reversals when the statistic depends on those reversals. This issue is measurable: compare direct covariance estimates, block estimates and independent starts, and extend the run where those diagnostics disagree.

A block-average variance exposes a correlation length without fitting a tail

Divide a large zero-field spin image into square blocks of side b and record the average spin in each block. Its variance is the sum of all connected correlations within the block divided by the squared number of sites. For independent spins of variance one, it is exactly $1/b^2$ in two dimensions. If neighbouring spins are correlated over a length ξ much smaller than the block, the variance still falls approximately with inverse block area, but with a larger prefactor proportional to the integrated correlation. Below block sizes comparable to ξ, this averaging becomes much less effective.

At a critical point the connected correlation decays algebraically rather than exponentially over the accessible range. If in dimension d it behaves as $G(r) \propto r^{-(d-2+\eta)}$, integrating pair separations to the block scale gives $\operatorname{var}(m_b) \propto b^{-(d-2+\eta)}$ in the scaling regime. For the two-dimensional short-range Ising benchmark with $\eta = 1/4$, doubling the block side reduces the variance by only $2^{-1/4} = 0.841$, compared with 0.25 for independent spins. This is an observable distinction even when a full bulk thermodynamic limit is unavailable.

Block analysis of one image does not generate independent thermodynamic systems. Neighbouring blocks share correlated surroundings, overlapping blocks are statistically dependent, and every block inherits global constraints or drift in the original preparation. Increasing block size also reduces the available count sharply. Repeat across independent configurations, compare several whole-image sizes and study block-size windows well above a lattice spacing and well below the full image. An apparently accurate critical exponent from the final two largest blocks of one image is particularly unreliable.

A real magnetic image may measure a blurred field projection rather than the local spin itself. Spatial point spread creates detector-induced correlations at small separations. Determine it with a reference pattern and omit or model blocks comparable to the blur. Conversely spatially varying illumination adds long-range covariance unrelated to the sample. Flat-field calibration and sample rotations can separate these artefacts from material correlations. A second-moment structure-factor estimate and a block-variance estimate respond differently to these measurement effects, providing a useful cross-check.

This method also links microscopic and macroscopic noise. The scaled block variance in the large-block regime approaches the same integrated connected correlation controlling uniform susceptibility, with boundary terms that shrink as block dimensions exceed ξ. Comparing a weak-field response with block statistics therefore tests equilibrium fluctuation–response and spatial calibration simultaneously. If the discrepancy grows only as blocks approach the full system size, a fixed total magnetisation constraint or incomplete sampling of global sign reversals deserves examination.

For a concrete noise budget, consider a camera averaging adjacent pairs of independent sites before reporting them as displayed pixels. Those displayed pixels can have a positive nearest-neighbour correlation if their averaging footprints overlap, even when the original spins are completely uncorrelated. A measurement of the apparent short-range correlation length would then reflect the instrument kernel. Blocking on sizes much larger than the overlap can recover the correct large-scale inverse-area trend, while

direct comparison of single-pixel variance with independent-spin theory fails. Calibrating only the detector's global gain cannot fix this spatial response. A measured point-spread map, an uncorrelated reference preparation and tests of blocks at several image resolutions can separate the resolution scale from the material's growing correlation length as temperature is varied.

Compare a correlation length with a finite specimen

Determine how the subleading transfer-matrix eigenvalue controls the length seen in a colour measurement. A finite chain much shorter than this length can look ordered.

Open this experiment online

Try the calculation

Reduced temperature kBT/J **1**

Reduced colour field H/J **0.15**

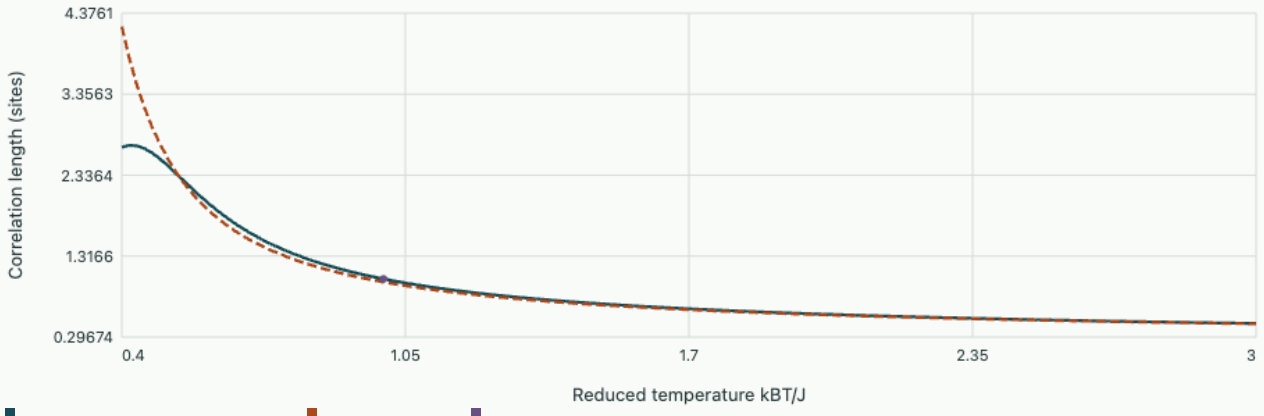

Leading eigenvalue 4.9881, second 1.8884, longest length 1.0295 sites.

Download the OneSimulation model

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

```
python spectra-and-orbits.py chain --temperature 1 --bias 0.15
```

Tutorial · Extract two exponent ratios from a size doubling

At a candidate critical temperature susceptibilities for sizes 32 and 64 are 110 and 360 in common units. Absolute magnetisations are 0.310 and 0.282. Estimate the corresponding exponent ratios and compare with the two-dimensional Ising benchmarks.

▼ **Use logarithmic ratios and inspect the small differences**

The susceptibility ratio gives

$$\gamma/\nu = \log(360/110)/\log 2 = 1.71.$$

For magnetisation $\beta_m/\nu = -\log(0.282/0.310)/\log 2 = 0.137$. These are near 1.75 and 0.125, but no error bars or corrections have been supplied. An additive susceptibility background or a slight critical-temperature offset can bias both effective exponents. Use several sizes and a dimensionless crossing diagnostic before assigning a universality class.

Tutorial · Determine whether a finite sample really switched phases

A zero-field record contains 20,000 sweeps of positive magnetisation near $+0.6$, followed by one reversal and 20,000 sweeps near -0.6. It yields a near-zero mean and large variance. Why does the nominal record length not establish a precise symmetric susceptibility? What extra records or controls would help?

▼ **Identify the slow event determining the uncertainty**

Within each sign sector there may be many locally decorrelated fluctuations, yet the relative weight of the two peaks is estimated from essentially one reversal. Treating 40,000 sweeps as independent underestimates the error of the peak weights and of statistics sensitive to global sign. Starting multiple independent runs with both signs can improve coverage, but deliberately balancing starting signs is not a measurement of true equilibrium weights if a weak residual field exists.

Record switching-time distributions over much longer runs or use an equilibrium-preserving algorithm that efficiently samples both sectors, while distinguishing computational updates from physical kinetics. Apply small positive and negative fields to test susceptibility and residual bias. For a within-phase susceptibility condition on a clearly defined phase and report the changed observable.

Tutorial · Explain a crossing that drifts with the boundary

Two groups study the same lattice Hamiltonian. One uses square periodic samples; the other uses twice-as-long rectangular samples with open ends. Their Binder curves cross near similar temperatures but at different cumulant values. Is this evidence against universal critical behaviour?

▼ **Separate universal exponents from finite-geometry functions**

Changing aspect ratio and boundary conditions changes the finite-size scaling function and hence the crossing value of a dimensionless ratio. Universal bulk exponent comparisons do not require identical Binder values across different finite geometries. Match boundary conditions, shape, operator and cumulant convention before comparing purported universal amplitude ratios.

Each group can still extrapolate its own crossing temperatures, examine moment slopes proportional to $L^{1/\nu}$ and test bulk exponent ratios with geometry fixed across its own size series. If one dimension is held fixed while the other grows, however, the family approaches a strip rather than a bulk two-dimensional limit. The thermodynamic sequence itself is an experimental variable, not just a plotting preference.

Cambridge Statistical Physics · phase transitions, criticality and scaling

Statistical mechanics from microscopic counting to material response · 11. Phase nucleation, metastability and waiting-time statistics

A cooled liquid can remain liquid below its equilibrium freezing temperature and then transform abruptly. How can a bulk phase be thermodynamically favoured while a small embryo of it tends to shrink? Balance interfacial cost against bulk gain, convert a barrier into a probabilistic waiting-time prediction and test which measurements can distinguish surface nucleation from a bulk mechanism.

Choose a driving free-energy density at fixed controls

At specified temperature and pressure, phases in equilibrium have equal molar Gibbs free energies. Below a melting temperature T_m the stable solid has lower chemical potential than the liquid. For modest undercooling, use $d\mu = -s\,dT + v\,dp$. At fixed pressure the chemical-potential difference per molecule is approximately $\mu_l - \mu_s \simeq \Delta h_f (T_m - T)/T_m$ if latent enthalpy of fusion per molecule varies little. Multiplying by the number density of the new phase gives a positive driving free-energy density Δg. State whether each quoted latent heat is per particle, per mole or per unit volume; converting it twice gives enormous radius errors.

The underlying free-energy landscape has stable and metastable minima separated by a barrier. In a phase-field picture a local bulk density penalises intermediate order within the metastable region, and a gradient penalty charges a finite energy for an interface. A finite amount of stable phase can therefore cost free energy before its bulk advantage wins. A coexistence construction determines which bulk phases minimise free energy; it does not provide the rate at which an existing metastable preparation can reach them.

Compare a bulk coexistence construction with the cost of an interface

Derive the critical radius and barrier

In the simplest capillarity model an embryo is a sphere of radius r, with isotropic positive interface energy γ per area and a bulk gain Δg per volume. Treat its interface as thin and its bulk properties as those of the macroscopic phases. Neglect elastic strain, long-range fields and depletion of the surrounding reservoir, then the reversible work to create it is

$$\Delta G(r) = 4\pi r^2 \gamma - \frac{4\pi}{3} r^3 \Delta g,$$

$$r_* = \frac{2\gamma}{\Delta g},$$

$$\Delta G_* = \frac{16\pi\gamma^3}{3\Delta g^2}.$$

The derivative vanishes at r_* and the second derivative there is $-8\pi\gamma$, so this is a maximum along the size coordinate. A slightly smaller embryo tends to shrink and a slightly larger one tends to grow, subject to noise and kinetic constraints. The radius where the total formation free energy first drops below its initial zero is

$$3\gamma/\Delta g = 1.5 r_*.$$

A supercritical nucleus can grow downhill even while its cumulative creation free energy is still positive. Confusing these two radii yields a fifty percent size error.

Take illustrative values $\gamma = 0.010$ J per square metre and $\Delta g = 1.0 \times 10^{7}$ J per cubic metre. Then $r_{*} = 2.0$ nm and $\Delta G_{*} = 1.68 \times 10^{-19}$ J, about 40.5 thermal energies at 300 K. Doubling driving density halves critical radius and quarters the barrier if γ is unchanged. At a two-nanometre radius the embryo may be just a few molecular spacings across. This calculated barrier is a sensitivity estimate; the assumed bulk interface tension, sharp interface and spherical shape become increasingly questionable there.

An alternative derivation uses the pressure jump $2\gamma/r$ across a curved interface together with a bulk chemical-potential difference. Near coexistence the condition of zero growth recovers the same critical curvature. In a vapour close to ideal behaviour the driving chemical-potential difference involves $k_B T \log S$ with supersaturation $S = p/p_{\mathrm{sat}}$, not simply $S - 1$ at arbitrary supersaturation. Near $S = 1$ the two are close. A quantitative cloud-chamber inference must also account for the evolving supersaturation and evaporation or growth after nucleation.

Background · Why does a faster temperature ramp move the apparent freezing point?

Assume independent homogeneous nucleation with rate density $J(T)$ in volume V, and prompt detection of the first nucleus. The chance that the untransformed sample survives the next interval dt is $1 - V J(T)\, dt$. Multiply these conditional probabilities along a cooling ramp with constant positive speed $v = -dT/dt$. Starting at T_0 gives

$$S(T) = \exp\!\left[-\frac{V}{v}\int_{T}^{T_0} J(T')\, dT'\right].$$

At fixed undercooling the faster ramp has accumulated less nucleation probability. The median observed freezing temperature satisfies an integrated hazard of $\ln 2$. Doubling the sample volume and the ramp speed together leaves the predicted distribution unchanged. Failure of that scaling can signal surface-dominated nucleation, thermal lag or a growth time long enough to affect detection.

Derive the barrier reduction for a spherical cap

A nucleus forming on a flat inert wall can replace one existing interface with another. Let it be a spherical cap of curvature radius r and contact angle θ through the new phase. The liquid-facing area is $2\pi r^2(1 - \cos\theta)$, contact area $\pi r^2 \sin^2\theta$ and cap volume $\pi r^3(2 - 3\cos\theta + \cos^3\theta)/3$. Young's contact-angle condition relates the wall energy difference to minus $\gamma\cos\theta$ per replaced contact area. Combining all terms gives the homogeneous spherical result multiplied by

$$f(\theta) = \frac{(2 + \cos\theta)(1 - \cos\theta)^2}{4}.$$

Both surface and bulk coefficients receive the same factor. Critical curvature radius therefore remains $2\gamma/\Delta g$ in this simple model, while the barrier becomes $f(\theta)\Delta G_{*}$. At 180°, $f = 1$ as the full sphere barely touches the wall. At 90°, $f = 1/2$ for a hemisphere. At 60°, $f = 0.15625$, lowering the illustrative 40.5-thermal-

energy barrier to 6.33 thermal energies. Even a relatively small number of active wall sites can dominate observed events.

Complete wetting, a chemically patterned surface, a very small defect or appreciable energy per contact-line length can invalidate the simple contact geometry. Contact angle measured on a macroscopic equilibrium droplet need not equal the local microscopic value for a tiny forming embryo. A fitted effective angle that varies systematically with undercooling can signal a missing mechanism rather than a temperature-dependent property of a clean ideal wall.

Distinguish an attempt rate from an equilibrium barrier

A classical rate expression is $J = J_0 e^{-\Delta G_*/(k_B T)}$ per unit volume per unit time for homogeneous events. The prefactor contains the supply of available configurations and the kinetics for traversing the critical region. A barrier alone is insufficient to predict a waiting time. Lowering temperature can raise thermodynamic driving yet also reduce mobility and attachment, giving a nonmonotone total rate as the two effects compete.

If a well-mixed sample of volume V has constant independent hazard JV, the survival probability of the metastable state follows from $P_0(t+dt) = P_0(t)(1 - JV dt)$: $P_0 = e^{-JVt}$. Mean wait is $1/(JV)$, median $\log 2/(JV)$. A 1.0 microlitre sample with $J = 2.0 \times 10^{8}$ per cubic metre per second has hazard 0.20 per second and median 3.47 seconds. After five seconds 63.2 percent have transformed at least once, if the detector registers nucleation immediately.

If surfaces contribute independent rate per area J_s, the combined first-event hazard is $J_v V + J_s A$. Geometrically similar droplets have $A \propto V^{2/3}$, so changing sample size can distinguish these contributions. Vary active surface chemistry at fixed geometry as a second probe. If most samples contain a single rare defect not scaling smoothly with area, survival can be a mixture of exponentials rather than one exponential. A straight-line fit to mean waits alone obscures such heterogeneity.

A cooling ramp integrates many different hazards

Under a deterministic temperature history $T(t)$, survival is $P_0(t) = \exp[-\int_0^t J(T(s)) V ds]$ for homogeneous, conditionally independent events. In a constant cooling ramp $q = -dT/dt > 0$, the exponent becomes $-(V/q)\int_T^{T_0} J(T')dT'$. The observed freezing-temperature histogram is weighted by survival as well as the instantaneous rate. The most probable freezing temperature need not minimise any free energy, and different ramp speeds can shift it even when all material properties are unchanged.

For a ramp calculation on gridded temperatures, update the remaining fraction by $e^{-\Delta H}$ where ΔH is the integrated hazard in that interval. Using a transition probability $JV\Delta t$ even when it exceeds one violates probability bounds. Subdivide the interval and check convergence of median and distribution width. At steep temperature dependence, a midpoint estimate over a coarse interval can move a reconstructed barrier enough to dominate the apparent contact-angle uncertainty.

The detected event may be a threshold in light scattering, latent heat or visual crystal size. A nucleus must grow before producing that signal. At fixed growth speed u_g, reaching detection radius r_d adds an approximate lag $(r_d - r_*)/u_g$; under cooling, both growth and hazard vary. Compare multiple detection thresholds or resolve the early growth trajectory and extrapolate a starting time. A shift in waiting time due entirely to the imaging threshold could otherwise be misinterpreted as altered microscopic nucleation.

Test how changing undercooling alters a barrier in the course's phase-equilibrium example

A finite reservoir can stabilise a droplet that would otherwise grow

The standard critical-radius derivation supposes that the driving free-energy density remains fixed while the nucleus grows. In a small sealed droplet or isolated simulation box, formation of the new phase can consume a significant fraction of the supersaturation. The surrounding phase then approaches coexistence and the bulk driving force falls. A single critical radius computed from initial supersaturation may describe early stability but cannot describe the entire late-time growth or the final phase fraction.

As a transparent model let conversion of new-phase volume v reduce the driving density linearly to $\Delta g(v) = \Delta g_0 - a v$, with $a > 0$ incorporating finite-reservoir compressibility or depletion. The bulk free-energy gain is the integral of the driving density, not its final value times volume. For a spherical interface $A(v) = (36\pi)^{1/3} v^{2/3}$, the total model is

$$\Delta G(v) = \gamma (36\pi)^{1/3} v^{2/3} - \Delta g_0 v + \frac{a}{2} v^2.$$

The initial derivative is positive because a very small sphere has large area per volume. At large v the positive depletion derivative can make further growth unfavourable. If the derivative turns negative at intermediate volume there can be both a barrier top and a larger stable or metastable finite-droplet minimum. If it never turns negative the finite reservoir is too small or too weakly driven for a stable droplet in this model. This explains why a finite system can show a rare nucleus followed by arrested growth without contradicting the initial local barrier idea.

For a dimensionless illustration choose $\Delta G/(k_B T) = 4v^{2/3} - 4v + 0.25 v^2$, with a specified arbitrary reference volume absorbed into v. Its derivative is $(8/3) v^{-1/3} - 4 + 0.5 v$. At $v = 0.10$ it is about 1.80, at $v = 1$ about minus 0.83 and at $v = 10$ about 2.24. At least two stationary sizes therefore lie along the growth path. A sign test on neighbouring sizes distinguishes the early maximum from the later minimum without assuming every stationary root is a critical nucleus. Numerical root finding should bracket these branches separately instead of passing a single initial guess to a solver.

An experiment can test depletion by varying reservoir size at the same initial supersaturation and surface treatment. If early appearance times are comparable after correcting for trial volume but final crystal sizes grow with available material, finite feeding is plausible. Limited thermal transport, impurities blocking the interface or elastic stress can also arrest growth, so monitor the remaining parent-phase composition and temperature. A small system is precisely where the thermodynamic ensemble and the available amount of material need explicit definition.

A detection time can be dominated by diffusion after nucleation

After crossing the barrier a supercritical spherical domain may grow by diffusion of a dilute species. In a quasi-steady infinite exterior with surface concentration c_s and far concentration $c_\infty > c_s$, the concentration solves the radial Laplace equation. Its solution is

$$c(r) = c_\infty - (c_\infty - c_s) R / r.$$

The total inward solute flux is $4\pi D R(c_\infty - c_s)$. Equate this to solute incorporated into an advancing sphere, $4\pi R^2 (c_{\mathrm{new}} - c_s)\dot{R}$, to obtain

$$R^2(t) - R^2(0) = 2D\frac{c_\infty - c_s}{c_{\mathrm{new}} - c_s}t.$$

This square-root growth law assumes quasi-steady diffusion, a spherical shape, nearly constant far conditions and a surface near local equilibrium. If interface attachment instead limits growth, a roughly constant $\dot{R}$ may be a better first model over a short range. Measuring time to several optical thresholds distinguishes R_d^2 scaling from R_d scaling. Threshold variation therefore separates growth and nucleation more effectively than one image definition held fixed.

Take $D = 1.0 \times 10^{-9}$ square metres per second and concentration ratio 0.020. From a negligibly small initial radius, reaching 5.0 micrometres takes $25 \times 10^{-12}/(4.0 \times 10^{-11}) = 0.625$ s; reaching ten micrometres takes 2.50 s. If true median nucleation time is 1.0 s, the detected median at the larger threshold can be 3.5 s even when the microscopic barrier and initial supersaturation are unchanged. Calling its inverse a nucleation rate would seriously underestimate the true first-event hazard.

At very small radius a Gibbs–Thomson curvature shift increases the equilibrium surface concentration. The same curvature effect that opposes a small embryo can therefore slow its initial excess growth even after it has crossed the most probable barrier. At late times the bath may not be infinite and growing domains compete for material. These effects change the simple square-root law in different regions. Image a radius trajectory for each event instead of assigning the same unverified growth lag to all samples.

A thermal ramp with a tractable hazard shape

Over a restricted undercooling interval an experimentally measured hazard can rise approximately exponentially with undercooling $\Delta T = T_0 - T$: take $h(\Delta T) = h_0 e^{a\Delta T}$. This is a local fit to the full barrier and prefactor, not a universal replacement for their pronounced temperature dependence. At constant cooling rate $q > 0$, direct integration gives cumulative hazard $H = h_0(e^{a\Delta T} - 1)/(aq)$ and survival e^{-H}. The median solves

$$\Delta T_{1/2} = \frac{1}{a}\log\left[1 + \frac{aq\log 2}{h_0}\right].$$

For $h_0 = 0.010$ per second, $a = 1.0$ per kelvin and $q = 0.10$ K/s the median undercooling is 2.07 K. Doubling cooling rate raises it to 2.70 K. The shift is smaller than one kelvin but large relative to many thermometer precisions. It arises even with exactly the same thermodynamic coexistence temperature and material interfaces. Counting a median freezing temperature at only one ramp rate can therefore mix kinetics with the equilibrium phase boundary.

The probability density per unit undercooling is $[h(\Delta T)/q]e^{-H}$, where the first factor measures the rate conditional on surviving to that temperature. Taking a log derivative shows that an interior density maximum occurs when $h = aq$, provided this exceeds the initial hazard. The median and the most probable temperature are distinct. At slower ramps the maximum can even lie at the beginning of the observation range, depending on preparation and chosen T_0. A histogram peak should be interpreted together with the surviving population, not treated as a direct plot of microscopic rate.

To reconstruct hazard from a measured ramp, estimate the fraction still untransformed at each temperature and inspect minus the log of that fraction. Differentiate smooth cumulative hazards cautiously; differentiation amplifies noise and small late survivor counts. A joint fit of raw first-event and censored records over several ramp rates generally preserves more information than differentiating a single binned histogram. Varying sample volume as well helps separate a distribution of bulk rates from one associated with surfaces.

A committor test examines whether radius is a good reaction coordinate

The capillarity free-energy maximum uses one coordinate, radius. A practical microscopic test begins from candidate configurations near the proposed critical size, randomises thermal momenta or bath noise consistently with the preparation and launches many short independent trajectories. For each starting configuration estimate the fraction reaching a specified stable new-phase basin before returning to the parent phase. This fraction is the forward committor. A good transition-surface coordinate should identify configurations with committor near one half under the specified dynamics.

If configurations at the same radius show committors ranging from almost zero to almost one, that radius alone poorly locates the dynamical transition state. Shape, local order, strain or neighbour concentration may supply missing coordinates. Binomial uncertainty from finite trials also matters: 25 independent launches near one half give a standard error around 0.10 for an individual configuration. Increasing launches distinguishes true variation between configurations from sampling noise.

The committor depends on physically chosen basins and dynamics. An artificially accelerated Monte Carlo kinetic rule may preserve equilibrium free energies while changing which paths reach the new phase first. Use the same dynamics when comparing committor tests with measured waiting times. Equilibrium barrier heights, kinetic attempt rates and the adequacy of a proposed reaction coordinate are connected questions but require different data.

Distinguish first appearance from transformed volume fraction

A first-event detector stops caring about a sample once the earliest nucleus is registered. A calorimeter or diffraction measurement instead often follows the total fraction of material transformed. Even if nucleation is homogeneous at a constant rate and crystals grow at constant speed, these two signals have different time dependences. Treat an infinite initially untransformed sample, with nuclei placed randomly at rate J per volume and spherical domains growing radially at speed u. To ask whether a particular point remains untransformed at time t, count possible earlier nucleation events whose growth sphere would have reached it.

A nucleus born at time s reaches points within radius $u(t-s)$, corresponding to volume $4\pi u^3(t-s)^3/3$. Integrate the attempt rate over that backwards influence volume from zero to t. For an ideal Poisson space-time nucleation model with this growth rule the probability of no such event is $\exp[-\pi J u^3 t^4/3]$, giving transformed fraction $X = 1 - \exp[-\pi J u^3 t^4/3]$. The result uses the standard extended-volume construction with random nucleation and isotropic constant-rate growth. It is not the simple first-event law $1 - e^{-JVt}$ for a finite observation sample.

For sites all pre-existing at time zero with random density n_0, the same backwards-volume reasoning gives

$$X = 1 - \exp[-(4\pi/3) n_0 u^3 t^3].$$

A log plot of minus the log remaining fraction against log time therefore has ideal slope four for continual three-dimensional nucleation and three for instantaneous site saturation. In two dimensions the powers would both be reduced by one. A measured slope between these values is not automatically evidence for a fractional geometric dimension. Changing growth law, active-site exhaustion, temperature ramps or overlapping different mechanisms can alter the apparent slope.

If measured isothermal transformed fraction is one half at 100 s in the continuous model, the product Ju^3 equals $3\log 2/(\pi\,100^4) = 6.62 \times 10^{-9}$ per second to the fourth. This identifies the product, not the nucleation rate and growth speed separately. Independent imaging of isolated domain radii versus time, or a set of first-event measurements in small known volumes, can fix the remaining scale. A fit to the entire same ideal fraction curve cannot identify two quantities that enter it only through their product.

Actual calorimetric power is approximately latent enthalpy times dX/dt, convolved with thermal response. Its maximum occurs when transformation rate peaks rather than when the first nucleus forms or exactly when half the volume has transformed. Compare integrated calorimetric energy with the final transformed fraction and deconvolve the calibrated instrument response before interpreting any schedule of peaks as a nucleation-time distribution.

The influence-volume argument also predicts a useful geometry check in a thin film. If domains reach both film surfaces early and then grow as expanding disks in the plane, a three-dimensional spherical growth model will eventually overestimate newly transformed volume. The apparent exponent can cross from an initially three-dimensional value to a lower late-time value without any change of microscopic nucleation pathway. Compare films of several thicknesses, identify the time at which individual radius reaches the film thickness and then test the predicted crossover in transformed fraction. Altering film thickness usually changes wall-to-volume ratio as well, so site density and growth geometry can vary together. Independent first-event measurements and direct images of individual domains remain necessary if the objective is to separate those two effects.

Locate a critical cluster and its barrier

Vary surface energy and bulk driving force independently. Find the stationary radius before reading the height of the barrier.

Open this experiment online

Try the calculation

Interfacial energy (mJ/m²) 35

Bulk driving density (MJ/m³) 25

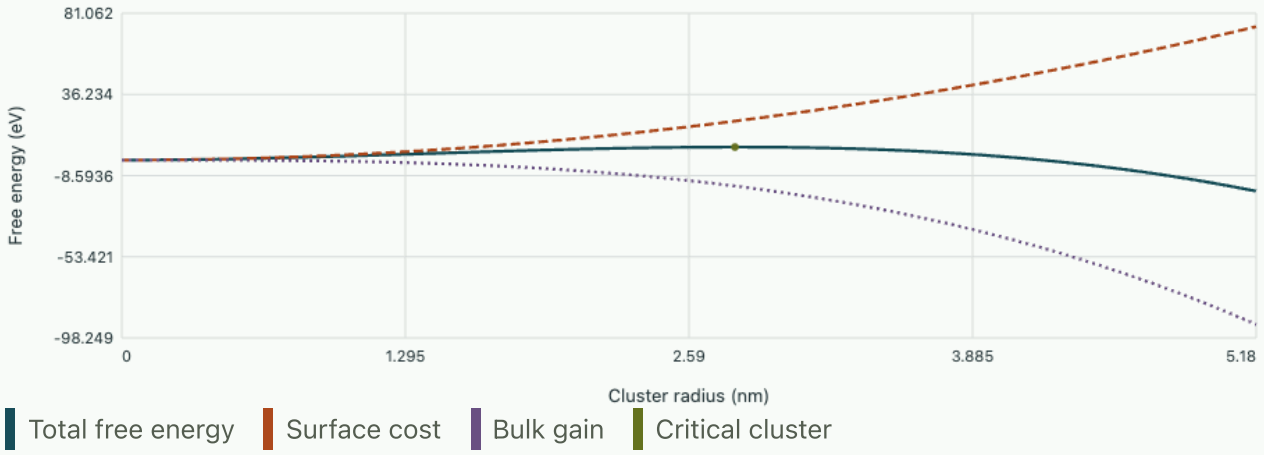

Critical radius 2.8 nm; barrier 7.174 eV. The free-energy curve crosses zero again at 4.2 nm.

Tutorial · Extract a rate from records with censored observations

Twelve equal 2.0 microlitre samples are held at constant temperature for at most 20 s. Eight show first events at times whose sum is 64 s; four remain untransformed at the end. Estimate the homogeneous rate and explain the bias from averaging only the observed event times.

▼ **Count events and sum the full time at risk**

Likelihood for independent exponential waits with right censoring is proportional to $h^{N_e} e^{-hT_{\text{exposure}}}$. The exposure is $64 + 4(20) = 144$ s, so $\hat{h} = 8/144 = 0.0556$ per second. Dividing by 2.0×10^{-9} cubic metres gives $J = 2.78 \times 10^{7}$ per cubic metre per second. The observed-only mean of 8 s would instead imply 0.125 per second and bias the rate upward by discarding samples known to survive a long time. With eight events, statistical uncertainty is substantial despite the precise arithmetic.

Tutorial · Distinguish area from volume nucleation

Two cubic samples have side lengths L and $2L$. Their measured first-event hazards are 0.060 and 0.360 per second. Assume homogeneous bulk and wall contributions $b = J_v L^3$, $s = J_s 6L^2$ in the smaller cube. Infer their relative importance and predict the hazard for a cube of side $3L$.

▼ **Use different geometric powers to separate the contributions**

The measurements give $b + s = 0.060$ and $8b + 4s = 0.360$. Subtracting four times the first equation gives $4b = 0.120$, hence $b = s = 0.030$ per second. For side $3L$, hazard is $27b + 9s = 1.08$ per second. This prediction requires the same temperature, wall state and sample history, and negligible detection lag. An

unmodelled population of corner-active sites adds a term independent of L, so measurements at a third size test the two-term assumption rather than merely improving its fit.

Tutorial · Propagate interfacial uncertainty into a rate

Suppose the barrier is estimated at $35 k_B T$, with five-percent relative uncertainty in interfacial tension while driving force, temperature and prefactor are treated as known. Estimate its effect on log rate and explain why the rate uncertainty is strongly asymmetric when expressed in ordinary units.

▼ **Differentiate the barrier first and exponentiate last**

The barrier scales as γ^3, so its first-order relative uncertainty is 15 percent. The uncertainty in $\log J$ is therefore about $35(0.15) = 5.25$, corresponding to multiplicative factors of roughly $e^{5.25} = 191$. A symmetric error in interfacial tension need not produce a symmetric additive error in rate. Computing the exact cubed upper and lower tensions produces further asymmetry.

Repeat the inference with a driving-force uncertainty, whose logarithmic coefficient is minus two in the barrier. Tension and driving parameters estimated from common coexistence data may be correlated. A direct comparison of predicted log rates across several undercoolings is often more stable than comparing raw rates spanning many orders of magnitude.

Cambridge Statistical Physics · metastability, phases and transitions

Statistical mechanics from microscopic counting to material response · 12. Work fluctuations and equilibrium free-energy differences

Repeated rapid changes of a laboratory control can give a broad distribution of work even when every run follows the same prescribed protocol. Derive why a particular exponential average retrieves an equilibrium free-energy difference, why the average work is generally larger, and why a finite experiment can miss the rare trajectories that make the equality true.

Specify the system, preparation and measured work

Let a classical system have Hamiltonian $H(z,\lambda)$, with microstate z and an externally set parameter λ. Initially it is sampled from the canonical equilibrium distribution at λ_A and temperature T. A prescribed protocol $\lambda(t)$ carries it to λ_B. Work done on the system along one trajectory is the inclusive parameter work

$$W=\int_0^{t_f}\dot{\lambda}\,\partial_\lambda H(z(t),\lambda(t))\,dt,$$

$$\Delta F=F(\lambda_B,T)-F(\lambda_A,T).$$

If λ is a trap centre, the generalised force in this convention is the derivative of energy with respect to that centre. A transducer reporting the opposite force exerted by the system on the actuator needs a sign change. For a stochastic subsystem in contact with a heat bath, a change of subsystem energy also includes heat. It is not generally equal to the measured protocol work. Specify the tip, trap, sample and bath boundary before turning a force trace into a work distribution.

Every run must begin from the same equilibrium preparation, not simply from whichever end state the previous protocol produced. The final state immediately after a fast drive need not be equilibrium. The free energy being inferred refers to equilibrium endpoint Hamiltonians at the same temperature. Equilibrating after the drive exchanges heat but involves no additional protocol work when the control is held fixed. Changing control during that relaxation would define a different protocol.

A Hamiltonian proof using phase-space volume

First treat an isolated driven Hamiltonian system after canonical preparation. Work along a trajectory equals $H_B(z_f)-H_A(z_i)$. Multiply $e^{-\beta W}$ by the initial canonical weight and integrate over initial phase space. The dependence on initial energy cancels. Hamiltonian dynamics preserves phase-space volume and is invertible, so change variables from initial to final point to obtain

$$\left\langle e^{-\beta W}\right\rangle_A=\frac{1}{Z_A}\int dz_i\,e^{-\beta H_B(z_f)}=\frac{Z_B}{Z_A}=e^{-\beta\Delta F}.$$

Thus the Jarzynski equality follows without assuming the driving is slow. This proof requires the mapping to cover the relevant final phase space with the stated measure; idealised instantaneous removal of infinitely high constraints can create subtleties of support. For a weakly coupled thermal environment include bath variables in the Hamiltonian derivation or use a stochastic dynamics satisfying appropriate local detailed balance. Arbitrary feedback, nonequilibrium initial preparation or an unmodelled strong bath coupling does not automatically obey this simple subsystem formula.

Jensen's inequality gives

$$\langle W \rangle \geq \Delta F,$$

with equality for a deterministically reversible work distribution in the relevant limit. An individual run may give less work than ΔF, sometimes even work of the opposite sign. This does not violate the ensemble inequality. The exponential average weights such small-work events more strongly than their raw frequency indicates. The larger the typical dissipated work, the more difficult a direct estimate often becomes.

Solve a complete finite-time example analytically

Take one overdamped particle equilibrated in a one-dimensional harmonic trap $H(x, \lambda) = k(x - \lambda)^2/2$ at $\lambda = 0$. Suddenly shift the centre to d while the particle position cannot respond. The endpoint free energy is unchanged by translation in an unbounded uniform bath. Initial x is Gaussian with mean zero and variance $k_B T/k$. During the instantaneous shift

$$W = \frac{k}{2}[(x - d)^2 - x^2] = \frac{kd^2}{2} - kdx.$$

Work is Gaussian with mean $kd^2/2$, variance $kd^2 k_B T$, and $\Delta F = 0$. Evaluating the Gaussian moment generates $e^{-\beta\langle W\rangle + \beta^2 \operatorname{var} W/2} = 1$, verifying the equality exactly at arbitrary shift distance. The mechanism is visible in the trajectories: particles already on the downstream side can lose potential energy in the sudden control shift, contributing negative work and disproportionately large exponential weights.

For $d = 2\sqrt{k_B T/k}$, the mean work is two thermal energies and the standard deviation is also two thermal energies. The probability of negative work is about 0.159. For $d = 5\sqrt{k_B T/k}$, mean work is $12.5 k_B T$, standard deviation $5 k_B T$, and negative work occurs only about 0.0062 of the time. Typical runs can all show substantial positive work even though the correct free-energy change for the whole translation is zero.

The most important work values for the exponential average are not simply all values below zero. Multiplying a Gaussian $P(W)$ by $e^{-\beta W}$ shifts its peak to $\langle W \rangle - \beta \operatorname{var} W$, which for this trap shift is $-\langle W \rangle$. In the five-width shift example that important region is five standard deviations below the ordinary work mean. Observing a few negative work values can still leave the estimate badly biased if the most heavily weighted part remains unsampled.

Compare forward and reverse trajectory probabilities

Define the reverse protocol by reversing the order of control values, and reverse momenta and any time-odd external parameters as required by microscopic time reversal. Prepare its initial system independently in equilibrium at B. Under microscopic reversibility the probability ratio of a forward path and its time-reversed counterpart has a boundary weight from the initial ensembles and a conditional path weight tied to heat exchange. These combine to give $e^{\beta(W - \Delta F)}$. Grouping paths by work produces the Crooks relation

$$\frac{P_F(W)}{P_R(-W)} = e^{\beta(W - \Delta F)}.$$

Where the two work densities coincide, $W = \Delta F$. Integrating the relation against $e^{-\beta W}$ recovers Jarzynski. The reverse work axis must be reflected: a crossing of two histograms plotted with the same unflipped work sign generally has no such interpretation. A finite bin estimates an integral over a work interval, so discretisation and sparse counts modify a naive pointwise log-ratio fit.

For a more general measured Gaussian forward work distribution with mean $5.0 k_B T$ and variance $4.0 (k_B T)^2$, the equality predicts $\Delta F = 3.0 k_B T$ if the Gaussian law including tails is justified. With equal-variance reverse statistics, the corresponding mean of $-W_R$ is $1.0 k_B T$. The two densities cross at $3.0 k_B T$. Plotting $\log[P_F(W)/P_R(-W)]$ against W should give slope β, providing a check on independent force and temperature calibration. Good agreement over a tiny overlap range does not constrain remote tails strongly.

Use the finite estimator honestly

From independent forward works W_i, estimate ΔF as minus $k_B T$ times the logarithm of their mean exponential weight. The sample mean of weights is unbiased under correct sampling, but its logarithm is nonlinear. Jensen implies the finite-sample free-energy estimate is biased upward on average under the ideal assumptions. A single dataset can fluctuate either way. An ordinary mean-work error bar does not describe uncertainty in an exponential moment.

For Gaussian work the squared coefficient of variation of one exponential weight is $e^{\beta^2 \operatorname{var} W} - 1$. With $d = 2$ thermal widths it equals

$$e^4 - 1 = 53.6.$$

To estimate the mean weight to ten percent relative standard error by independent averaging would require roughly 5360 runs under that distribution. At five thermal widths the required count is dramatically higher. The nominal number of runs is not a sufficient convergence diagnostic: inspect the fraction of weight carried by the largest events and compare independent batches.

For numerical stability shift all dimensionless logarithmic weights by their largest value before exponentiating, then add it back after taking the logarithm. This prevents overflow without repairing missing rare events. Bootstrapping a dataset with no observations in the relevant tail cannot generate them. More gradual driving, intermediate equilibrium windows or a bidirectional estimator using forward and reverse overlap can reduce the difficulty. Each approach requires its own equilibrium preparations and sufficient overlap, not just a different final formula applied to the same deficient sample.

Discretise a real force trace without changing the work definition

For a measured optical trap centre λ_j and positions x_j, actual control updates and acquisition clocks may be staggered. If the control moves instantaneously from λ_j to λ_{j+1} at fixed x_j, its exact contribution is $H(x_j, \lambda_{j+1}) - H(x_j, \lambda_j)$. A smoothly moving control instead uses an integral of the measured generalised force, with a quadrature matched to timestamps and sampling bandwidth. A midpoint convention applied to an unresolved abrupt jump can miss a finite quadratic term.

For the harmonic trap the instantaneous work above is exactly $k[(\lambda_j + \lambda_{j+1})/2 - x_j](\lambda_{j+1} - \lambda_j)$. Taking only the initial force would give a difference $k(\lambda_{j+1} - \lambda_j)^2/2$ per jump. Over many short smooth increments this bias may vanish under refinement, but for a deliberate quench it is the entire mean work if the initial mean

displacement is zero. The distinction can be checked by reversing the jump and independently measuring the equilibrium spatial variance.

Compare irreversible work and relaxation with the course's independently defined entropy-balance experiment

Derive the dissipated work for a constant-speed pull

Between infinitely slow displacements and a sudden quench lies an exactly soluble finite-duration example. For a harmonic trap translated at constant velocity u in an overdamped bath, the mean coordinate obeys $\zeta\, d\langle x\rangle/dt = -k(\langle x\rangle - \lambda)$. Begin equilibrated at zero and set $\lambda(t) = ut$ for $0 < t < t_f$. The mean lag $\lambda - \langle x\rangle$ obeys a first-order equation and grows from zero toward $u\tau$, with $\tau = \zeta/k$. The mean force exerted by the actuator in this sign convention is $k(\lambda - \langle x\rangle)$. Integrating the product of this force with control speed gives

$$\langle W\rangle = \zeta u^2[t_f - \tau(1 - e^{-t_f/\tau})].$$

The free-energy change for translating this uniform trap is zero. At a fixed displacement $d = ut_f$, taking $t_f \ll \tau$ gives $kd^2/2$, precisely the sudden-quench mean. As $t_f \gg \tau$ the main term is $\zeta d^2/t_f$, tending to zero for a quasistatic pull. Simply keeping speed fixed while taking a longer time is a different limit because the endpoint then moves farther.

The fluctuating part of $x(t)$ around its deterministic mean remains an equilibrium Ornstein–Uhlenbeck process in this additive linear model. Work is a linear functional of that Gaussian process and is therefore Gaussian. Its variance follows by integrating its covariance twice, giving $\operatorname{var} W = 2k_BT\langle W\rangle$. This recovers the exponential equality at every pull duration and supplies a nontrivial benchmark for a stochastic integrator. The probability of extracting work can be computed from the resulting Gaussian tail without assuming a second ensemble.

For a pull of two thermal widths completed in one relaxation time, use $\zeta d^2/t_f = kd^2 = 4k_BT$. The bracketed dimensionless fraction is $1 - (1 - e^{-1}) = e^{-1}$, so mean work is $1.47k_BT$ and standard deviation about $1.72k_BT$. A sudden pull of the same distance has higher mean $2k_BT$. A pull lasting five relaxation times reduces the mean to $4k_BT[5 - (1 - e^{-5})]/25 = 0.641k_BT$. These changes predict which protocol will yield a usable work-overlap region for a finite number of trials.

A measurement can test this crossover without fitting free energy at all. Calibrate stiffness from an equilibrium position distribution, obtain the relaxation time from a correlation or driven response and predict the entire mean-work curve versus duration. Then assess Jarzynski estimates independently. Failure of the predicted variance-to-mean relation can reveal nonharmonic confinement, non-Gaussian bath forces, actuator timing offsets or insufficient equilibration. It need not invalidate the fluctuation identity itself.

A two-outcome work distribution shows how rare events restore equality

A particularly transparent work distribution comes from a sudden change in a two-level system. Initially its ground energy is zero and upper energy $2k_BT$. In a sudden protocol move the upper level to zero while leaving the initial ground energy unchanged. The initial upper probability is

$$p = e^{-2}/(1+e^{-2}) = 0.1192.$$

Work can take only two values: zero for an initially ground system and minus two thermal energies for a particle initially upper. Both new energy levels are degenerate, so the exact equilibrium partition ratio is

$$2/(1+e^{-2}) = 1.762.$$

The exponential average equals $(1-p)+pe^{2} = 1.762$, giving $\Delta F = -0.566 k_B T$. Mean work is only $-0.238 k_B T$, greater than this equilibrium free-energy change as required. With five independent realisations the probability of observing no upper preparation is

$$(1-p)^{5} = 0.530.$$

More than half such tiny samples would report zero work every time and infer a free-energy change of zero. A bootstrap resampling those five identical observations would report no empirical uncertainty at all, spectacularly missing the known sampling limitation.

Increase the initial gap to $6 k_B T$ before the same quench to degeneracy. The initial upper probability falls to 0.00247, but its exponential weight rises to e^{6}. Its mean contribution to the exponential estimator remains near one. Even one hundred runs then miss this branch with probability about 0.78. Equality relies on the true preparation probability of a rare event combined with its enormous statistical weight. A finite dataset cannot certify absence of the branch by having no observations of it.

The reverse experiment begins with two equally populated degenerate states and raises one level. It samples the formerly rare branch far more easily. Reflect its raw works onto the forward work axis and compare point masses, not unrelated smooth histogram densities. A forward and reverse data combination then constrains the free-energy difference much more effectively. A fitted reverse sample must still begin at equilibrium in the degenerate endpoint; carrying the low excited-state fraction over from the forward final state would destroy its intended initial ensemble.

Instrumental work noise is amplified by the exponential

Suppose each genuine trajectory work is recorded with an independent Gaussian error of mean zero and variance s_W^2. The arithmetic mean work stays unbiased. The mean exponential weight is multiplied by $\exp(\beta^2 s_W^2/2)$, so a naive Jarzynski estimate of free energy shifts downward by $\beta s_W^2/2$ in energy units in the infinite-record limit. This systematic shift is distinct from the finite-record logarithm bias, which tends upward on average for ideal correctly measured weights.

For rms measurement error $0.5 k_B T$, this Gaussian-noise effect is $-0.125 k_B T$. A small positive finite-sampling bias can accidentally cancel it, producing an apparently correct number with both the work measurement and ensemble estimate wrong. Vary the number of independent trajectories and the instrument noise separately. Calibrate work reconstruction on repeated known control or force signals to estimate the error distribution rather than assuming it is Gaussian from its rms alone.

If measurement error correlates with the trajectory—for example a saturated force transducer clips the largest restoring forces—the factorisation above fails. Errors in low-work tails can matter most even if a force trace looks accurate over the bulk of ordinary values. Compare independent sensors or forward and reverse force

calibrations across the full visited range. A threshold discarding negative work as unphysical would systematically delete exactly the values the equality needs.

A related problem occurs when temperature itself is inferred from the same finite fluctuation record. Temperature uncertainty affects both the exponent and conversion of the final logarithm to energy. A Gaussian work model gives an analytic sensitivity, but outside that model propagate uncertainty by re-evaluating the full weighted estimator over independently plausible temperatures and work calibrations. Treat correlations between calibration and work estimates appropriately; merely quoting the precision of the bath set point does not establish the temperature of a fast, weakly equilibrated preparation.

Verify a stochastic work protocol on a two-state network

A small simulated system with discrete energies and heat-bath jumps supplies an independent test of a coupled dynamics and work implementation. At each step change a control value instantaneously while preserving current state; record the corresponding energy difference as parameter work. Then allow a period of fixed-control stochastic relaxation whose transition probabilities preserve the new canonical distribution. Energy change during that relaxation is heat, not additional protocol work. Define the reverse process by precisely reversing the sequence of controlled changes and using an independent equilibrium start at its endpoint.

An exact proof for such a process does not require a continuum limit. After one energy change, multiplying the previous equilibrium weights by the exponential work converts them into unnormalised canonical weights of the new levels. A stochastic transition preserving these new weights leaves that weighted distribution unchanged. Apply the same argument successively and sum the final weighted probabilities to retrieve the ratio of endpoint partition functions. Each trajectory can jump many times, but the weighted identity follows from a short sequence of matrix operations.

This proof is a practical verification tool. Enumerate all paths of a two-state system over a few control steps and compare the exact weighted sum with a Monte Carlo estimate. Test a deliberately coarse but precisely specified sequence before refining the time mesh. If direct enumeration fails, inspect signs, equilibrium preparation, order of control and relaxation, and whether the transition kernel truly preserves the prescribed endpoint weights. If enumeration succeeds while sampled averages disagree, inspect rare weights and independent batch variability before modifying the underlying dynamics.

Why postselection changes a work experiment

The standard work equality assumes a specified forward protocol and its initial equilibrium ensemble. Suppose an operator begins with an equilibrium distribution in the harmonic trap but keeps only trajectories whose initial displacement already points toward the intended final centre. Those retained starts are not a canonical ensemble for the original unrestricted trap. The conditional work distribution is systematically shifted toward smaller work, so inserting it into the ordinary exponential average generally produces a biased apparent free-energy change. This selection problem arises even if no macroscopic control action visibly changes between the retained runs.

For a sudden positive shift d, work decreases linearly with initial x. Conditioning on $x > 0$ deletes precisely the high-work half of a symmetric Gaussian preparation. More generally, groups selected with probabilities depending on the measured initial state should be recombined using the true group frequencies if the aim is

to estimate the unconditional work moment. An equal-weight average of two conditionally estimated free energies is not equivalent to averaging their exponential weights, because the logarithm is nonlinear.

If the feedback actually chooses the direction of a subsequent shift after measuring the initial position, the protocol itself is now correlated with the state. For example, shifting the trap toward the initial displacement on each run reduces average mechanical work relative to a fixed-direction protocol. The additional information and the measurement/control apparatus are part of the thermodynamic accounting. A direct substitution into the no-feedback equalities derived above would address the wrong experiment. A full information-thermodynamic treatment needs specified measurement probabilities and a well-defined corresponding reverse procedure.

A simple experimental guard is to record every attempted run, including rejected starts, along with selection and trigger conditions. Verify the initial position and energy distributions both before and after data-quality cuts. A cut based on electronic saturation can remove a physically important work tail even if it has no deliberate feedback purpose. Conversely, discarding runs for an independent external electronics failure can be statistically benign when independence is demonstrated. The fact that a cut is operationally convenient is not evidence for such independence.

The same reasoning applies to chained protocols. A sequence of small equilibrium windows can improve overlap, but each initial distribution must be re-equilibrated or appropriately reweighted. Starting each window immediately after the previous rapid drive changes its preparation. A complete path-level method may still treat the sequence consistently; multiplying independent endpoint estimates derived under a false equilibrium assumption generally does not. An explicit record of state preparation, control timing and retained trajectories is as central to a fluctuation-relation measurement as an accurate mechanical force calibration.

A useful stationary control keeps the trap centre fixed but processes position and force data through the same analysis pipeline as a moving-trap run. Correct inclusive parameter work is then zero on every trajectory because the control never changes, even though the particle exchanges random heat with the bath and its potential energy fluctuates. A processed nonzero work distribution in this test reveals that the analysis integrated a different quantity, confused particle displacement with actuator displacement or introduced artificial control noise. Conversely a real translating-trap experiment can finish with the same mean potential energy it had initially after equilibration, while having delivered positive net work and dissipated heat. Tracking both control displacement and particle displacement makes these statements testable trajectory by trajectory where the specified stochastic energy convention permits.

Determine which rare work values dominate a free-energy estimate

Change the preparation gap and the number of repetitions. Compare the average work, the exact free-energy change and the chance of missing the rare branch entirely.

Open this experiment online

Try the calculation

Initial gap (kBT) 4

Independent preparations 20

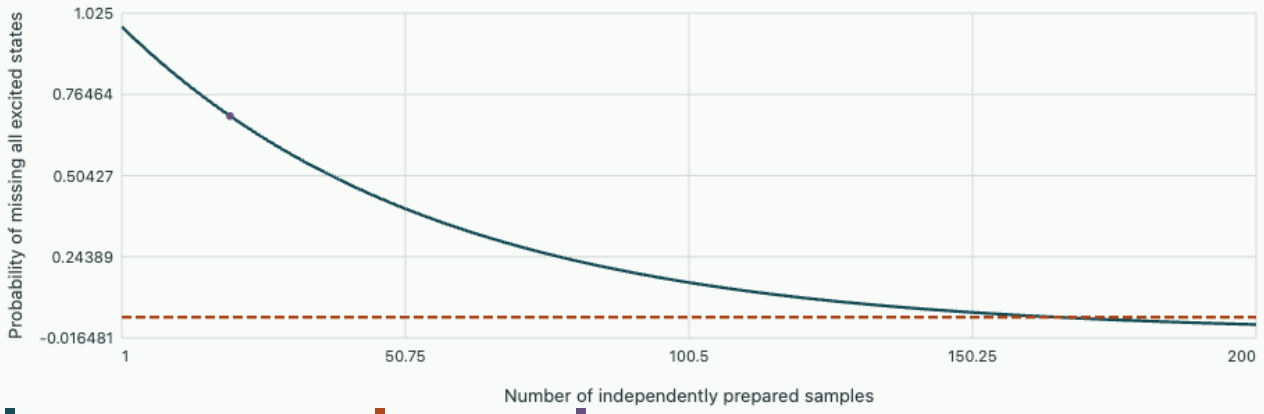

No initially excited sample | Five percent | Selected experiment size

P(no rare work sample) = 0.69559. Mean work / kBT = -0.071945, ΔF/kBT = -0.675.

Tutorial · A stiffness quench changes the equilibrium free energy

An equilibrated classical harmonic trap has stiffness k. It is suddenly changed to $2k$ at the same centre. Derive mean work, exact free-energy difference and verify the exponential equality without assuming Gaussian work.

▼ **Integrate the Gaussian position rather than a guessed Gaussian work law**

Work is $W = kx^2/2$, so mean work is $k_BT/2$. The configurational partition function scales as $k^{-1/2}$, giving $\Delta F = (k_BT/2)\log 2 = 0.3466 k_BT$. The exponential-weight average is the integral of $e^{-\beta k x^2/2}$ against the original position density proportional to $e^{-\beta k x^2/2}$. The combined Gaussian integral yields

$$1/\sqrt{2} = e^{-\beta \Delta F}.$$

The work is nonnegative and chi-square distributed, not Gaussian. Its mean excess above ΔF is about $0.1534 k_BT$.

Tutorial · Estimate a crossing with uncertainty

Measured forward and sign-reflected reverse densities have equal-variance Gaussian shapes with means 6.0 and 2.0 thermal energies. Their standard deviations are both 2.0 thermal energies. Find the crossing, verify the expected logarithmic slope and explain how an unmodelled force offset would affect a displacement protocol.

▼ **Complete the square and track the offset through both protocols**

The ratio of the two Gaussian densities has logarithmic slope $(6-2)/4=1$ per thermal energy and vanishes at their midpoint $4.0k_BT$. This is consistent with a free-energy change of four thermal energies. If fitted means have independent standard errors 0.2 and 0.3 thermal energies, their midpoint has standard error $\sqrt{0.2^2+0.3^2}/2=0.18k_BT$ under the equal-variance model. Unknown unequal variances or poor histogram overlap require a more general fit.

An additive force calibration offset integrated over a known displacement adds a constant work shift proportional to that displacement. The reverse protocol has opposite displacement and hence opposite offset in its raw work, which becomes the same shift after reflecting its work axis. The crossing moves with both histograms and their difference need not reveal the calibration error. Independent force-zero checks matter.

Tutorial · Why an unchanged average dissipated work is insufficient

Two protocols with the same endpoints have the same mean work above equilibrium free-energy change. One has a narrow distribution near that mean and a very rare low-work tail; the other has a broader nearly Gaussian shape. Can equal mean dissipation guarantee comparable convergence of a unidirectional free-energy estimate?

▼ **Compare the quantities that govern the estimator**

The mean work determines only the first moment. The exponential estimator depends on the entire distribution and its variance depends on the second exponential moment $\langle e^{-2\beta W}\rangle$, which can differ greatly at fixed $\langle W\rangle$ and fixed Jarzynski average. A rare low-work tail can dominate both the true mean weight and its sampling error. Identical mean dissipation does not determine its scale or frequency.

Compare forward and reverse overlap over the actual important work interval, analyse independent batches and vary protocol speed. In a proposed experiment verify initial equilibrium at each endpoint, acquisition of the full work history and calibration of temperature and control force. Convergence of a mean mechanical work with repeated trials answers a different question from convergence of an exponential estimate of equilibrium free energy.

Jarzynski · original nonequilibrium equality for free-energy differences

Crooks · entropy production and the forward/reverse work relation

Statistical mechanics from microscopic counting to material response · 13. Put a thermometer on a bounded spin spectrum

A collection of identical two-level degrees of freedom has both a lowest and a highest energy. Count the accessible states at each energy and determine from adjacent entropy differences which way energy flows when two systems are brought into weak thermal contact.

Differentiate an entropy measured in counts

Let N distinguishable, weakly coupled sites each carry energy zero or ε. A state with n excitations has $E_n = n\varepsilon$. The multiplicity of the narrow, fixed-energy shell equals the number of ways to place the excitations, $W_n = \binom{N}{n}$. Use the Boltzmann shell entropy $S_n = k_B \ln W_n$. The temperature inferred from the population of two adjacent shells samples the derivative at an energy midway between them. Dividing their multiplicities before taking a logarithm gives an exact finite-system expression,

$$\beta_{n+1/2}\varepsilon = \frac{S_{n+1} - S_n}{k_B} = \ln \frac{N-n}{n+1},$$

$$\beta = \frac{1}{k_B T}.$$

Adding an excitation near the bottom initially opens more arrangements. Beyond the middle, it closes them. The decreasing entropy on this upper branch assigns a negative inverse temperature to the population. In the limit of large N at fixed excitation fraction p, the slope becomes $\ln[(1-p)/p]$. The same relationship follows from $p/(1-p) = \exp(-\beta\varepsilon)$ for a canonical two-level system.

The upper energy bound allows a distribution weighted towards high energy to normalise. In experiment the degrees of freedom considered must also equilibrate amongst themselves on a timescale shorter than their relaxation into ordinary, effectively unbounded environmental modes. This calculation describes the spin subsystem during that interval.

Predict the direction of an exchange

Take $N = 20$ and examine the transition from 15 to 16 excitations. The multiplicity ratio is $5/16$, the dimensionless inverse temperature is –1.1632, and $k_B T/\varepsilon$ is approximately –0.860. The reflected transition from 4 to 5 has ratio $16/5$ and the opposite slope. The same number of possible arrangements at energies reflected about the midpoint gives opposite entropy gradients.

Suppose system A initially has 16 excitations and the same-sized system B has 4. Moving one quantum from A to B multiplies the product of their multiplicities by $(16/5)^2 = 10.24$. A weak contact that samples the available microstates favours transfer out of the inverted specimen. More generally, when A gives up a small positive energy q, the entropy change is $\Delta S/k_B \simeq (\beta_B - \beta_A)q$. An inverted subsystem therefore gives up heat to an ordinary positive-temperature subsystem when the coupling permits energy exchange.

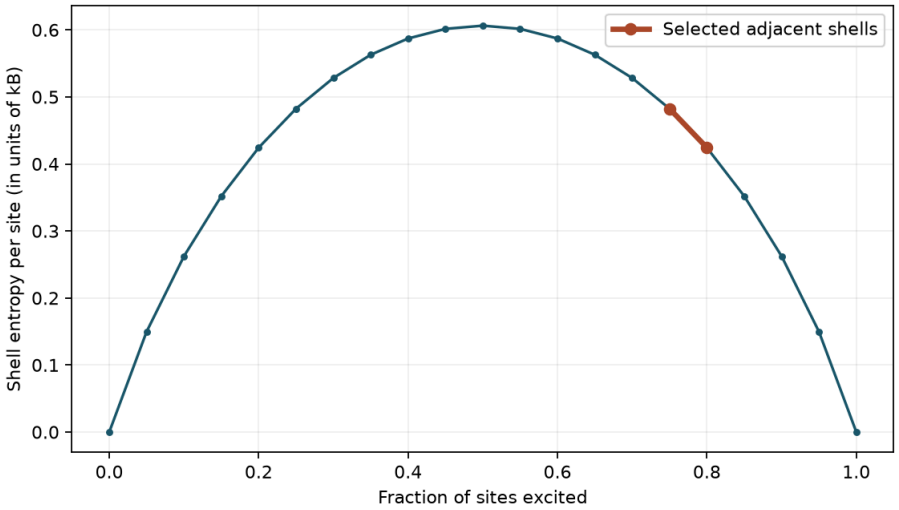

The segment joining the two highlighted shells measures the entropy difference used in the finite-system inverse temperature.

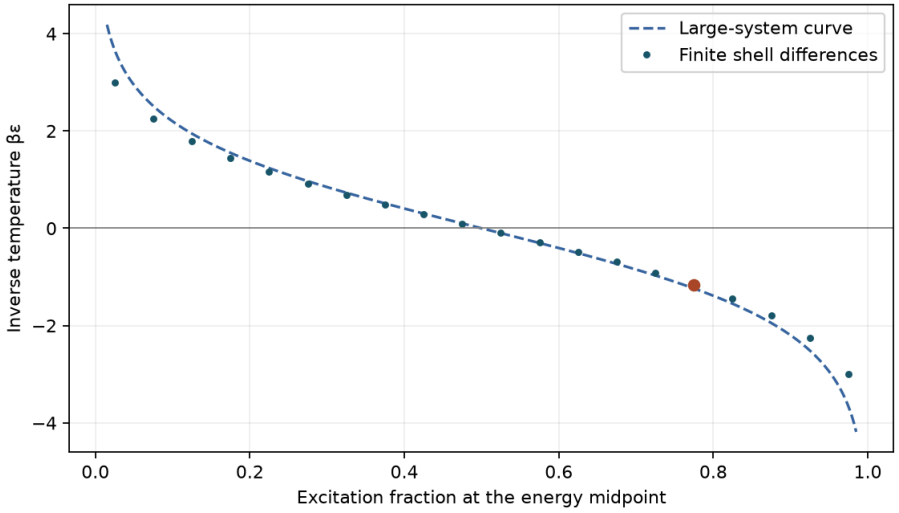

The discrete slope belongs between neighbouring energies. Compare finite sizes at the same fractional position to see where the large-system approximation converges first.

Exercise · the temperature halfway up an odd spectrum

For eleven sites find the adjacent shells with exactly zero inverse temperature. What site occupation does this represent in the canonical description?

▼ **Work through the solution**

Setting $11 - n = n + 1$ gives $n = 5$. Five and six excitations have the same multiplicity, giving zero entropy slope at $E = 5.5\varepsilon$. In the canonical ensemble $\beta = 0$ gives equal probabilities for both site energies. The average excitation number is 5.5 although individual measurements return integers.

Exercise · infer a population inversion

A spin measurement yields three excited sites for every site in the lower state, in a sample large enough to neglect finite-size corrections. Calculate the inverse temperature and the mean energy of 100 equivalent sites.

▼ **Work through the solution**

Here $p = 3/4$. The Boltzmann ratio gives $\beta\varepsilon = -\ln 3 \simeq -1.099$, and the mean energy is

$$100p\varepsilon = 75\varepsilon.$$

At the positive inverse temperature $+\ln 3/\varepsilon$ the mean falls to 25ε. The two occupations have the same binary mixing entropy and opposite response to an infinitesimal energy exchange.

Try it · Negative temperature

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 bounded two-level sites	**20**
Selected excitation fraction	**0.75**

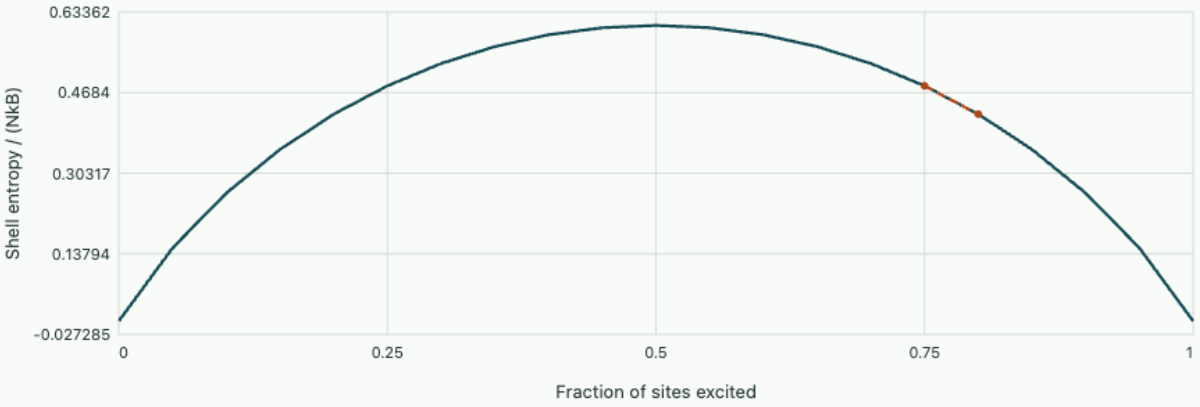

Exact log binomial multiplicity Adjacent shells used for temperature

Finite difference βε = ln[(N−n)/(n+1)] = -1.1632. Above half filling entropy decreases with energy, giving negative temperature for this bounded spectrum.

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

Statistical mechanics from microscopic counting to material response · 14. Follow a refrigerant through four devices and a gas through a liquefier

A thermometer on a valve tells us little about the cooling capacity of a refrigerator. We need to follow one kilogram through its entire cycle and distinguish flash evaporation inside the valve from heat absorbed from the cold load.

Assign the four states and derive the cooling performance

Use a steady one-inlet, one-outlet control volume for each component and neglect kinetic and potential energy changes. The flow work of the inlet and outlet accompanies internal energy as $h = u + pv$. With q positive into the fluid and w_{in} positive as shaft work supplied, the local first law is $h_{\mathrm{out}} - h_{\mathrm{in}} = q + w_{\mathrm{in}}$.

Label 1 the evaporator exit at low pressure, 2 the compressor discharge at high pressure, 3 the condenser exit liquid at high pressure and 4 the valve exit mixture at low pressure. Assume an adiabatic compressor, an adiabatic valve with no shaft and negligible pressure drops in the heat exchangers. Then the following equations come directly from four separate first laws.

$$w_c = h_2 - h_1,$$

$$q_{\mathrm{out}} = h_2 - h_3,$$

$$h_4 = h_3,$$

$$q_c = h_1 - h_4.$$

$$\mathrm{COP}_c = \frac{h_1 - h_4}{h_2 - h_1},$$

$$q_{\mathrm{out}} = q_c + w_c.$$

The last equality checks the whole loop. For an adiabatic real compressor use its isentropic efficiency $\eta_c = (h_{2s} - h_1)/(h_2 - h_1)$ with $s_{2s} = s_1$ at the discharge pressure. Solve for the larger actual enthalpy rise. Shaft losses and heat leakage would require further stated measurements.

Work one cycle with measured and tabulated quantities

For R-134a take saturated vapour leaving at −10 °C and saturated liquid leaving the condenser at 30 °C. Rounded ASHRAE saturated-fluid data give $p_1 = 0.20060$ MPa, $h_1 = 392.66$ kJ/kg and $h_{f,1} = 186.70$ kJ/kg. At 30 °C the pressure is 0.77020 MPa and $h_3 = 241.72$ kJ/kg. Suppose the measured discharge state corresponds to $h_2 = 430.00$ kJ/kg on the same reference. ASHRAE Handbook, R-134a saturation data.

The valve leaves 241.72 kJ/kg. At the low pressure that energy lies between the saturated liquid and vapour values. Writing $h_4 = (1 - x_4)h_f + x_4 h_g$ gives $x_4 = (241.72 - 186.70)/(392.66 - 186.70) = 0.267$. Thus 26.7% by mass is already vapour before any externally supplied cooling load arrives.

Each kilogram removes 150.94 kJ and costs 37.34 kJ in compressor work. The COP is 4.04. At 0.010 kg/s this is 1.509 kW of cooling, 0.373 kW of compressor power, and 1.883 kW of condenser rejection. The reversible upper bound between isothermal reservoirs at 263.15 and 303.15 K is 6.58. A real load and heat sink need approach temperatures, which can lower the applicable upper bound.

Recall · An enthalpy zero is a convention

Adding the same constant to all the specific enthalpies changes neither performance nor vapour fraction. Mixing tables with different reference states can violate the four-state energy balance. Two temperatures at saturation alone do not determine a two-phase enthalpy because the mass fraction also matters.

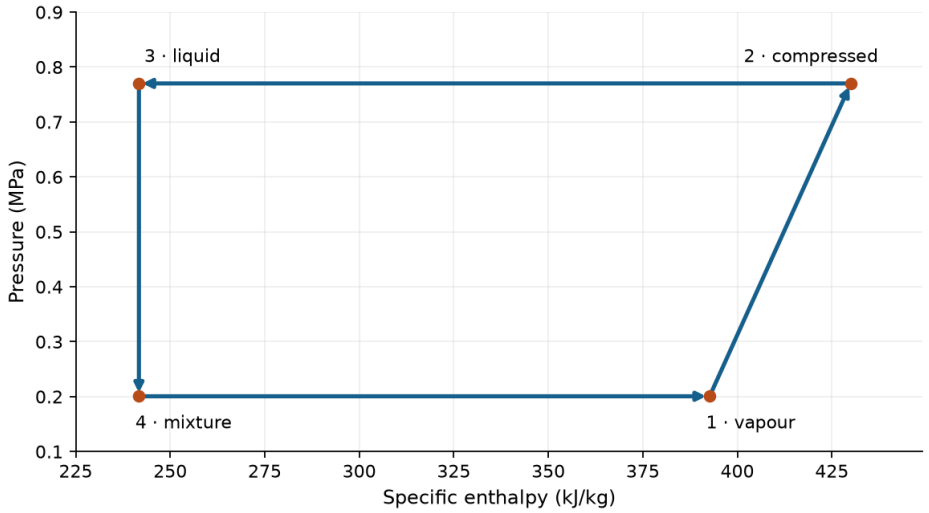

Four R-134a states on pressure–enthalpy axes. Arrows give component order. Their connecting lines do not resolve the internal path through each device. Evaporation removes heat along the bottom edge, while throttling leaves enthalpy fixed.

Try it · Refrigeration

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

Discharge enthalpy h2 (kJ/kg) **430**

Liquid subcooling (K) **0**

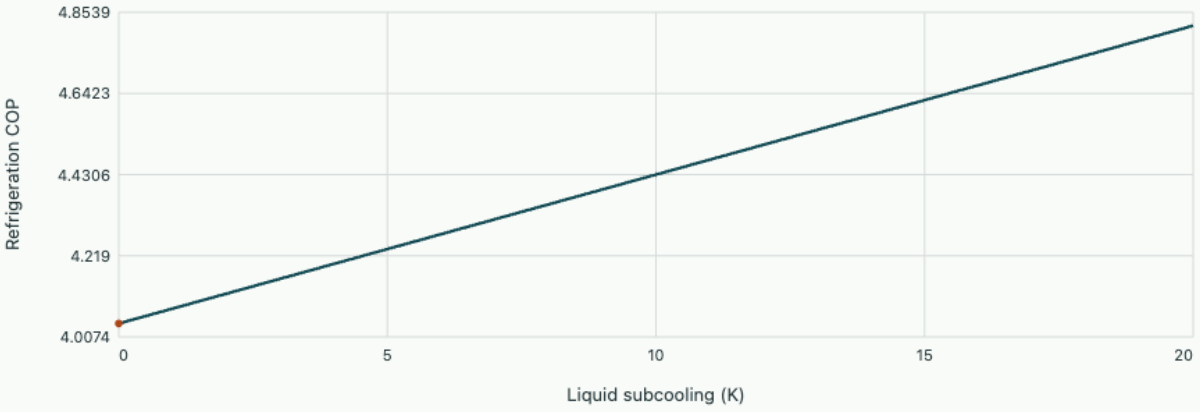

Four-state enthalpy balance Selected subcooling

At selected states: cooling 150.94 kJ/kg, work 37.34 kJ/kg, valve vapour fraction 0.26714. Saturated endpoint data are held fixed; liquid cp=1.45 kJ/(kg K) is an illustrative local approximation.

Close the balance around a regenerative gas liquefier

Throttle a gas at fixed enthalpy. In a single phase $dh = c_p\,dT + [v - T(\partial v/\partial T)_p]\,dp$. The initial cooling or heating slope is therefore $\mu_{\mathrm{JT}} = (T(\partial v/\partial T)_p - v)/c_p$. For an ideal gas this vanishes. A real gas must enter on the appropriate side of its inversion curve for an initial pressure drop to cool it.

Send high-pressure gas down a counterflow heat exchanger before the valve. Separate the cold liquid from the flashed vapour and return that vapour through the exchanger, cooling the next incoming stream. Around the combined exchanger, valve and separator, let the inlet specific enthalpy be h_H, the warm low-pressure return h_R and the separated liquid h_L. Per unit mass entering, an adiabatic outer balance is

$$h_H = y h_L + (1 - y) h_R,$$

$$y = \frac{h_R - h_H}{h_R - h_L}.$$

For illustrative values $h_H = 260$, $h_R = 275$, $h_L = 40$ kJ/kg the liquid yield is 0.0638. A heat leak of 3 kJ per kg fed changes the numerator to $h_R - h_H - 3$ and the yield to 0.0511. An ideal-gas throttle alone cannot realise these end states. Check attainable pressures and a positive temperature difference everywhere in the exchanger before interpreting an enthalpy balance as a feasible device.

Tutorial on diagnosing a cycle

1 · Recover a flow rate from electrical input

▼ Work through the solution

If the compressor actually receives 0.560 kW of shaft work at the same states, $\dot{m} = 0.560/37.34 = 0.0150$ kg/s. Cooling is $\dot{m}(150.94) = 2.26$ kW. An electrical meter upstream of a lossy motor also needs the motor efficiency.

---

2 · What does five kelvin of subcooling change?

▼ Work through the solution

At an illustrative liquid heat capacity of 1.45 kJ/(kg K), h_3 falls about 7.25 kJ/kg at fixed pressure. With unchanged inlet and discharge states, q_c increases to 158.19 kJ/kg and the flash fraction falls by about $7.25/205.96 = 0.0352$. A measured change in compressor inlet conditions would change this simple comparison.

---

3 · Detect an impossible liquefier specification

▼ Work through the solution

A claimed yield of 0.10 with the stated warm and liquid enthalpies would require $h_H = 0.10(40) + 0.90(275) = 251.5$ kJ/kg before any heat leak. The proposed inlet at 260 kJ/kg exceeds it. More regeneration, precooling or different pressure levels must be specified for that yield.

For further study see Schroeder, An Introduction to Thermal Physics, chapter 4 on real refrigerators, throttling and liquefaction.

Statistical mechanics from microscopic counting to material response · 15. Eutectic mixtures, freezing-point depression and boiling-point elevation

Adding a nonvolatile solute can lower the temperature at which a solvent begins to freeze. In a binary eutectic both species can solidify. The two observations follow from the same equality of chemical potentials, with different phases available for equilibrium.

Derive the boundary between liquid and a pure solid

At fixed pressure take a liquid ideal solution and immiscible pure solids A and B. Let x_i be a liquid mole fraction. The liquid potential is $\mu_i^{l,*} + RT \ln x_i$ where the star denotes the pure liquid at the same temperature. Equilibrium with pure solid i requires $\mu_i^{l,*} + RT \ln x_i = \mu_i^{s,*}$. At its pure melting point $T_{m,i}$ the standard potentials coincide.

Use the Gibbs–Helmholtz identity $d(\mu/T)/dT = -h/T^2$ at fixed pressure on the pure-phase difference. Assuming the molar enthalpy of fusion stays constant over the temperature interval gives

$$\ln x_i = -\frac{\Delta h_{\mathrm{fus},i}}{R}\left(\frac{1}{T} - \frac{1}{T_{m,i}}\right).$$

Cooling a liquid enters a two-phase region when one species first satisfies this equality. As that pure solid forms, material balance forces the remaining liquid composition along its liquidus. Both pure solids can coexist with the liquid only where the two predicted liquid mole fractions sum to one. The isobaric phase rule gives $F = C - P + 1 = 0$ at this three-phase point.

Read amounts before and after reaching the eutectic

Use model components with $T_{m,A} = 900$ K, $T_{m,B} = 700$ K and fusion enthalpies 14.5 and 10.0 kJ/mol. The eutectic lies at 569.65 K and liquid $x_B = 0.67493$. Start at overall $z_B = 0.40$. Solid A forms first. Just above the eutectic all B resides in the liquid, so its mole-basis amount fraction is $z_B / x_{B,l} = 0.593$ and the remaining 0.407 is already solid A.

At constant pressure the remaining eutectic liquid then converts isothermally into two solids as latent heat leaves. Below the eutectic the final amounts of pure B and A are 0.40 and 0.60 of the total moles. A sufficiently fast quench can supercool the sample or prevent equilibrium phase compositions. A micrograph alone does not locate the equilibrium temperature.

OneSimulation experiment · Where the two liquidus branches meet

Try changing the phase diagram

Change B's melting point or fusion enthalpy, then move the overall composition across the eutectic. Read which solid appears first and use the reported liquid fraction to check the lever rule. The calculation bisects the composition-closure equation and redraws both liquidus branches for each input.

The model assumes an ideal liquid, two immiscible pure solids, constant fusion enthalpies and equilibrium at fixed pressure. Set `melting_b` and `fusion_b` (in J/mol) in the OneSimulation notebook below to reproduce a chosen diagram. The composition slider selects a cooling path; it does not change the equilibrium boundaries.

Open this experiment online

Try the calculation

Melting point of B (K)	700
Fusion enthalpy of B (kJ/mol)	10
Overall mole fraction of B	0.4

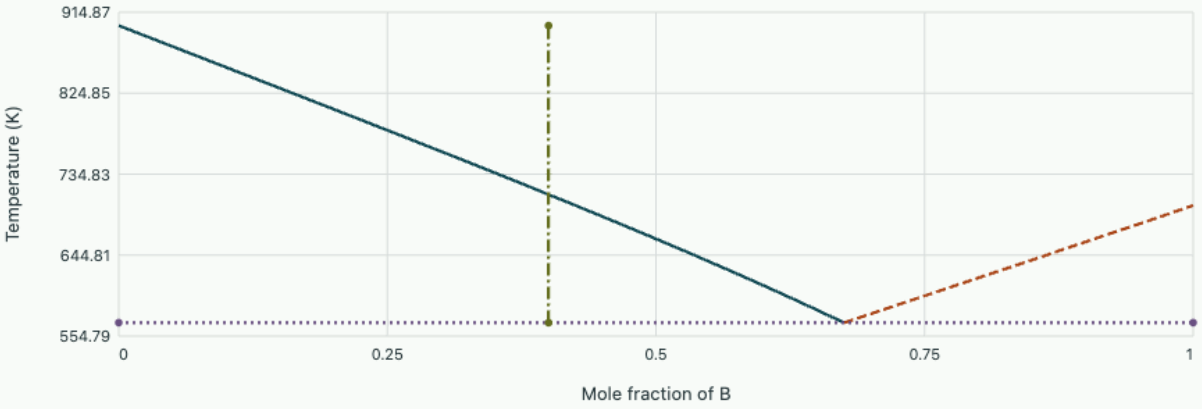

Eutectic: 569.653 K, xB = 0.674927. First solid: A. Liquid fraction immediately above the eutectic = 0.592657 (mole basis). A remains fixed at 900 K and 14.5 kJ/mol.

Reproduce with OneSimulation

Change one fusion enthalpy and predict which liquidus moves. OneSimulation solves the monotone composition closure equation. Check the returned composition against each single-component equation before interpreting the figure.

Run the downloaded notebook with OneSimulation, NumPy and Matplotlib installed. The output below was recorded with the displayed calculation.

```
"""Ideal liquid in equilibrium with two immiscible pure solids."""
import numpy as np
import matplotlib.pyplot as plt
from onesim.numerics.nonlinear.kernels.local_solvers.inversion.bracketed_inverse import bisect_increasing

gas_constant = 8.314462618
melting_a, melting_b = 900.0, 700.0
fusion_a, fusion_b = 14500.0, 10000.0  # J per mol

def liquid_fraction(t, melting, heat):
    return np.exp(-heat / gas_constant * (1/t - 1/melting))

def fraction_sum(t):
    return liquid_fraction(t, melting_a, fusion_a) + liquid_fraction(t, melting_b, fusion_b)
```

```
temperature = float(bisect_increasing(fraction_sum, 1.0, 250., melting_b))
x_b = liquid_fraction(temperature, melting_b, fusion_b)
print(f"Eutectic temperature = {temperature:.2f} K")
print(f"Eutectic liquid mole fraction of B = {x_b:.5f}")
print(f"Composition closure error = {abs(fraction_sum(temperature)-1):.2e}")
ta = np.linspace(temperature, melting_a)
tb = np.linspace(temperature, melting_b)
plt.plot(1-liquid_fraction(ta, melting_a, fusion_a), ta, label="Liquid + solid A boundary")
plt.plot(liquid_fraction(tb, melting_b, fusion_b), tb, label="Liquid + solid B boundary")
plt.axhline(temperature, linestyle=":", color="grey", label="Three-phase temperature")
plt.scatter([x_b], [temperature], color="black")
plt.xlabel("Mole fraction of B")
plt.ylabel("Temperature (K)")
plt.legend()
plt.tight_layout()
plt.savefig("physics-eutectic.png", dpi=160)
```

Computed output

```
Eutectic temperature = 569.65 K
Eutectic liquid mole fraction of B = 0.67493
Composition closure error = 2.22e-16
```

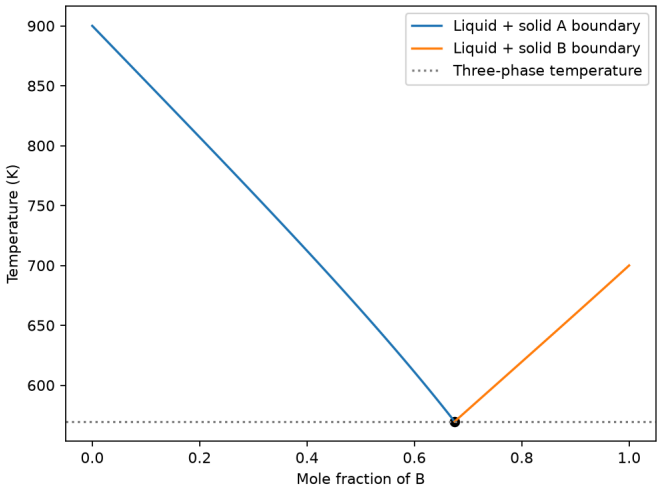

Where the two liquidus branches meet.

Download editable notebook · Download Python script

Expand the solvent equilibrium about its pure transition temperatures

Take a small mole fraction x_s of nonvolatile solute. For freezing into a pure solvent solid use $x_A = 1 - x_s$ and $\ln(1 - x_s) \simeq -x_s$. Expanding the reciprocal temperature around the pure freezing point T_f gives

$$T_f - T_{\text{freeze}} \simeq \frac{R T_f^2}{\Delta h_{\text{fus}}} x_s.$$

For boiling into pure solvent vapour at fixed external pressure the equality is $\mu_g^* = \mu_l^* + RT \ln x_A$. The entropy difference between vapour and liquid at T_b is positive. Expanding $\mu_g^* - \mu_l^* \simeq -(\Delta h_{\mathrm{vap}}/T_b)(T - T_b)$ gives the opposite temperature shift.

$$T_{\mathrm{boil}} - T_b \simeq \frac{R T_b^2}{\Delta h_{\mathrm{vap}}} x_s.$$

In the dilute limit $x_s \simeq m_s M_A$ with solute molality m_s in mol/kg and solvent molar mass M_A in kg/mol. For water, $M_A = 0.0180$, $T_f = 273.15$ K, $\Delta h_{\mathrm{fus}} = 6.01$ kJ/mol. A 0.10 mol/kg nonelectrolyte gives a freezing-point depression near 0.186 K. Using $T_b = 373.15$ K and $\Delta h_{\mathrm{vap}} = 40.7$ kJ/mol gives a boiling-point elevation near 0.051 K.

> **Recall · Count dissolved particles and choose the stable phases**
>
> Dissociation changes particle count. Activities replace simple mole fractions when solutes interact. Solid solubility requires a separate chemical potential for solute in the solid. A shift caused by nucleation delay depends on cooling rate and can be distinguished by comparing heating and cooling scans.

Tutorial on a cooling curve

1 · Follow a B-rich sample

▼ Work through the solution

For $z_B = 0.80$ solid B forms first. Immediately above the eutectic all A remains liquid. The liquid amount is $(1 - 0.80)/(1 - 0.67493) = 0.615$. A lever rule requires specifying whether compositions are mole or mass fractions.

---

2 · Estimate the apparent molar mass

▼ Work through the solution

Dissolving 1.00 g of an ideal nonelectrolyte in 0.100 kg of water depresses freezing by 0.186 K. With $K_f = 1.86$ K kg/mol the molality is 0.100 mol/kg, so 0.0100 mol weighs 1.00 g and the molar mass is 100 g/mol.

---

3 · Check a phase-diagram inference

▼ Work through the solution

A liquid containing only 40% B cannot first precipitate solid B in the model when its liquidus is on the A side. Solids seen after complete equilibrium solidification can contain both A and B even though the first nuclei

were solely A. Determine the first phase from the crossing of the cooling path with the liquidus.

For further study see Schroeder, An Introduction to Thermal Physics, sections 5.4–5.5 on phase changes in mixtures and solutions.

Statistical mechanics from microscopic counting to material response · 16. From reaction affinity to a reversible cell voltage

An electrochemical cell transfers reaction free energy into electrical work through a controlled charge transfer. A measured voltage can therefore tell us about a reaction quotient, provided the electrodes are in reversible equilibrium.

Identify which thermodynamic potential controls useful work

Consider an isothermal, isobaric cell releasing reactants into products according to stoichiometric numbers ν_i with products positive. For an extent change $d\xi$, $dn_i = \nu_i\, d\xi$. The chemical contribution to the change in Gibbs free energy is $\Delta_r G\, d\xi$ where $\Delta_r G = \sum_i \nu_i \mu_i$. The maximum non-expansion work delivered reversibly to an external circuit is $-dG$ for the reaction step. If it sends n moles of electrons through the external circuit per mole of reaction, the charge magnitude is $nF\, d\xi$.

$$nF E_{\mathrm{rev}} = -\Delta_r G,$$

$$\Delta_r G = \Delta_r G^{\circ} + RT \ln Q,$$

$$Q = \prod_i a_i^{\nu_i}.$$

Here $F = N_A e$ and the activities are dimensionless relative to standard states. Substituting directly yields the Nernst equation. Reversing the written chemical reaction reverses the signed cell voltage and reciprocates Q.

$$E_{\mathrm{rev}} = E^{\circ} - \frac{RT}{nF} \ln Q,$$

$$\ln K = \frac{nF E^{\circ}}{RT}.$$

Make voltage from the same material in two environments

Take two identical reversible monovalent metal electrodes in ideal solutions with different ion activities and negligible liquid-junction potential. Combining oxidation in the dilute compartment with reduction in the concentrated one transfers an ion effectively from high to low activity. The standard contributions cancel.

$$E = \frac{RT}{F} \ln \frac{a_{\mathrm{high}}}{a_{\mathrm{low}}}.$$

At 298.15 K a tenfold ratio gives 59.16 mV. Doubling both activities leaves the voltage unchanged for ideal activities. In a nonideal solution measure activity coefficients or use a calibrated reference. A voltmeter that draws significant current changes the surface concentrations and thus the very chemical potentials under study.

Distinguish reversible work from power delivered under load

For a model cell with $n = 2$ and $E^\circ = 0.80$ V at 298 K, a reaction quotient of 100 lowers the voltage by 59.2 mV. The reversible work at this composition is $2F(0.7408) \simeq 143$ kJ/mol. Under an external load suppose a 10 mA current flows with measured terminal voltage 0.68 V. The electrical output is then 6.8 mW, with about 0.61 mW less output than the reversible power at that instant. Ohmic resistance, interfacial kinetics and transport all contribute. A single operating point cannot distinguish them.

Recall · Why can equilibrium electrodes show a voltage?

Individual electrode reactions can be in local equilibrium at open circuit while the bulk reactants and products remain separated. An external path would allow the overall chemical reaction to move toward equilibrium, where the full-cell reversible voltage for that reaction vanishes.

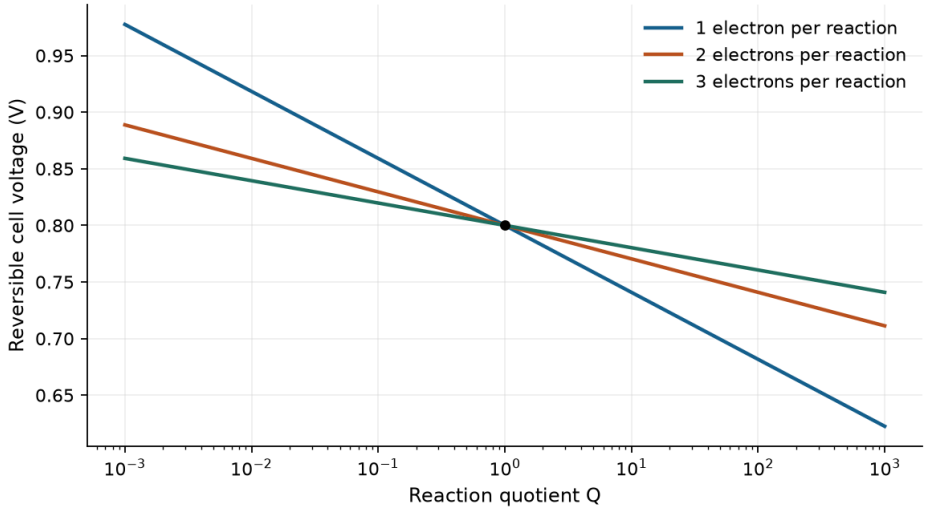

Increasing the product-to-reactant activity quotient lowers the reversible voltage. At the same standard voltage the slope per decade falls as the number of electrons per reaction increases.

Try it · Nernst

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)	300
Electrons per reaction	2
Standard voltage (V)	1.1

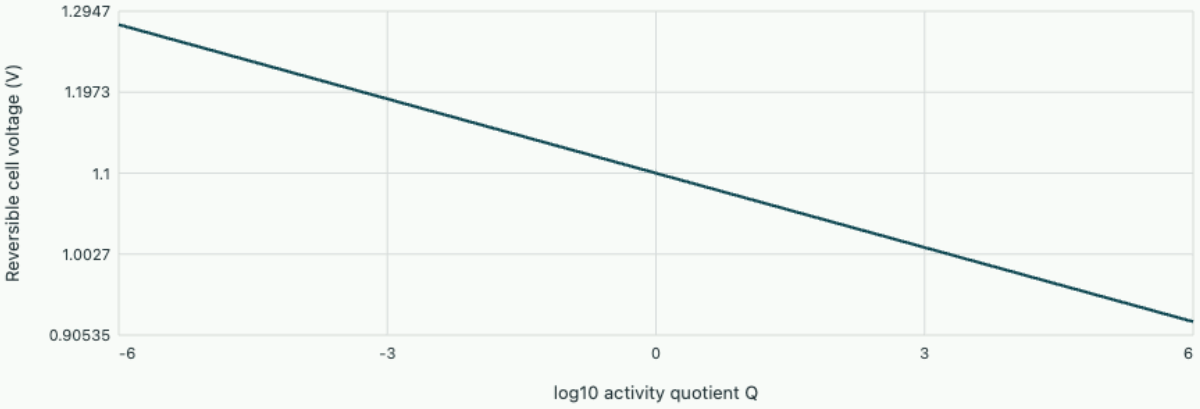

Nernst equation

Slope -0.029763 V per decade. Activities are dimensionless; irreversible overpotentials are excluded.

Tutorial on activities and polarity

1 · Find an activity ratio from a voltage

▼ Work through the solution

A monovalent concentration cell gives 35.0 mV at 298.15 K. Inverting gives $a_{\mathrm{high}}/a_{\mathrm{low}} = \exp(0.0350/0.02569) = 3.91$. Reversing the leads changes the sign while the inferred larger-to-smaller ratio remains above one.

2 · Which reaction transfers more charge per released joule?

▼ **Work through the solution**

If two reactions at the same temperature each deliver 80 kJ/mol reversibly but carry respectively one and two moles of electrons, their voltages are 0.829 V and 0.415 V. Adding a stoichiometric multiple of an entire reaction multiplies both n and $\Delta_r G$, leaving voltage invariant.

---

3 · Use a temperature sweep for calorimetry

▼ **Work through the solution**

Since $(\partial \Delta_r G/\partial T)_p = -\Delta_r S$, reversible voltage at fixed pressure and consistently specified composition gives $\Delta_r S = nF(\partial E/\partial T)$. A two-electron cell with slope -0.20 mV/K has $\Delta_r S \simeq -38.6$ J/(mol K). Compare at the same reaction extent and control changes in activity coefficients.

For further study see Schroeder, An Introduction to Thermal Physics, chapters 3 and 5 on chemical potential, free energy and electrochemical work.

Counting, correlations and thermal measurements

Calculate an observable from an explicitly countable microscopic model. Determine which finite-size, sampling or instrumental effect changes a measured equilibrium or relaxation law.

1. ## How many lattice sites remember the colour of a neighbour?

   A chain of three-state spins gains energy when adjacent colours agree. A weak field favours one colour over the other two. Use a transfer matrix to calculate both free energy and the distance over which correlations survive. Recover the same numbers from OneSimulation eigenvalues without relying on symmetry reduction.

2. ## Measure the heat capacity of a reservoir from the width of an energy histogram

   Two small Einstein solids share a fixed number of excitation quanta. A calorimeter records how many are in the first solid. Count the allowed microstates, compare the exact distribution with a reservoir approximation and see how energy conservation correlates the readings.

3. ## Recover a free-energy difference from a distribution of work

   A two-state system starts in equilibrium, then a control jumps its energy gap. Individual work readings differ although every run uses the same control protocol. Use the elementary model to derive an exponential work identity, compare it with the average work and estimate how many runs it takes to see the events that dominate the estimate.

4. ## Why a molecular beam and its source have different measured speeds

   A tiny hole lets a dilute gas escape into a vacuum. A rotating chopper admits a short pulse and a downstream detector records arrival times. Derive the source speed density, the effusive speed density and the actual time-of-flight histogram. Each samples the same Maxwellian physics with a different weighting.

5. ## Detect irreversible steady motion from transition counts

   A three-state system appears equally likely to occupy each state. Is it at thermal equilibrium? Build two different stationary processes with identical populations, then determine what a time-resolved instrument must count to distinguish them.

6. ## Why macroscopic measurements become reproducible

   The same collection of microscopic switches can teach us two different facts about a macroscopic sample. Its mean becomes predictable, and the fluctuations about that mean acquire a particular shape. Derive both statements, then find an experiment in which only the first prediction is reliable at the available sample size.

7. ## Weigh a colloid by counting where it spends its time

Brownian motion spreads a dilute suspension through a cell whilst buoyancy biases its vertical position. Derive the equilibrium concentration, predict counts in finite camera bins and estimate a buoyant mass without measuring how quickly any individual sphere falls.

Counting, correlations and thermal measurements · 1. How many lattice sites remember the colour of a neighbour?

A chain of three-state spins gains energy when adjacent colours agree. A weak field favours one colour over the other two. Use a transfer matrix to calculate both free energy and the distance over which correlations survive. Recover the same numbers from OneSimulation eigenvalues without relying on symmetry reduction.

Find a local weight whose product counts each energy once

Each site has colour $a = 0, 1, 2$. A matching pair contributes energy $-J$, with $J > 0$, and a site of colour zero earns an additional $-H$. In a periodic ring each site belongs to two adjacent bonds. Give each of those bonds half the local field energy so the product of their weights counts each field term exactly once. Define $K = J/(k_B T)$ and $h = H/J$.

$$M_{ab} = \exp\left[K\left(\delta_{ab} + \frac{h}{2}(\delta_{a0} + \delta_{b0})\right)\right],$$

$$Z_N = \operatorname{tr}(M^N).$$

To see the trace, write the partition function as a sum over $a_1, \ldots, a_N$ of the product $M_{a_1 a_2} \cdots M_{a_N a_1}$. Summing an intermediate index is matrix multiplication; closing the initial and final index is the trace. This is exact even for a short periodic chain defined by that bond count. The transfer matrix contains no time step: multiplying it advances spatial position.

Compare the leading eigenvalue with the next one

Without bias the matrix has diagonal e^K and all other entries one. The uniform vector has eigenvalue $e^K + 2$. Every vector with components summing to zero has eigenvalue $e^K - 1$, twice degenerate. Thus $Z_N = (e^K + 2)^N + 2(e^K - 1)^N$. The thermodynamic free energy per site is $-k_B T \ln(e^K + 2)$. To measure a connected spatial correlation, insert a diagonal colour observable at sites separated by r transfer operations. The uniform contribution cancels from the connected part; each nonuniform mode contributes a factor $(\lambda_j/\lambda_1)^r$ when its overlap is nonzero. The longest correlation length in sites is

$$\xi^{-1} = \ln(\lambda_1/|\lambda_2|).$$

At $k_B T = J$ and zero bias, $\lambda_1 = 4.71828,\ \lambda_2 = 1.71828$ and $\xi = 0.9900$ sites. At low temperature e^K is large and $\xi \simeq e^K/3$. A finite chain can then appear ordered simply because it is shorter than its correlation length. There is no finite positive temperature at which these two one-dimensional eigenvalues cross.

Keep track of a mode that a chosen detector cannot see

At $h \neq 0$, colours one and two remain interchangeable. Their difference vector is an eigenmode of eigenvalue $e^K - 1$. In the orthonormal basis comprising colour zero and the sum of one and two, the remaining block is

$$\begin{pmatrix} e^{K(1+h)} & \sqrt{2}e^{Kh/2} \\ \sqrt{2}e^{Kh/2} & e^K + 1 \end{pmatrix}.$$

Its eigenvalues follow from the trace and the square root of the squared diagonal difference plus four times the squared off-diagonal element. Sort all three eigenvalues. A detector sensitive only to whether the spin is zero is invariant under interchange of colours one and two. It has no overlap with their odd difference mode, even if that mode sets the mathematically longest length. A resolved measurement of colours one and two can see it. Quoting a correlation length from eigenvalues requires specifying which observation can actually couple to that eigenvector.

Change temperature and examine the spectral gap

Vary reduced temperature and reduced field H/J independently. The plot compares the longest colour-correlation length to the exact zero-field result. At a chosen point the downloadable model constructs the full three-by-three matrix and computes its modes using OneSimulation's contour eigensolver. It reports free energy per site and all three eigenvalues. In this small system the direct reduction above is available as an independent check. The same matrix construction generalises to more states even when no colour symmetry reduces the size.

Compare two chains at the same temperature. Set the colour bias to zero and apply a positive bias of 0.30 J. Use the transfer-matrix eigenvalue ratio to explain the change in the longest correlation length.

Open this experiment online

Try the calculation

Reduced temperature kBT/J **1**

Reduced colour field H/J **0.15**

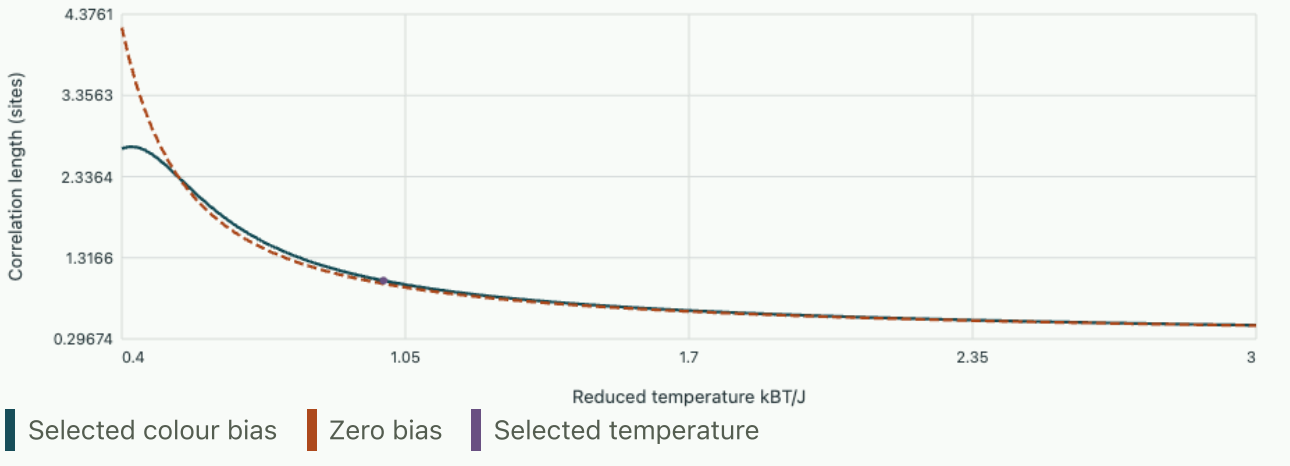

Leading eigenvalue 4.9881, second 1.8884, longest length 1.0295 sites.

Download the OneSimulation model

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

```
python spectra-and-orbits.py chain --temperature 1 --bias 0.15
```

Tutorial

At $k_B T = J,\ h = 0$ what fraction of the asymptotic leading contribution to a connected correlation survives after three lattice sites?

▼ **Work through the solution**

The factor is $(1.71828/4.71828)^3 \simeq 0.0483$. The correlation observable sets its initial amplitude; the eigenvalue ratio sets this fraction.

Why does squaring a single bond weight to represent the effect of two bonds give the wrong two-step transition amplitude?

▼ **Work through the solution**

Two steps can pass through any of the three intermediate colours. Matrix multiplication sums all possibilities $M_{ab}M_{bc}$ over b. Squaring the single number M_{ac} loses those intermediate configurations.

Suppose only the leading eigenvalue is reliably resolved. Which thermodynamic-limit prediction is possible and which remains undetermined?

▼ **Work through the solution**

One can obtain free energy per site, and with its parameter dependence the derivatives giving energy or magnetisation. A connected spatial length requires a subleading eigenvalue and overlap of the chosen observable. A precise leading free energy alone cannot give that numerical length.

Compare this transfer calculation with the two-state Ising chain experiment.

Counting, correlations and thermal measurements · 2. Measure the heat capacity of a reservoir from the width of an energy histogram

Two small Einstein solids share a fixed number of excitation quanta. A calorimeter records how many are in the first solid. Count the allowed microstates, compare the exact distribution with a reservoir approximation and see how energy conservation correlates the readings.

Count the ways to store an integer number of quanta

Let solid A contain a independent harmonic modes and solid B contain b, all with energy spacing ε. Omit their fixed zero-point energy. The total excitation count Q is fixed. In one solid a configuration is an ordered list of nonnegative mode occupations. An arrangement of q identical excitation symbols and $a-1$ separators specifies such a list. The multiplicity is $\binom{q+a-1}{a-1}$. The count for distinguishable packets would be different, because exchanging two excitations between identical occupations introduces no new configuration here.

$$P(q) = \frac{\binom{q+a-1}{a-1}\binom{Q-q+b-1}{b-1}}{\binom{Q+a+b-1}{a+b-1}},$$

$$q = 0, \dots, Q.$$

The denominator counts all occupations of the $a+b$ modes; summing the numerator over every partition must reproduce it. As a simple check, set $a=2, b=1, Q=2$. The unnormalised weights for $q=0,1,2$ are 1, 2 and 3. The most probable count is 2, but the mean is $4/3$. “Thermal equilibrium” does not require every observed energy to coincide with the most likely energy.

Separate the mean from the fluctuations

By permutation symmetry each individual mode holds mean $Q/(a+b)$. Therefore $\langle q \rangle = Qa/(a+b)$. For the second moment, introduce an auxiliary variable x marking quanta in solid A. The coefficient of z^Q in $G(x,z) = (1-xz)^{-a}(1-z)^{-b}$ is the unnormalised probability generating function. Differentiating once in x weights by q, twice by $q(q-1)$. At $x=1$ the second derivative is $a(a+1)z^2(1-z)^{-(a+b+2)}$. Divide its coefficient by the total multiplicity and combine first and second moments. Writing $s = a+b$ gives

$$\operatorname{var}(q) = \frac{Qab(s+Q)}{s^2(s+1)}.$$

The second solid contains exactly $Q-q$. Its variance is identical, and the covariance of the two measured energies is minus their common variance in energy units. A pair of independently fitted canonical distributions cannot reproduce this anticorrelation. At $Q=24,\ a=4,\ b=8$ the mean number in A is 8 and the variance is about 14.77. These are fluctuations of a perfectly isolated combined system, not a leak in the calorimeter.

Background · Can two noisy calorimeters expose the true energy variance?

The two solids exchange energy while their sum is fixed. If the first energy fluctuates upward by δE, the other fluctuates downward by exactly the same amount. Let the two recorded energies be $X = E_1 + n_1$ and $Y = E_2 + n_2$. Suppose each readout noise is independent of the physical energies and of the other noise. The same-channel variance and the cross covariance then measure different combinations.

$$\operatorname{var}(X) = \operatorname{var}(E_1) + \operatorname{var}(n_1),$$
$$\operatorname{cov}(X, Y) = -\operatorname{var}(E_1).$$

A single histogram that is too broad can mimic a reservoir with different heat capacity. Simultaneous measurements in the two channels allow the exchanged-energy variance to be recovered from the negative covariance. A drifting bath or common electronic pickup violates the assumptions. The measured variance of $X + Y$ is a direct diagnostic for such extra fluctuations when the readout noise levels are known.

Compare with a canonical ensemble at the same mean

In a canonical description each mode has a geometric distribution. Match its mean to $\bar{n} = Q/s$. A single mode then has variance $\bar{n}(1 + \bar{n})$, and a independent modes have variance $a\bar{n}(1 + \bar{n})$. Dividing the exact microcanonical variance by this canonical variance leaves $b/(s + 1)$. At the example point this is $8/13$. Increasing a along with b at fixed fraction does not make a macroscopic subsystem independent of the reservoir; the ratio stays below one. Holding a fixed, taking b large and keeping Q/s fixed does recover the canonical width.

The negative-binomial canonical distribution has support on every nonnegative integer, even those above Q. This is another sign that it describes another experiment. Merely deleting the tail without changing its other weights does not recover the true finite-reservoir conditional probability.

Calculate the distribution without enormous factorials

The ratio of successive unnormalised weights is $(q + a)(Q - q)/[(q + 1)(Q - q + b - 1)]$. Start with a positive arbitrary initial weight, march through the allowed q and normalise once. The graph compares this exact probability to a canonical law at the same mean. Adjust reservoir size and total energy; use the displayed first two moments to distinguish a shift in the average from a change in the width. The plotted canonical values cover the physically accessible count interval, so their displayed partial sum can be less than one.

Open this experiment online

Try the calculation

Reservoir modes 8

Excitation quanta 24

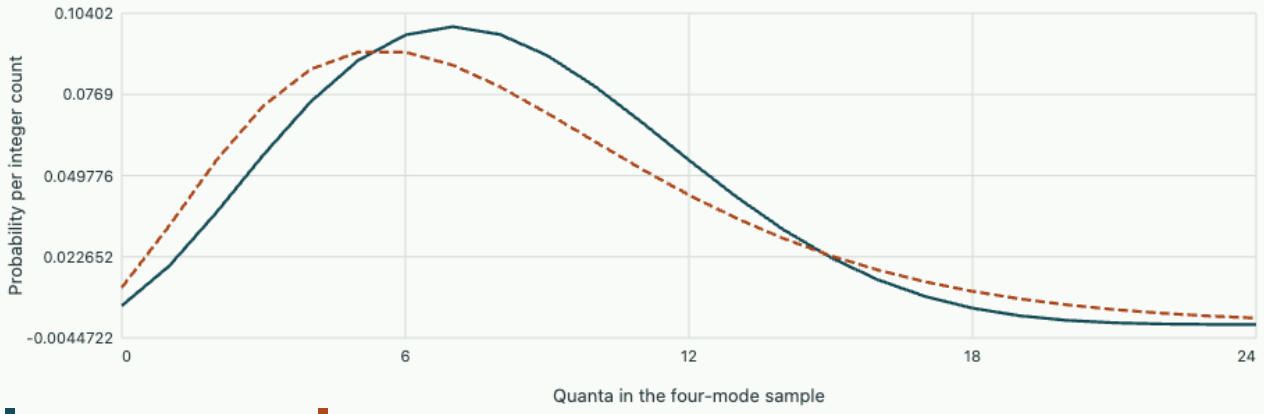

Isolated pair of solids Canonical sample

Mean 8 quanta; closed-pair variance 14.769. Canonical variance 24.

Tutorial

If A and B contain identical mode counts, is $Q/2$ necessarily the most likely observed value?

▼ **Work through the solution**

The distribution is symmetric. For odd Q there is no integer value at exactly $Q/2$; two equally likely central counts are possible. With one mode in each solid, every admissible value is equally likely. Symmetry fixes the mean exactly even when it does not identify a unique most probable observation.

For $a = 4, b = 8, Q = 24$ what variance would an unconditioned canonical fit predict?

▼ **Work through the solution**

Here $\bar{n} = 2$. The canonical count variance is $4(2)(3) = 24$. Multiplying by $8/13$ returns 14.769 for the closed pair. Count-to-energy conversion multiplies variances by ε^2.

An experiment finds a noisy sum of the two reported energies. What can one infer before relaxing the energy-conservation model?

▼ **Work through the solution**

In the ideal model $E_A + E_B$ is constant. Check both readouts' measurement noise and energy calibration. Independently fluctuating detector errors add variance in the sum even for a closed sample. Record a reference configuration or estimate detector noise from repeated measurements of the same state.

The finite-reservoir chapter connects this exact count to the curvature of the entropy.

Counting, correlations and thermal measurements · 3. Recover a free-energy difference from a distribution of work

A two-state system starts in equilibrium, then a control jumps its energy gap. Individual work readings differ although every run uses the same control protocol. Use the elementary model to derive an exponential work identity, compare it with the average work and estimate how many runs it takes to see the events that dominate the estimate.

Specify how work is measured

Set the lower energy to zero and the upper energy initially to ϵ_i. Prepare a canonical ensemble at inverse temperature β, isolate it and suddenly change the gap to ϵ_f without changing the energy eigenvectors. A run that started low stays low and requires zero work. A run that started high stays high and requires work $\epsilon_f - \epsilon_i$, positive if the external control gives energy to the sample. No heat flows during the jump. Returning the system to equilibrium afterwards can exchange heat, but it is a separate part of the protocol.

The initial high-state probability is

$$p_i = e^{-\beta\epsilon_i}/(1 + e^{-\beta\epsilon_i}).$$

Therefore

$$\langle e^{-\beta W}\rangle = (1 - p_i) + p_i e^{-\beta(\epsilon_f - \epsilon_i)} = \frac{1 + e^{-\beta\epsilon_f}}{1 + e^{-\beta\epsilon_i}} = \frac{Z_f}{Z_i}.$$

Taking a negative logarithm and multiplying by $1/\beta$ gives the equilibrium free-energy difference of two ensembles at the same temperature. The actual final population after the jump is still p_i, which generally differs from the final equilibrium occupation. This distinction explains why an equilibrium state function can be recovered from a nonequilibrium final distribution.

Compare the two averages

The mean work is $p_i(\epsilon_f - \epsilon_i)$. Convexity of the exponential gives

$$\langle W\rangle \geq \Delta F.$$

In units of $k_B T$, take an initial gap of 4 and a final gap of 0. The high state occurs with probability 0.01799 and gives work −4 when it occurs. Mean work is about −0.07194, yet $\Delta F = -\ln[2/(1 + e^{-4})] \simeq -0.6750$. Both numbers are negative. The greater number is the one closer to zero. Checking only the sign of average work would not reveal the large difference.

Each rare upper-state event contributes e^4 to the exponential average. Its probability times its weight is nearly one, comparable to the entire contribution of the common lower state. This is why replacing an average of exponentials by an exponential of an average loses essential information. Repeating the same protocol more slowly with contact to a bath would describe a different path through state space.

Ask what a finite dataset can support

For N independent preparations, the probability of seeing no initial upper state is $(1-p_i)^N$. With 20 runs in the example this is about 0.696. In such a dataset every measured work is zero and the empirical free-energy difference is zero. This is fully compatible with the model. To have at least a 95 percent chance of seeing even one upper-state event requires $N \geq \ln(0.05)/\ln(1-p_i)$, or about 166 runs. Seeing only one such event still gives a noisy estimate of its frequency.

The graph holds the final gap at zero. Change the initial gap and ensemble size. It shows the probability of missing the rare branch at each sample size, together with the five-percent level. The reported exact free energy and mean work use all initial states. These predictions can be checked by explicitly enumerating a two-outcome histogram, without drawing a random sample and silently treating its fluctuations as a new law.

Open this experiment online

Try the calculation

Initial gap (kBT) 4

Independent preparations 20

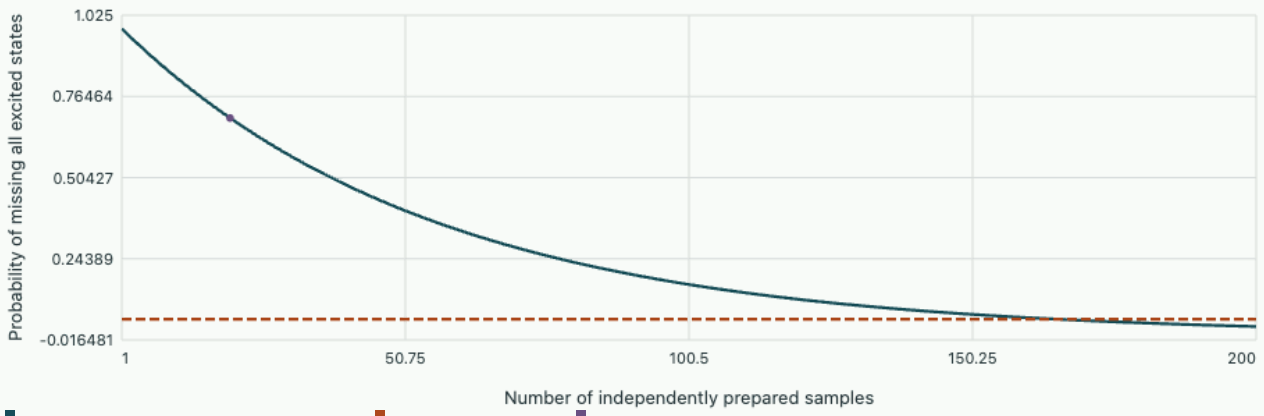

P(no rare work sample) = 0.69559. Mean work / kBT = -0.071945, ΔF/kBT = -0.675.

Tutorial

For an initial gap zero and a final gap $2k_BT$, calculate $\langle W \rangle$ and ΔF.

▼ **Work through the solution**

Initially $p_i = 1/2$. The average work is k_BT. The free-energy difference is $k_BT \ln[2/(1+e^{-2})] \simeq 0.5662 k_BT$. Some energy supplied by the control is unavailable as recoverable equilibrium free energy at this temperature.

A preparation procedure always starts the sample in its lower state independent of the specified initial gap. Can the same equality be used by changing only the sample count?

▼ **Work through the solution**

No. The canonical initial probabilities were used algebraically to turn the weighted sum into Z_f/Z_i. With deterministic preparation the exponential average is one in this protocol. Repeated measurements approach that different expectation.

At fixed initial and final gaps, what does reducing temperature do to the number of runs required to see an initial excitation?

▼ **Work through the solution**

The excitation probability is approximately $e^{-\beta\epsilon_i}$ for a large positive reduced gap. Required sample count grows approximately as $1/p_i$ for any fixed detection probability. Controlling preparation and understanding tails become more important as temperature falls.

Connect this model to the partition function and to the role of free energy in a reversible isothermal process.

Counting, correlations and thermal measurements · 4. Why a molecular beam and its source have different measured speeds

A tiny hole lets a dilute gas escape into a vacuum. A rotating chopper admits a short pulse and a downstream detector records arrival times. Derive the source speed density, the effusive speed density and the actual time-of-flight histogram. Each samples the same Maxwellian physics with a different weighting.

State what crosses a surface

For a gas with particle mass m and temperature T, write $a = m/(2k_BT)$. Cartesian velocity components have density proportional to $e^{-a|\mathbf{v}|^2}$. Integrating over a velocity shell gives an equilibrium speed density proportional to $v^2e^{-av^2}$. Sampling gas particles uniformly in a volume measures this density. For particles crossing a plane aperture in unit time, a particle with normal speed $v_z > 0$ can reach the plane from a depth $v_z\,dt$. Multiply by this flux factor and integrate over directions in the outward hemisphere. In polar coordinates about the aperture normal the angular integral of $\cos\theta\,\sin\theta$ is independent of speed and an extra factor v remains.

$$p_{\text{gas}}(v) = \frac{4a^{3/2}}{\sqrt{\pi}}v^2e^{-av^2},$$

$$p_{\text{beam}}(v) = 2a^2v^3e^{-av^2}, \qquad v > 0.$$

This model assumes a hole small compared with the molecular mean free path, negligible collisions between aperture and detector, and acceptance that selects a fixed narrow angular cone with speed-independent efficiency. A nozzle that produces an expanding collisional jet needs a different velocity distribution.

Transform speed into arrival time

Let L be the flight distance and $t = L/v$ the transit time along the accepted direction. Probability conservation for a histogram bin requires

$$p_t(t)\,dt = p_{\text{beam}}(v)|dv|.$$

The derivative $|dv/dt| = L/t^2$ is essential:

$$p_t(t) = 2a^2L^4t^{-5}e^{-aL^2/t^2}.$$

The mode of the equilibrium speeds is $1/\sqrt{a}$, that of beam speeds is $\sqrt{3/(2a)}$, and the mode of arrival times is $L\sqrt{2a/5}$. The speed calculated as L/t_{mode} is $\sqrt{5/(2a)}$; it is neither of the speed modes. Peaks of densities are not invariant under nonlinear changes of variable. Integrated probabilities for corresponding intervals are invariant. Changing the horizontal labels on a histogram while retaining its original bin heights gives the wrong likelihood.

Fit a temperature from counted arrivals

Define thermal speed $v_T = \sqrt{2k_B T/m}$ and time $t_0 = L/v_T$. The dimensionless variable $s = t/t_0$ has density $2s^{-5}e^{-1/s^2}$. An argon atom of mass 39.95 atomic mass units at 300 K has $v_T \simeq 353\,\mathrm{m\,s^{-1}}$. For a flight of 0.50 m the most common arrival time is about 0.895 ms; interpreting its reciprocal as the most probable gas speed would overestimate the temperature by a factor of 2.5. Finite gate width convolves the ideal distribution with the opening-time distribution. Record the gate profile or vary L to separate instrument timing from thermal spread.

Rather than fitting one maximum, consider ideal independent arrival times t_i. The a-dependent part of the log likelihood is $2N\ln a - aL^2\sum_i t_i^{-2}$. Differentiation gives

$$\hat{a} = 2N/[L^2\sum_i t_i^{-2}].$$

Consequently $\hat{T} = mL^2\sum_i t_i^{-2}/(4Nk_B)$. This estimator is just a beam-weighted mean squared speed divided by the coefficient appropriate to an effusive beam. A velocity-dependent detector efficiency would alter that coefficient.

Compare density shapes in the measurement coordinate

Change source temperature and flight length. The graph compares arrival times for particles sampled from a volume with arrivals from the effusive flux, both transformed with their correct Jacobians to equal-width time bins. Compare the most probable times and widths. Diluting an ideal gas at fixed temperature changes the signal rate, while the normalised ideal shapes stay the same. A shift of shape with pressure would suggest collisions or a change in how the source is sampled.

Open this experiment online

Try the calculation

Argon temperature (K) 300

Flight length (m) 0.5

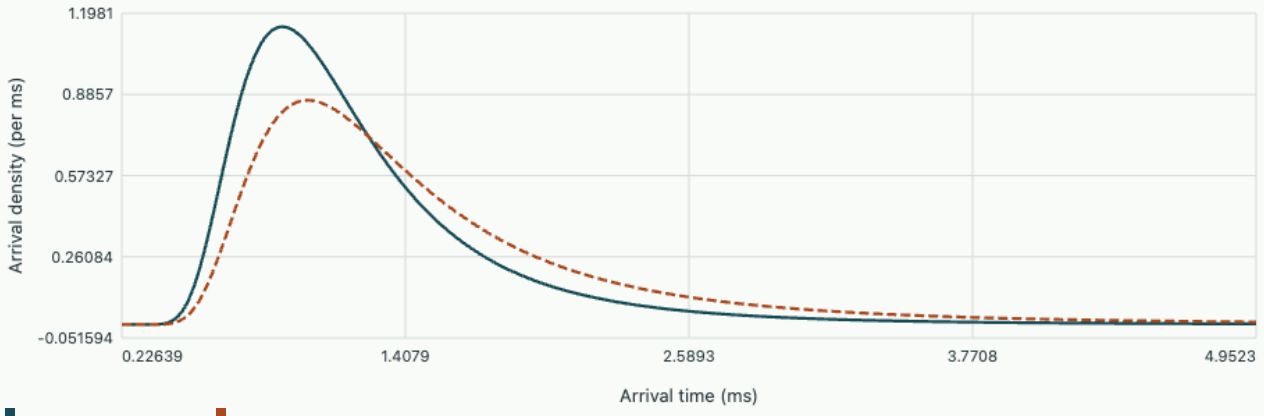

Aperture flux | Volume-sampled source

Thermal speed 353.37 m/s. Most probable effusive arrival 0.89488 ms. Interpreting L/t as a gas speed would give 558.73 m/s.

Tutorial

A detector records twice the flight time after doubling the path length while all dimensionless histogram quantiles coincide. Which interpretation is consistent with the data?

▼ **Work through the solution**

Flight is ballistic at unchanged speed distribution. Under this model t_0 scales linearly with L. If a delay were dominated by an electronic timing offset it would not scale this way.

Suppose the detector efficiency is inversely proportional to speed across the relevant band. How does this affect the effusive weighting?

▼ **Work through the solution**

The source flux multiplies the gas density by v; the detector multiplies the result by $1/v$. After normalisation the detected speed density has the same shape as the gas density. The time histogram still requires the Jacobian L/t^2.

Why is averaging the directly transformed speeds L/t_i different from dividing flight length by the mean arrival time?

▼ **Work through the solution**

The reciprocal is nonlinear. The latter is related to a harmonic mean of speeds. A broad distribution enhances the difference, even with perfect timing.

Compare a phase-space sample with a flux measurement in the kinetic theory chapter.

Counting, correlations and thermal measurements · 5. Detect irreversible steady motion from transition counts

A three-state system appears equally likely to occupy each state. Is it at thermal equilibrium? Build two different stationary processes with identical populations, then determine what a time-resolved instrument must count to distinguish them.

Construct a stationary distribution without assuming equilibrium

Arrange states $0, 1, 2$ on a directed ring. From each state the clockwise rate is k_+, the anticlockwise rate is k_-, both positive. For example $\dot{p}_0 = k_+ p_2 + k_- p_1 - (k_+ + k_-) p_0$, with the other equations obtained by cycling the indices. Uniform $p_i = 1/3$ makes every derivative vanish for any values of the two rates. The normalised stationary populations contain no information on their ratio.

Detailed balance would require $p_i k_{i\to j} = p_j k_{j\to i}$ across each individual edge. Here the stationary net clockwise current on each edge is $J = (k_+ - k_-)/3$. Summing over all edges gives mean signed jump count per time $k_+ - k_-$. One complete rotation contains three net clockwise jumps. A stationary histogram does not determine whether those currents exist because a divergence-free current leaves each state's population unchanged.

Read the direction of time from a trajectory

At stationarity compare the probability of one clockwise jump across an edge with its reverse. Their logarithmic ratio is $\ln(k_+/k_-)$. Multiplying this increment by k_B and the net flux, then summing the three oriented edges, gives the irreversible entropy-production rate for the observed Markov model.

$$\frac{\dot{S}_{\mathrm{prod}}}{k_B} = (k_+ - k_-)\ln(k_+/k_-) \geq 0.$$

Both factors change sign together if the circulation is reversed. There is no $1/2$ here because each undirected edge was included once with its chosen orientation. Counting all ordered pairs instead would require compensating for double counting. An interpretation in terms of actual heat delivered to a reservoir requires thermodynamic information such as local detailed balance. The transition record alone measures time-reversal asymmetry of this model.

At $k_+ = 6\,\mathrm{s}^{-1}$, $k_- = 2\,\mathrm{s}^{-1}$, each edge carries $4/3$ net jumps per second. The entropy-production rate is

$$4\ln 3 = 4.394\,k_B/\mathrm{s}.$$

An unbiased process with $k_+ = k_- = 4\,\mathrm{s}^{-1}$ has the same uniform occupations and the same total jump rate of eight per second, yet zero net current and zero entropy production. Counting only total events is insufficient too.

Resolve what a slow camera misses

Because every state has the same outgoing rates, the number of clockwise and anticlockwise jumps in duration t can be constructed from independent Poisson processes of means $k_+ t, k_- t$. Their difference has

mean $(k_+ - k_-)t$ and variance $(k_+ + k_-)t$. At the example rates a 10 s record has expected signed count 40 and standard deviation $\sqrt{80} \simeq 8.94$. If the electronics wraps angles modulo three, a sequence of three consecutive clockwise jumps between frames looks like zero displacement. High temporal resolution and a sufficiently long record solve different experimental problems.

Parameterise the rates by total activity $r = k_+ + k_-$ and bias $b = (k_+ - k_-)/r$. Then $\dot{S}_{\mathrm{prod}}/k_B = rb\ln[(1+b)/(1-b)]$. Close to equilibrium this grows as $2rb^2$, while current grows linearly in b. A weak current may be detectable in a signed record when estimates of an entropy-production rate remain noisy. The logarithmic estimate from raw counts is particularly unstable if a short record has no detected backward jumps.

Try equal stationary histograms with different histories

Change the bias and total transition rate. The graph gives the expected unwrapped signed count and one-standard-deviation band. The predicted occupation of each state stays fixed at one third. Reverse the bias and compare the mean drift, fluctuation width and entropy production. This experiment plots analytical ensemble statistics, so repeated changes of the controls give reproducible values.

Open this experiment online

Try the calculation

Clockwise bias **0.5**

Total jump rate (s⁻¹) **8**

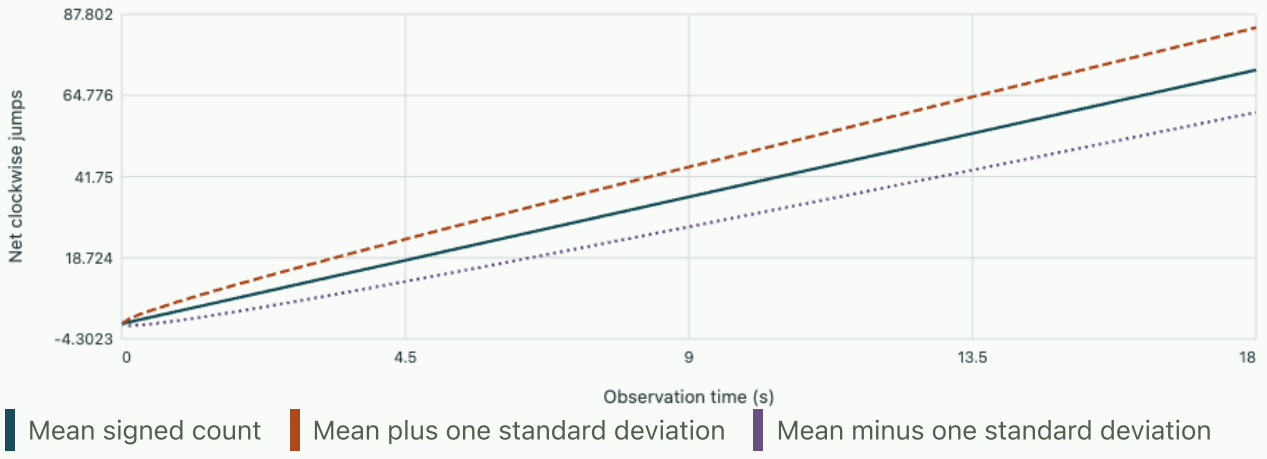

Every stationary population is 1/3. Net clockwise edge current 1.3333 s⁻¹. Entropy production 4.3944 kB/s.

Tutorial

A stationary 100 s record contains 480 clockwise and 320 anticlockwise jumps across all edges. Estimate the two rates and the entropy production under the stated homogeneous model.

▼ **Work through the solution**

The rate estimates are 4.8 and 3.2 per second. The inferred rate divided by k_B is $1.6 \ln(1.5) = 0.649\,\mathrm{s}^{-1}$. Estimating separate edge rates for an inhomogeneous ring would also require time spent in each source state.

Can an observed two-state continuous-time chain sustain a nonzero stationary edge current with one reversible edge?

▼ **Work through the solution**

Stationarity for either state directly sets the net edge current to zero. Hidden states or multiple unresolved transition channels can still dissipate. Absence of a current in a coarse record does not determine every microscopic flux.

How long must the idealised process with rates 6 and 2 per second be observed for the expected drift to be five times its standard deviation?

▼ **Work through the solution**

Solve $4t = 5\sqrt{8t}$, giving $t = 12.5$ s. This signal-to-noise criterion compares an ensemble mean to its spread and does not guarantee a particular outcome in one record.

Relate this inference to equilibrium reversibility in the master-equation chapter.

Counting, correlations and thermal measurements · 6. Why macroscopic measurements become reproducible

The same collection of microscopic switches can teach us two different facts about a macroscopic sample. Its mean becomes predictable, and the fluctuations about that mean acquire a particular shape. Derive both statements, then find an experiment in which only the first prediction is reliable at the available sample size.

Start with an experiment that has only two outcomes

A detector records whether one two-level site is excited. Write $X_i = 1$ for an excitation at site i and $X_i = 0$ for its ground state. For now take the sites to be equivalent and independent, with fixed excitation probability $0 < p < 1$. Sites with identical level spacings in a heat bath can share a bias different from $p = 1/2$. For a nondegenerate gap ϵ in thermal equilibrium, the ratio $p/(1-p) = \exp(-\epsilon/k_B T)$ fixes the bias. Even when every site has the same preparation, a single record can contain any count from zero to N.

To obtain exactly k excitations, first specify one pattern with k occupied sites. Independence gives its probability $p^k(1-p)^{N-k}$. There are $\binom{N}{k}$ ways to select the occupied positions. Thus

$$\mathbb{P}(K = k) = \binom{N}{k} p^k (1-p)^{N-k},$$

$$K = \sum_{i=1}^{N} X_i.$$

Distinguish this probability from the multiplicity. With four sites and $p = 0.2$, the probability of exactly two excitations is

$$6(0.2)^2(0.8)^2 = 0.1536.$$

Six is a count of patterns. The factor multiplying it carries information about the physical preparation. A change of temperature changes that factor even though the available patterns stay the same.

Derive a law of large numbers from one inequality

Linearity of expectation gives

$$\mathbb{E}K = Np,$$

with no requirement of independence. For the variance, expand the square of the deviation as the sum of all pair products. Distinct sites give zero covariance under the present assumptions. Each diagonal term contributes $p(1-p)$, because $X_i^2 = X_i$. Consequently the measured excited fraction has variance $p(1-p)/N$. The width of the count grows with N. The width of the fraction shrinks.

Choose an absolute tolerance $a > 0$. On the event that the measured fraction differs from p by at least a, its squared deviation is at least a^2. Keeping only that event in the expectation of the nonnegative square gives

$$\mathbb{P}(|K/N - p| \geq a) \leq \frac{\mathbb{E}(K/N - p)^2}{a^2} = \frac{p(1-p)}{Na^2}.$$

At fixed a the upper bound tends to zero. This proves convergence in probability for the measured fraction. It gives a conservative guarantee even if little else is known about its distribution. At $p = 1/2$, $a = 0.05$, $N = 1000$, the bound is 0.1. It is a limit on a probability, and substituting it as the probability of failure would discard information we actually possess about these sites.

Find the shape hidden by the concentration bound

For the count itself, magnify deviations on their natural scale. Set $\sigma = \sqrt{p(1-p)}$, $Y_i = (X_i - p)/\sigma$ and $Z_N = N^{-1/2} \sum Y_i$. Now $\mathbb{E}Y_i = 0$ and $\mathbb{E}Y_i^2 = 1$. The characteristic function $\phi(t) = \mathbb{E}e^{itY_i}$ encodes moments in its Taylor series. In this bounded two-state example its first three terms yield

$$\phi(t) = 1 - t^2/2 + O(|t|^3).$$

The missing linear term follows from centring. At $p = 0$ or $p = 1$ the count is already fixed and its zero standard deviation makes this rescaling undefined.

Independence turns the characteristic function of the sum into a product. Each argument gets divided by $\sqrt{N}$, while each of the N factors is identical. Taking logarithms turns this product into

$$\ln \mathbb{E}e^{itZ_N} = N \ln \phi(t/\sqrt{N}) = -t^2/2 + O(|t|^3/\sqrt{N}) \quad \text{for fixed } t.$$

The limit $e^{-t^2/2}$ is the characteristic function of the unit normal density. The continuity theorem for characteristic functions therefore gives convergence in distribution to a standard normal. This statement concerns the probabilities of scaled intervals. A normal random variable assigns zero probability to any single value, whereas a finite count assigns positive probability to integer bins. To compare them in an experiment, integrate the normal density over each bin from $k - 1/2$ to $k + 1/2$ in the original count coordinate. Probability outside the physical range is a visible defect of the approximation for small samples.

The first correction predicts which experiment will look asymmetric

For one Bernoulli site the centred third moment equals $p(1-p)(1-2p)$. In the sum, terms with any singly occurring independent centred factor vanish. Thus the third centred moment of the count grows as N, while its standard deviation cubed grows as $N^{3/2}$. Their ratio is

$$\gamma_1 = \frac{1 - 2p}{\sqrt{Np(1-p)}}.$$

A rare excitation produces a longer upper tail. For $N = 20, p = 0.05$, even an entirely empty detector has probability $0.95^{20} \simeq 0.3585$. Here a symmetric curve extends into a physically impossible negative count and may miss the leftmost bin. At $N = 20, p = 1/2$ the skewness vanishes exactly by reflection symmetry. Finite support and discreteness still matter, particularly in the far tails. Increasing N while reducing the excitation probability to $p_N = 1/N$ describes a different limit. The mean count stays at one and the empty-bin probability $(1 - 1/N)^N$ tends to e^{-1}. A Poisson count of mean one describes this rare-event limit.

Open this experiment online

Try the calculation

Independent sites N 20

Excitation probability p 0.2

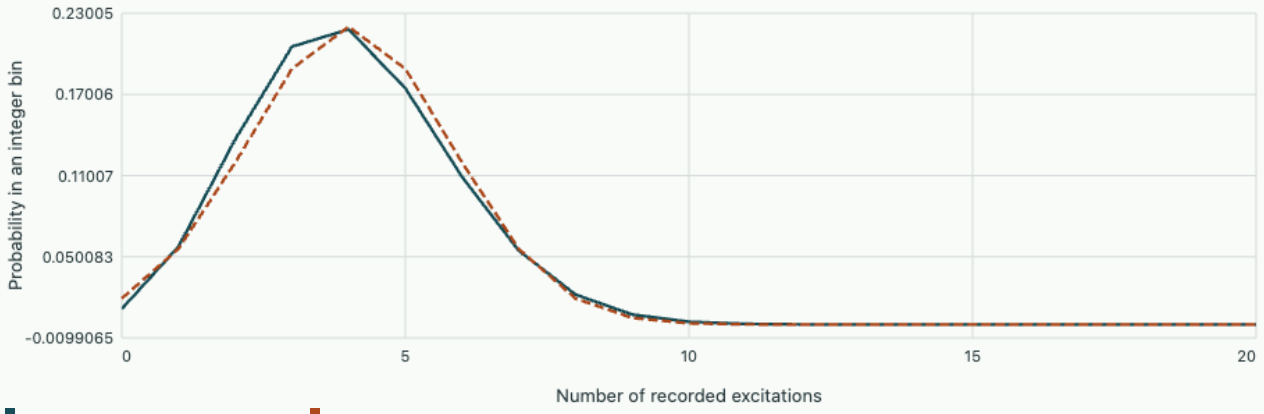

Exact binomial count Normal probability in the bin

Mean 4, standard deviation 1.7889. Largest absolute bin error 0.017167 at count 2. Skewness 0.33541.

Compare a detector with one expected excitation and a detector with twenty. The displayed normal-bin probabilities use the same mean and variance as the exact count. Identify the bin with the largest absolute discrepancy. A small central discrepancy alone gives little information about a very small tail probability.

Changing the physics can change the scaling

Now imagine observing the same slowly fluctuating site at evenly spaced times. Centring removes its mean without removing memory. If the sequence is stationary with lag-h covariance $C(h)$, then expanding the variance of the sample mean once more gives

$$\operatorname{var}(\bar{X}) = \frac{1}{N}\left[C(0) + 2\sum_{h=1}^{N-1}(1 - h/N)C(h)\right].$$

There are $N - h$ pairs at separation h in either order. Positive correlations inflate the variance compared with independent observations. For $C(h) = C(0)\rho^h$, $0 \le \rho < 1$, and a record long compared with the correlation time, the bracket tends to $C(0)(1+\rho)/(1-\rho)$. With $\rho = 0.9$, roughly nineteen records contain the information of one independent record for estimating the mean. Dividing the same observation interval into more samples need not buy nineteen times the information.

A different failure comes from the single-site law itself. A standard Cauchy variable has characteristic function $e^{-|t|}$ and has no finite variance or mean. The characteristic function of the arithmetic average of N independent Cauchy readings is $(e^{-|t|/N})^N = e^{-|t|}$. The average has exactly the original spread at every sample size. Identify the physical observable and its tail behaviour before assigning the familiar square-root error bar.

Tutorial on counting, fluctuations and experimental design

1. Size an equilibrium sample. For sites prepared with $p = 0.2$, the measured excited fraction should have standard deviation no greater than 0.01. An instrument can instead measure repeated copies with the same gap. How many independent readings suffice, and how does this differ from a guarantee that every report is within 0.01?

▼ **Calculate the required number**

The variance criterion is $0.2(1-0.2)/N \le 10^{-4}$, giving $N \ge 1600$. It supplies one standard deviation. It leaves a finite probability of larger deviations. Chebyshev's inequality alone would require 32,000 readings to guarantee that the probability of a deviation of at least 0.01 is at most 0.05. If 0.2 is only a pilot estimate, a guarantee valid for any preparation uses $p(1-p) \le 1/4$, giving 2,500 and 50,000 readings for the two requirements.

---

2. Test a tail. At $p = 0.05, N = 20$, approximate the probability of no excitations in two ways. First calculate it exactly. Then use the normal probability contained between –0.5 and 0.5.

▼ **Separate the empty bin from a negative count**

The exact value is 0.35849. The mean is 1 and the standard deviation is

$$\sqrt{0.95} = 0.97468.$$

The standardised boundaries are about –1.539 and –0.513, yielding approximately 0.242. The same discrete event can be written as $K \le 0$. Applying a one-sided continuity correction to that cumulative probability gives $\Phi((0.5-1)/0.97468) \simeq 0.3040$. The difference of 0.0619 is the Gaussian probability below –0.5. The exact detector has no bins there. The plot integrates over individual finite bins, so its empty-bin value is 0.2421.

---

3. Identify shared noise. One hundred identical thermometers each have independent zero-mean noise of standard deviation 0.20 K. All share a random calibration offset with zero mean and standard deviation 0.10 K. What is the standard deviation of the mean reading under repeated calibration?

▼ **Keep the common contribution**

Write $X_i = T + B + \eta_i$, with B independent of the individual noises. The average contains B exactly once and the mean of all η_i. Its variance is $0.10^2 + 0.20^2/100 = 0.0104\,\mathrm{K}^2$, so the standard deviation is about 0.102 K. Collecting more readings under the same calibration can only shrink the second term.

---

4. Keep the units when converting a fluctuation. Fifty independent sites have an energy gap of 3 meV. At one temperature, one quarter of the sites are typically excited. Find the mean excitation energy of the whole sample and its standard deviation.

▼ **Convert the count to energy**

The energy relative to the all-ground state is $E = \epsilon K$. Hence

$$\langle E \rangle = 3(50)(0.25) = 37.5\,\mathrm{meV},$$

$$\sigma_E = 3\sqrt{50(0.25)(0.75)} = 9.19\,\mathrm{meV}.$$

The standard deviation of the sample energy per site is 0.184 meV. Which of these the detector reports determines whether the measured width grows or shrinks as additional sites are included.

Fließbach places mathematical statistics before the construction of the statistical ensembles. The publisher's contents for *Statistische Physik* identify this starting point. Continue to the finite-reservoir experiment to see how conservation couples the recorded counts.

Counting, correlations and thermal measurements · 7. Weigh a colloid by counting where it spends its time

Brownian motion spreads a dilute suspension through a cell whilst buoyancy biases its vertical position. Derive the equilibrium concentration, predict counts in finite camera bins and estimate a buoyant mass without measuring how quickly any individual sphere falls.

Balance drift against diffusion

Measure z upwards from the bottom of a closed cell. A sphere of volume V displaces fluid of density ρ_f. The downward gravitational force on the particle exceeds the upward buoyancy force by $m_b g$, where $m_b = V(\rho_p - \rho_f)$. Signed buoyant mass is what a concentration profile can reveal. Let μ denote mobility, the ratio of speed to applied force. For dilute, independently moving particles at uniform temperature, the total vertical number flux is

$$J = -D\frac{dn}{dz} - \mu m_b g n.$$

At steady state the flux is constant, and impermeable cell ends set that constant to zero. The Einstein relation $D = \mu k_B T$ eliminates the mobility. Integrating $d\ln n/dz = -a$ with $a = m_b g/(k_B T)$ gives

$$n(z) = n(0)e^{-az}.$$

The signed slope changes sign for buoyant particles and vanishes at exact density matching. Viscosity affects the time to reach equilibrium.

Suppose the camera counts centres in equal-width horizontal slices of a column with uniform cross-section. A bin centred at z_i has expected count proportional to $e^{-az_i}\int_{-h/2}^{h/2} e^{-au}du$. The integral is common to all bins, hence normalising the exponential values at their centres yields the exact bin probabilities for this model. Wall exclusion sets the interval accessible to centres. Define the specified cell height as that accessible interval.

Infer a mass with its sign

Use a radius of 0.50 µm, particle excess density 50 kg/m³ and temperature 293 K. The volume is 5.236×10^{-19} m³, the buoyant mass is 2.618×10^{-17} kg and $k_B T/(m_b g) = 15.75$ µm. Counts in two equal slices separated by 10 µm should have lower-to-upper ratio $e^{10/15.75} \simeq 1.89$. Doubling the radius at fixed density ratio multiplies the mass by eight and divides the gravitational length by eight.

Sampling is important when using the upper part of a steep profile. Suppose bins separated by 10 µm contain 120 and 60 approximately independent particle observations. The estimated slope is $\ln(2)/10 = 0.0693$ µm⁻¹. For independent Poisson counts the standard error from propagating the two count variances is $\sqrt{1/120 + 1/60}/10 = 0.0158$ µm⁻¹. Repeatedly imaging the same particle at closely spaced times introduces correlations. Counting all those frames as independent would make the quoted sampling precision too optimistic.

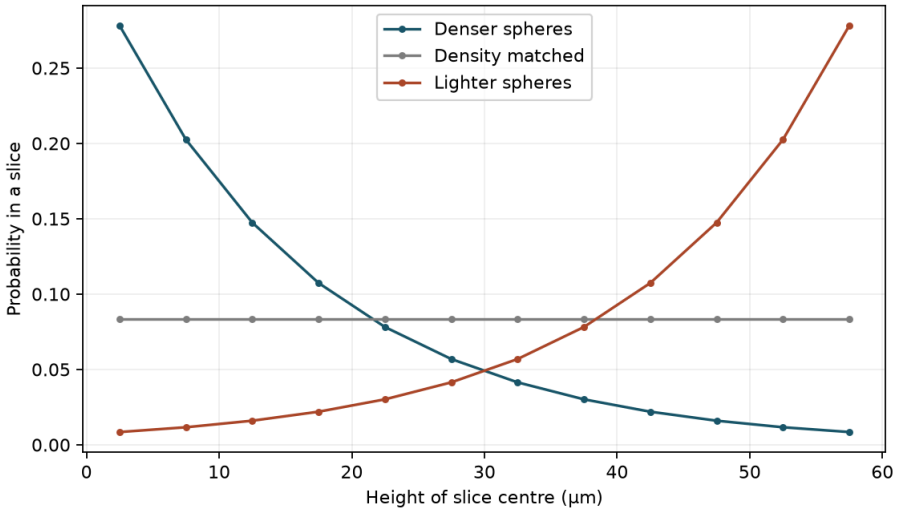

The total probability in the cell is one for each curve. Matching the fluid density removes the gravity-induced concentration gradient.

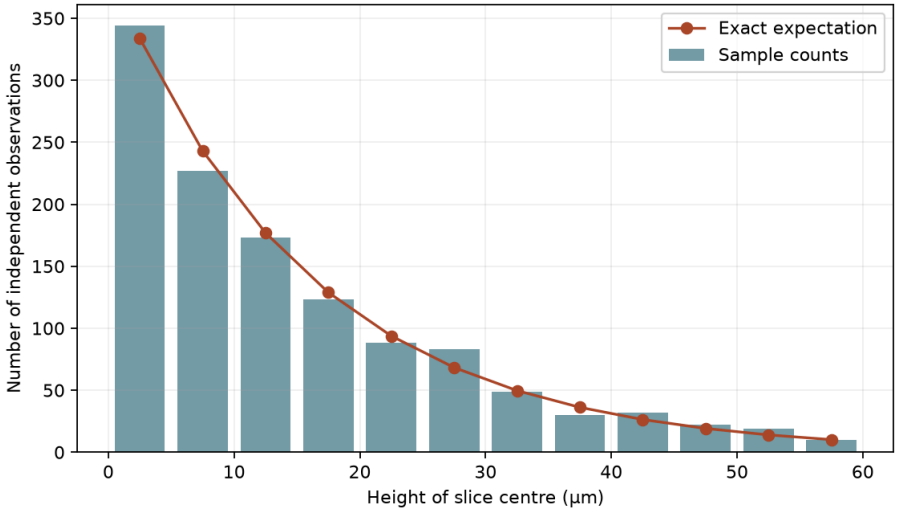

Reduce the number of independent particle observations and change the random seed. Noise is particularly large relative to the expected number in sparsely occupied slices.

Exercise · determine the size of a density-matched particle

A uniform profile persists in a well-equilibrated, isothermal cell. Can gravity alone distinguish two different radii if the particles have the same density as the fluid?

▼ Work through the solution

At exact density matching, $m_b = V(\rho_p - \rho_f) = 0$ for either radius, so both profiles are uniform. Their diffusion times can still differ because hydrodynamic mobility depends on size. A time-resolved measurement carries information absent from the equilibrium profile.

Exercise · reverse the measured gradient

Two equally wide slices separated vertically by 12 μm contain 80 and 160 observations, in order of increasing height. Estimate m_b at 293 K.

▼ **Work through the solution**

The signed slope is $\ln(80/160)/(12\,\mu\mathrm{m}) = -0.05776$ µm⁻¹. Converting to metres and multiplying by $k_B T/g$ gives $m_b \simeq -2.38 \times 10^{-17}$ kg. The sign indicates that the displaced fluid is heavier than the particle.

Try it · Sedimentation

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

Particle radius (µm)	0.5
Particle minus fluid density (kg/m³)	50
Independent observations	1200

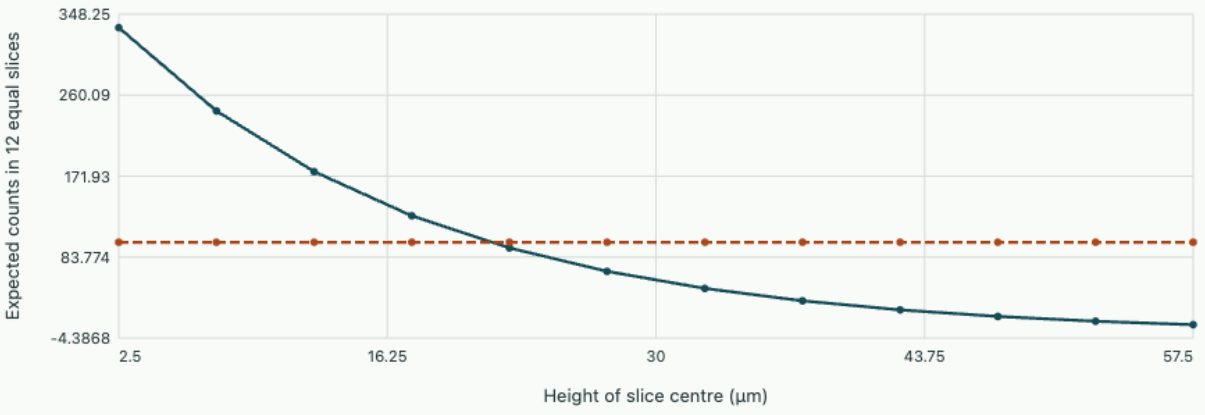

Temperature 293 K; signed gravitational slope 0.063466 µm⁻¹. Counts sum to 1200. The notebook adds seeded multinomial sampling to these expectations.

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

Reservoirs, useful work and mixtures

Keep track of which reservoir fixes temperature, chemical potential or pressure. Close the energy and entropy balances before inferring the useful work or the coexisting phase fraction.

1. State, process and the first law for a closed system
2. Open systems, enthalpy and steady-flow devices
3. Entropy, irreversibility and useful-work loss
4. Phase equilibrium and steam quality
5. Gas mixtures, humidity and condensation
6. Useful work and exergy destruction
7. Why a real gas cools or warms across a throttle
8. Finding the vapour fraction of an equilibrium flash
9. When does a polar crystal change its preferred state?
10. Moving constraints, boundary work and a changing energy

State, process and the first law for a closed system

Engineering thermodynamics and heat transfer

Distinguish properties from path-dependent transfers and determine heat and work for a piston-cylinder process.

FROM FIRST PRINCIPLES

Obtain closed-system work from the motion of a boundary

Choose a fixed mass of gas enclosed by a piston and cylinder. No mass crosses this boundary. The first law is the empirical balance of energy: with heat positive into the gas and work positive out, change of internal, kinetic and potential energy equals net heat minus net work. For a piston moving normally through dx against a uniform resisting pressure p_{ext}, force times displacement gives

$$\delta W_b = p_{\mathrm{ext}} A\, dx = p_{\mathrm{ext}}\, dV,$$
$$\Delta(U + KE + PE) = Q - W.$$

For slow internal equilibration with tiny pressure differences, one can use gas pressure in the integral. If bulk kinetic and potential changes are negligible and no shaft or electrical work crosses the boundary, the integrated heating requirement is

$$Q_{12} = U_2 - U_1 + \int_{V_1}^{V_2} p_{\mathrm{ext}}\, dV.$$

The area under the specified pressure–volume path is therefore part of the heat calculation. For a rigid vessel it vanishes. For a constant-pressure piston heating an ideal gas, $p\Delta V = mR\Delta T$ adds to the internal-energy change. The end temperatures alone determine the ideal-gas internal-energy change, while the moving boundary determines the work.

DERIVATION AND APPLICATIONS

Determine the heat input when a spring loads the piston

A freely sliding piston does not necessarily hold the gas pressure constant. Attach a linear spring to a frictionless piston of area A and let its resisting force increase as the gas expands. Consider a closed mass m of ideal gas, heated slowly enough that piston force balance holds at each intermediate state. Neglect piston kinetic energy, heat leakage and changes of elevation except as already included in the constant external pressure p_1. At the initial position the gas is in equilibrium at absolute pressure p_1 and volume V_1.

The displacement x produces a volume increment Ax and an extra resisting force kx. Dividing by piston area gives the pressure as a function of the evolving volume:

$$V - V_1 = Ax,$$
$$p(V) = p_1 + \frac{k}{A^2}(V - V_1).$$

Boundary work is performed at this changing external pressure. Integrating the straight-line path over the swept volume $\Delta V = V_2 - V_1$ separates work against the initial pressure from work stored in the spring:

$$W_{12} = \int_{V_1}^{V_2} p(V)\,dV = p_1 \Delta V + \frac{k}{2A^2}(\Delta V)^2.$$

For constant specific heat c_v and specific gas constant R, the first law gives the heat required, positive into the gas. State temperatures follow from the ideal-gas relation, including the final spring pressure:

$$Q_{12} = m c_v (T_2 - T_1) + W_{12},$$
$$T_i = p_i V_i / (mR).$$

Only the temperature change determines the internal-energy change; the mechanical law selects the work term. Setting $k = 0$ recovers expansion against constant pressure. Abrupt heating that accelerates the piston would invalidate the assumed quasistatic pressure path.

WORKED EXAMPLE

Calculate the heater input for an air cylinder

Change the inputs and test this calculation.

Take $m = 0.0100$ kg, $V_1 = 0.0100$ m³, $p_1 = 100$ kPa, $A = 0.0100$ m² and $k = 1000$ N/m. Increase volume by 0.00200 m³. The final pressure is 120 kPa; the work is

$$(100 + 120)(0.00200)/2 = 0.220\,\text{kJ}.$$

For $R = 0.287$ and $c_v = 0.718$ kJ/(kg K), $\Delta T = (1.440 - 1.000)/(0.0100 \times 0.287) = 153$ K. The gas receives $1.10 + 0.220 = 1.32$ kJ of heat.

SECOND APPLICATION

Separate actuator work from stored spring energy

Change the inputs and test this calculation.

On the same stroke the piston travels 0.200 m. The initially unloaded incremental spring stores $kx^2/2 = 20.0$ J, exactly the triangular part of the pressure–volume work. The constant load receives 200 J. If the spring were removed while retaining the same initial state and swept volume, the final pressure would stay at 100 kPa and the temperature rise would be only 69.7 K. The two processes end at different thermodynamic states despite having the same final volume.

Try it · Spring piston

Set the inputs to the values in the worked example, then vary one quantity while holding the others fixed. Compare the curve and numerical balance with the second application. The calculation uses the assumptions derived above; its readout identifies bounds such as choking, phase-region changes and inconsistent measurements.

Open this live calculation

Try the calculation

Spring stiffness (N/m) **1000**

Volume increase (litres) **2**

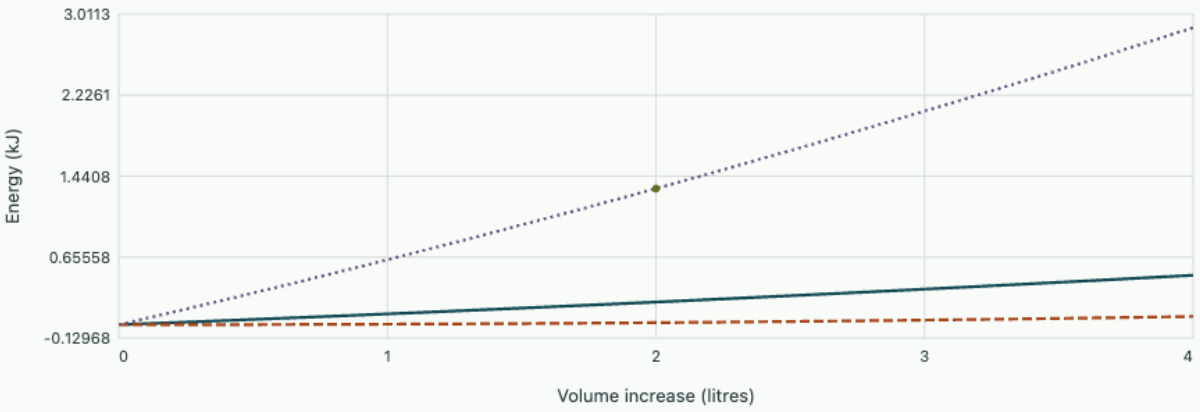

Boundary work | Incremental spring energy | Required heat | Selected heat

Air mass 0.010 kg, initial p=100 kPa, V=10 litres, piston area 0.010 m². Selected final pressure 120 kPa; temperature rise 153.31 K; work 0.22 kJ; heat 1.32077 kJ. Quasistatic ideal gas, cv=718 J/(kg K).

Decide what crosses the boundary

Draw a boundary around the gas in a piston cylinder. If no gas crosses its surface, the system is closed even though the piston moves. A rigid vessel is a special closed system with fixed volume. An insulated wall prevents heat transfer; a shaft or electrical lead can still transmit work through it. These descriptions impose different constraints and should be identified before choosing an equation.

At equilibrium a simple compressible substance can be described using independent state variables such as temperature and specific volume. Pressure p, volume V, temperature T and internal energy U are properties. An equation of state relates some of them. For a dilute ideal gas $pV = mRT$, where R is the specific gas constant. Always use absolute pressure and absolute temperature in this relation. Two initial and final states determine changes in properties, but do not specify the heat and work transfers between them.

Construct the process energy balance

Take work done by the system as positive and heat entering it as positive. The first law is

$$\Delta U + \Delta KE + \Delta PE = Q - W.$$

For a stationary vessel ignore changes in kinetic and gravitational energy of the system as a whole. A moving piston of area A against external pressure p_{ext} performs incremental work $p_{\text{ext}} A\,dx = p_{\text{ext}} dV$. For a quasistatic mechanically reversible motion the boundary pressure can be represented by the gas pressure throughout the process.

If gas expands from V_1 to V_2 at constant external pressure,

$$W = p_{\text{ext}}(V_2 - V_1).$$

If it instead follows a reversible isothermal ideal-gas path, inserting $p = mRT/V$ and integrating gives

$$W = mRT \ln(V_2/V_1).$$

Both expressions have units of energy. Their different volume dependence makes a pressure-volume diagram useful: the area under each pressure path has a direct work interpretation when its boundary-force assumptions hold.

Background · Which pressure enters the boundary-work integral?

The traction at a moving boundary determines the mechanical work across it. For a piston use the pressure exerted against the load. Replacing it by a uniform equilibrium gas pressure additionally assumes slow mechanical evolution. A limiting comparison makes this distinction measurable.

$$W_{\text{by}} = \int_{V_1}^{V_2} p_{\text{ext}}\, dV.$$

Remove a partition in an insulated rigid vessel so that an ideal gas expands into a vacuum. The external pressure against the advancing gas is zero. There is no heat transfer or boundary work, so the gas returns to the same equilibrium temperature. A reversible isothermal expansion between those same gas volumes requires heat from a bath and yields positive work $nRT \ln(V_2/V_1)$. The two endpoint states alone do not define a work measurement.

Heat air below a freely moving piston

A weight-loaded, frictionless piston keeps 0.50 kg of ideal air at 150 kPa absolute as a heater raises its temperature from 300 to 360 K. Use constant properties $R = 0.287$ and $c_v = 0.718$ kJ/(kg K). The ideal-gas law gives the volumes at the two ends of the process.

$$V_1 = \frac{mRT_1}{p} = 0.2870\ \mathrm{m^3},$$
$$V_2 = \frac{mRT_2}{p} = 0.3444\ \mathrm{m^3}.$$

The pressure is constant, so the work is the area of a rectangle on a pressure–volume diagram. Internal energy depends on the gas temperature. These are separate contributions to the heat input.

$$\begin{aligned} W &= p(V_2 - V_1) = 150(0.0574)\ \mathrm{kJ} = 8.61\ \mathrm{kJ}, \\ \Delta U &= mc_v(T_2 - T_1) = 0.50(0.718)(60)\ \mathrm{kJ} \\ &= 21.54\ \mathrm{kJ}. \end{aligned}$$

The required heat is

$$Q = \Delta U + W = 30.15\,\mathrm{kJ}.$$

The same result follows from $mc_p\Delta T$ because for this constant-pressure ideal-gas process with only boundary work $c_p = c_v + R = 1.005$ kJ/(kg K). The extra heat above the internal-energy increase moved the piston weight. Using $Q = mc_p\Delta T$ without inspecting the constraint would give the wrong heat input to a rigid tank.

Compare a rigid-wall experiment

Lock the piston at its initial volume before heating an identical mass from 300 to 360 K. Now $W = 0$, so only 21.54 kJ of heat is needed. The final pressure is

$$150(360/300) = 180\,\mathrm{kPa}.$$

Equal temperature changes in an ideal gas imply equal changes in internal energy, even though final pressures and energy transfers differ. For a real fluid near a phase boundary, constant ideal-gas heat capacities may fail; use tabulated state properties or a suitable equation of state. In either case the first-law accounting is unchanged once the initial and final internal energies and the actual boundary work are known.

A different path between specified volumes

An ideal gas initially at 300 kPa and 0.020 m³ expands quasistatically and isothermally to twice its initial volume. Since $pV = p_1V_1$, work out is

$$W = p_1V_1\ln(V_2/V_1) = 6\ln 2 = 4.16\,\mathrm{kJ}.$$

Ideal-gas internal energy depends only on temperature, so $\Delta U = 0$ and heat entering equals 4.16 kJ. A freely expanding ideal gas in an insulated rigid outer vessel can reach the same final temperature and occupied volume while doing zero boundary work. The state variables alone do not record which process occurred.

Determine heat and work along a measured compression path

A measured pressure–volume trace contains more information than its endpoints. For some slow piston processes an empirical polytropic law $pV^n = \text{constant}$ fits the path. Here n is a fitted process exponent and need not equal the gas heat-capacity ratio. Assume a closed cylinder with negligible kinetic and potential energy change, a quasistatic pressure load and a calorically perfect ideal gas. Work out of the gas is the area $W = \int_{V_1}^{V_2} p\,dV$, which is negative for compression. Substitute $p = p_1 V_1^n V^{-n}$ and integrate:

$$W = \frac{p_2 V_2 - p_1 V_1}{1 - n} \quad (n \neq 1),$$

$$W = p_1 V_1 \ln(V_2/V_1) \quad (n = 1).$$

The logarithmic case is the continuous limit as n approaches one. The ideal-gas equation also gives $T_2/T_1 = (V_1/V_2)^{n-1}$, so $\Delta U = m c_v (T_2 - T_1)$. With $Q - W = \Delta U$ and $pV = mRT$, the heat input is $Q = m[c_v + R/(1-n)](T_2 - T_1)$ for $n \neq 1$. Setting $n = \gamma = c_p/c_v$ eliminates Q. This inference relies on the ideal-gas, constant-property, quasistatic assumptions. An arbitrary irreversible compression with the same endpoints cannot be called isentropic from the fitted exponent alone.

Recover a cylinder heat loss from three measurements

A test traps air at $p_1 = 100\,\text{kPa}$ absolute, $V_1 = 0.0100\,\text{m}^3$, $T_1 = 300\,\text{K}$, then compresses it to $V_2 = 0.0040\,\text{m}^3$. The fitted exponent is 1.25. Using $R = 0.287,\ c_v = 0.718\,\text{kJ/(kg K)}$, the trapped mass is

$$p_1 V_1/(RT_1) = 0.01161\,\text{kg}.$$

The predicted pressure is

$$100(2.5)^{1.25} = 314.4\,\text{kPa},$$

$$T_2 = 300(2.5)^{0.25} = 377.2\,\text{K}.$$

Note that absolute pressure is required in both the fit and the gas law. A pressure transducer referenced to atmosphere would need its reference restored.

The gas does $W = [314.4(0.0040) - 100(0.0100)]/(1 - 1.25) = -1.03\,\text{kJ}$. Its internal energy rises by $0.01161(0.718)(77.2) = 0.644\,\text{kJ}$, giving $Q = \Delta U + W = -0.386\,\text{kJ}$. Compression work enters the gas while part leaves as heat through the cylinder wall. If the process were isothermal at 300 K, the same volume change would require only $-W = 0.916\,\text{kJ}$ and the full amount would leave as heat. A mid-stroke measurement gives a stronger check than the end pressure: the fitted law predicts $100(0.010/0.007)^{1.25} = 156.2\,\text{kPa}$ at 7.0 litres. A substantial discrepancy would favour direct numerical integration of the measured trace over using a single exponent.

An apparent exponent can also be estimated directly from pressure ratios:

$$n = \ln(p_2/p_1)/\ln(V_1/V_2).$$

Compare estimates from the first and second halves of the stroke. Agreement means the same log–log slope describes both halves, while disagreement may indicate changing heat transfer, leakage or phase lag between the pressure and position sensors. It does not identify which of those mechanisms is responsible. As

a separate check, suppose a water jacket gains the estimated 0.386 kJ each stroke and the compressor completes ten such strokes per minute. The mean heat removal from these strokes would be $0.386(10)/60 = 0.0643\,\mathrm{kW}$. A jacket measurement of the complete machine includes intake and exhaust processes and mechanical friction as well, so experimental comparison requires either an isolated compression interval or a balance around the rest of the cycle.

A change in operating speed can also change the fitted exponent because it changes the time available for heat exchange during a single compression stroke.

Open systems, enthalpy and steady-flow devices

Engineering thermodynamics and heat transfer

Use an inlet-outlet balance to distinguish shaft work, kinetic-energy conversion and pressure reduction.

FROM FIRST PRINCIPLES

Derive the steady-flow energy balance including flow work

Surround a turbine, nozzle or mixing chamber with a stationary control surface. In addition to transporting its internal, kinetic and gravitational energies, an entering fluid is pushed across the inlet by upstream pressure. A mass increment dm occupies volume $v\,dm$ and requires pressure work $pv\,dm$ to enter. The outlet carries the corresponding work away. Combining internal energy u and flow work defines enthalpy:

$$h = u + pv,$$

$$e_{\text{stream}} = h + V^2/2 + gz.$$

Here V denotes speed, v specific volume and z elevation. In steady operation, sum these stream contributions and include heat into and shaft work out of the control volume. For one inlet and one outlet with equal mass rate, rearrangement gives

$$\dot{Q} - \dot{W}_s = \dot{m}\,[h_2 - h_1 + (V_2^2 - V_1^2)/2 + g(z_2 - z_1)].$$

All terms have units of power. A nozzle can accelerate flow without shaft work by lowering static enthalpy. A throttle with no heat, no shaft work and negligible speed and elevation changes is instead approximately isenthalpic. Omitting the pressure work while also using internal energy would miss both conclusions.

Background · Specific enthalpy

$$h = u + pv$$

Specific quantities are per unit mass. Enthalpy accounts for internal energy and flow work when mass passes through a control surface.

DERIVATION AND APPLICATIONS

Convert a nozzle energy balance into a mass-flow limit

A steady nozzle converts static enthalpy into bulk kinetic energy. A converging passage has a maximum possible mass flow at a specified inlet stagnation state. Derive this for a steady, one-dimensional, adiabatic ideal gas with constant c_p and γ, reversible acceleration and negligible inlet-reservoir velocity. Write reservoir conditions as T_0, p_0 and the local static conditions as T, p. With no shaft work the energy balance gives

$$c_pT_0 = c_pT + u^2/2,$$

$$u = \sqrt{2c_pT_0(1 - T/T_0)}.$$

Isentropic acceleration links temperature and pressure. Let $r = T/T_0$, so $p/p_0 = r^{\gamma/(\gamma-1)}$. The mass flux is ρu and ideal-gas density is $p/(RT)$. Substitution gives a function of just one local variable:

$$G(r) = \frac{p_0}{RT_0}\sqrt{2c_pT_0}\, r^{1/(\gamma-1)}\sqrt{1-r}.$$

To find its maximum differentiate $\ln G = (\ln r)/(\gamma-1) + \ln(1-r)/2 + \text{constant}$. Equating the two derivative terms gives the sonic temperature and pressure ratios:

$$r_* = \frac{2}{\gamma+1},$$

$$p_*/p_0 = \left(\frac{2}{\gamma+1}\right)^{\gamma/(\gamma-1)},$$

$$u_*^2 = \gamma R T_*.$$

The flow through a converging exit is choked when downstream pressure is low enough for this state to be reached. Further lowering downstream pressure changes the flow after the exit; the ideal mass flux at that exit remains fixed. The model omits wall friction, heat transfer and discharge-coefficient corrections.

WORKED EXAMPLE

Size a sonic air metering aperture

Change the inputs and test this calculation.

For $p_0 = 200$ kPa absolute, $T_0 = 300$ K, $\gamma = 1.40$ and $R = 287$ J/(kg K), the critical state has $T_* = 250$ K and $p_* = 105.7$ kPa. Then $\rho_* = 1.47$ kg/m³ and $u_* = 317$ m/s, giving $G_* \simeq 467$ kg/(m² s). A 20.0 mm² effective aperture meters about 9.33 g/s if the downstream pressure allows choking.

SECOND APPLICATION

Decide whether a higher receiver pressure matters

Change the inputs and test this calculation.

Raise the receiver pressure to 160 kPa. Its ratio to reservoir pressure is now 0.800, above the critical 0.528. An ideal converging nozzle discharges at this pressure with $r = 0.800^{(0.4/1.4)} = 0.938$. The energy equation gives $u \approx 193$ m/s, static density about 1.98 kg/m³, and mass flux about 382 kg/(m² s). Through the same aperture the predicted flow falls to 7.64 g/s. A model assuming choked flow at both settings would miss this change.

Try it · Nozzle

Set the inputs to the values in the worked example, then vary one quantity while holding the others fixed. Compare the curve and numerical balance with the second application. The calculation uses the assumptions derived above; its readout identifies bounds such as choking, phase-region changes and inconsistent measurements.

Open this live calculation

Try the calculation

Stagnation pressure (kPa)	200
Stagnation temperature (K)	300
Receiver / stagnation pressure	0.8

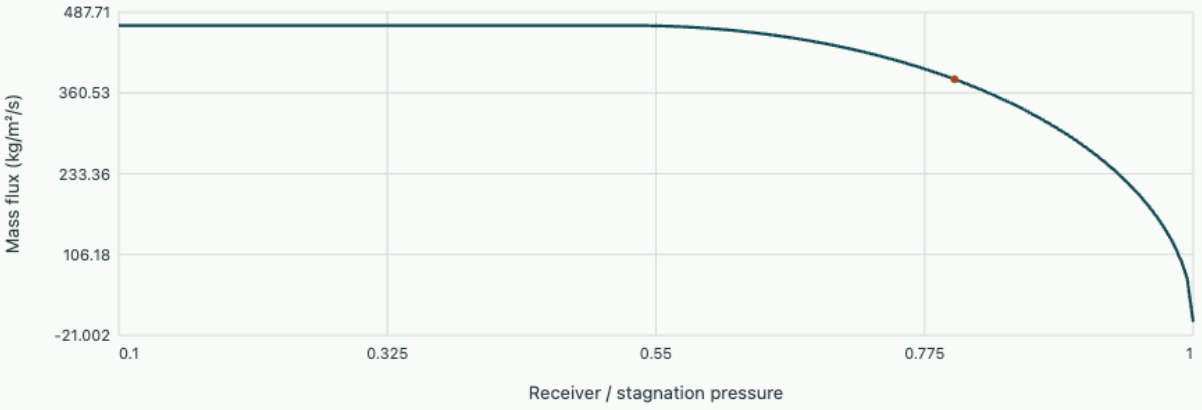

Converging ideal nozzle | Selected receiver

Critical pressure ratio 0.528282; choked mass flux 466.712 kg/(m² s). A 20 mm² aperture gives 7.64291 g/s at the selected receiver pressure. The converging exit is unchoked.

Account for fluid entering and leaving

A closed-system piston can be followed with its gas. A continuously operating turbine has fresh fluid crossing its boundary every second. Enclose the hardware and its internal flow in a stationary control volume. Flowing mass carries internal energy u, kinetic energy $v^2/2$ and gravitational energy gz per kilogram. Upstream pressure also does work pushing fluid in, while downstream pressure opposes its exit. Per unit mass this flow work is p/ρ, motivating the specific enthalpy $h = u + p/\rho$.

Adding all incoming and outgoing energy rates gives, for one inlet, one outlet and constant stored energy,

$$\dot{Q} - \dot{W}_s = \dot{m}\left[h_2 - h_1 + (v_2^2 - v_1^2)/2 + g(z_2 - z_1)\right].$$

Here $\dot{W}_s$ is shaft power out of the control volume. Pressure work at the flowing boundaries is already contained in enthalpy; adding it once more would count it twice. One inlet and outlet have the same steady mass rate even when their volumetric flow rates differ dramatically because gas expands.

Read the dominant energy conversion

In a short, insulated nozzle there is no shaft, the elevation difference is small and falling enthalpy raises velocity. The property $h + v^2/2$ is then constant, often called stagnation enthalpy. In an insulated turbine with small inlet and outlet speeds, the enthalpy drop instead appears mainly as positive shaft work. An insulated compressor takes work in, raising enthalpy. An insulated throttling valve has neither shaft power nor a large net change of bulk kinetic energy between distant ports, so its inlet and outlet enthalpies are approximately equal. These components can all drop pressure during expansion, yet their energy effects differ.

For ideal-gas air with nearly constant c_p, an enthalpy difference may be approximated by $c_p \Delta T$. Keeping h in kJ/kg while using $v^2/2$ in J/kg is a common factor-of-one-thousand error. At 100 m/s the kinetic energy is 5 kJ/kg. That estimate often decides whether an inlet velocity may be omitted from a turbine work calculation.

Evaluate heat loss and exhaust speed

A turbine passes 2.0 kg/s of gas from $h_1 = 820$ to $h_2 = 600$ kJ/kg. Inlet speed is 40 m/s, exhaust speed 160 m/s and the casing loses 10 kW of heat. Elevations are equal. The gas loses $2(820 - 600) = 440$ kW of enthalpy flow. Its kinetic-energy rate grows by $2(160^2 - 40^2)/2000 = 24.0$ kW. The unit delivers $440 - 24 - 10 = 406$ kW shaft power. Counting all 440 kW as shaft output would overstate the result by both the residual exhaust energy and the heat leak.

A separate nozzle using gas at $h_1 = 820$ kJ/kg and 40 m/s could reach $h_2 = 800$ kJ/kg without a shaft if its speed rose to $\sqrt{40^2 + 2(20\,000)} = 204$ m/s. This estimate is an energy condition. The required area ratio, pressure drop and possibility of choking also depend on compressible-flow relations.

Extend the same balance to a mixer

Two insulated water lines mix at negligible speed change. One supplies 0.20 kg/s at 360 K and the other 0.30 kg/s at 300 K. A mass balance gives 0.50 kg/s at exit. With equal, constant c_p and no shaft, summing the two inlet enthalpy rates yields an outlet temperature $[0.20(360) + 0.30(300)]/0.50 = 324$ K. The temperature lies between the inlet values, closer to the higher mass rate. The mixer conserves energy while creating entropy. The next lecture explains how both statements can hold simultaneously and why reversing the mixer without another effect is impossible.

Interpret a measured stagnation temperature

For an insulated diffuser with negligible shaft work, reducing speed from 200 to 40 m/s raises static enthalpy by $(200^2 - 40^2)/2 = 19.2\,\mathrm{kJ/kg}$. For air with $c_p = 1.005\,\mathrm{kJ/(kg\,K)}$ that corresponds to a 19.1 K static temperature rise. Stagnation enthalpy $h + V^2/2$ stays constant even if internal friction produces entropy. A

stationary temperature probe that slows the arriving air may approach a stagnation reading, so comparison with a predicted static temperature requires knowing where and how the measurement was taken.

Minimise compression work by selecting an intermediate pressure

A continuous air supply can be compressed in two stages and cooled at intermediate pressure. For a steady adiabatic compressor with negligible kinetic and potential changes, shaft work input is the stream enthalpy rise. With an ideal gas, constant c_p, γ, and a reversible stage,

$$w = c_p T_{\mathrm{in}} [(p_{\mathrm{out}}/p_{\mathrm{in}})^{a} - 1],$$

where $a = (\gamma - 1)/\gamma$. This follows by integrating $ds = c_p\,dT/T - R\,dp/p = 0$, then applying $dh = c_p\,dT$. If both stages use the same isentropic efficiency, divide each ideal rise by that efficiency.

Assume complete intercooling to the original inlet temperature T_1, no pressure losses and prescribed first and last pressures. Put $r_1 = p_i/p_1,\ r_2 = p_3/p_i$, so $r_1 r_2 = p_3/p_1 = r$. Total ideal work is $c_p T_1 [r_1^{a} + (r/r_1)^{a} - 2]$. Differentiate with respect to $\ln r_1$; the derivative is $c_p T_1 a (r_1^{a} - r_2^{a})$. The sum has a minimum at equal ratios, giving intermediate pressure $p_i = \sqrt{p_1 p_3}$. This is a geometric mean of absolute pressures. A gauge-pressure ratio would give a completely different design.

Compare two cooled stages with a single compression

Supply air at 100 kPa and 300 K and deliver it at 900 kPa absolute. For $\gamma = 1.4,\ c_p = 1.005\,\mathrm{kJ/(kg\,K)}$, an ideal single stage gives $T_3 = 300(9^{0.2857}) = 562.0\,\mathrm{K}$ and work about 263 kJ/kg. Two equal-ratio stages operate through 300 kPa. Each reaches $300(3^{0.2857}) = 410.6\,\mathrm{K}$ before its outlet or intercooler, requiring $1.005(110.6) = 111.2\,\mathrm{kJ/kg}$. Total work is about 222.4 kJ/kg, roughly 15.5% lower. The intercooler must remove 111.2 kJ/kg between the stages under the ideal complete-cooling assumption.

With equal stage efficiencies of 0.80, each real enthalpy rise becomes 139.0 kJ/kg and the intermediate discharge temperature is about 438 K. Cooling back to 300 K then removes 139.0 kJ per kilogram, with a corresponding load of 27.8 kW at 0.20 kg/s. The equal-ratio choice survives because both stage work expressions acquired the same multiplier. An intercooler that can cool only to 330 K increases the weight of second-stage compression work; the new optimum shifts towards a larger first-stage pressure ratio. A design comparing only compressor shaft power would also need to account for coolant pumping, fan power and pressure losses before choosing the number of stages.

For an incomplete intercooler the optimal shift can be calculated. Take stage two inlet temperature $T_m = 330$ K independently maintained and identical efficiencies. The variable portion of total work is proportional to $T_1 r_1^{a} + T_m (r/r_1)^{a}$. Differentiating with respect to log first-stage ratio gives $T_1 r_1^{a} = T_m r_2^{a}$, or $r_1/r_2 = (330/300)^{1/a} = 1.40$. At overall ratio nine,

$$r_1 = \sqrt{9(1.40)} = 3.54,\ r_2 = 2.54.$$

Compressing farther while the air is colder is energetically preferable within this simplified comparison. These ratios presume the cooler can maintain 330 K over its revised duty and that both discharge temperatures stay

within their allowable ranges. A compressor map may impose further flow and speed limits beyond the thermodynamic minimum.

At fixed pressure endpoints, the measured intermediate pressure alone would thus be insufficient for diagnosing whether the cooler or a compressor caused an observed power increase.

How much static temperature does a nozzle convert to speed?

Air flows steadily through an insulated nozzle with negligible shaft work and elevation change. Take $c_p = 1005$ J/(kg K). Divide the conserved stagnation enthalpy $c_pT + u^2/2$ by c_p to find the exit temperature.

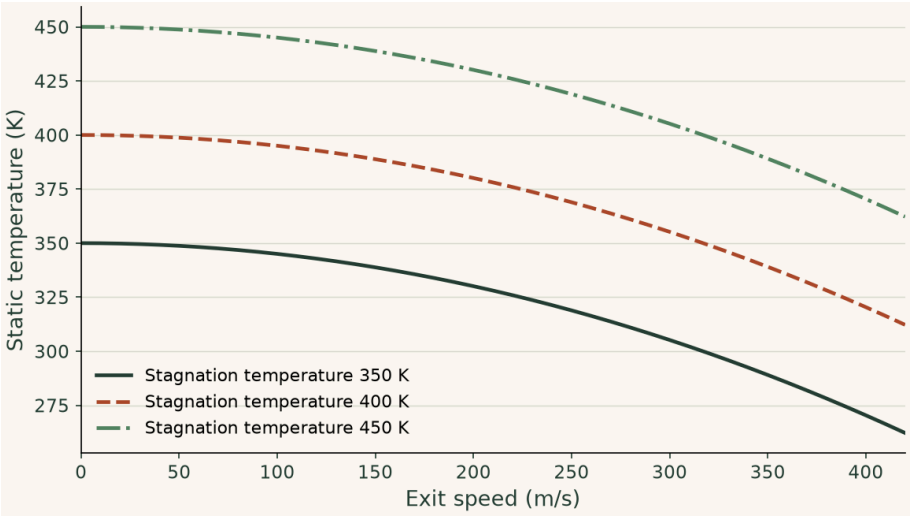

At 300 m/s the static temperature is about 44.8 K below stagnation temperature, regardless of pressure in this enthalpy calculation. Predicting exit pressure as well would require an entropy or loss model.

▶ **Read values from the plot**

Interactive calculation · change the inputs of this example

Change stagnation temperature and heat capacity independently. The speed-related temperature drop follows the energy balance.

The curves below are recalculated from this section's equations in your browser. The fixed reference plot above remains available for comparison.

Open this live calculation

Try the calculation

Stagnation temperature (K) **400**

Specific heat (J/kg K) **1005**

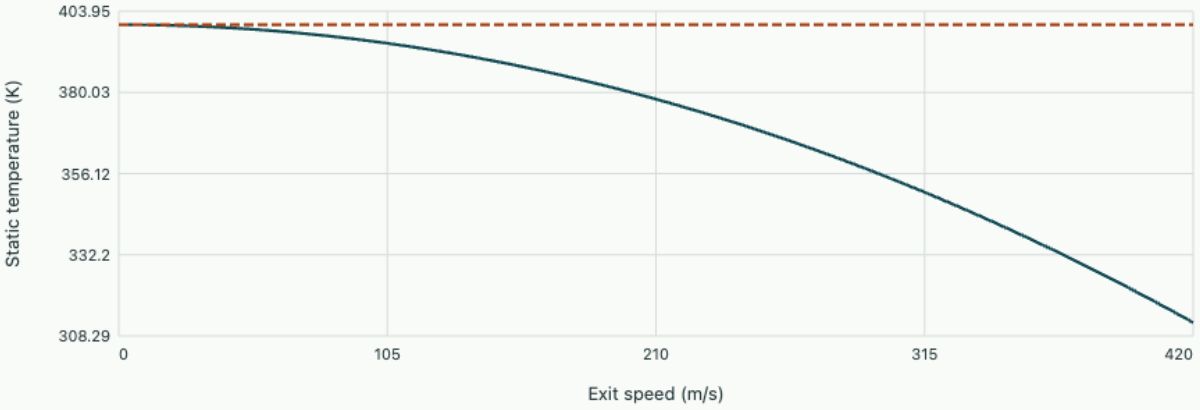

Selected stagnation state Inlet stagnation temperature

At 300 m/s the static temperature drop is 44.776 K. Pressure needs a separate entropy/loss relation.

Entropy, irreversibility and useful-work loss

Engineering thermodynamics and heat transfer

Check a proposed process with the second law and quantify the effect of a finite temperature difference.
FROM FIRST PRINCIPLES

Separate entropy transfer from entropy generated inside a boundary

For a reversible differential process the thermodynamic definition of entropy is

$$dS = \delta Q_{\mathrm{rev}}/T.$$

Irreversible changes reach the same state-function difference but can generate entropy internally. For a closed system exchanging heats Q_j across boundary portions at known local temperatures T_j, the second law gives an experimentally testable inequality:

$$S_2 - S_1 = \sum_j \int \frac{\delta Q_j}{T_j} + S_{\mathrm{gen}}, \qquad S_{\mathrm{gen}} \geq 0.$$

To evaluate the left side independently of the actual path, use a reversible path between the same equilibrium states. For an ideal gas the fundamental relation $Tds = du + p\,dv$, together with $du = c_v dT$ and $pv = RT$, yields

$$s_2 - s_1 = c_v \ln(T_2/T_1) + R \ln(v_2/v_1)$$

for constant specific heat. Inserting this property change into the balance isolates entropy generation of the actual process. A gas might cool and lose entropy while the combined gas plus surroundings generates entropy. For heat exchanged across a wall separating two reservoirs, the two reservoirs' boundary transfer terms have opposite signs and different temperature denominators.

Background · Entropy generation

$$\dot{S}_{\mathrm{gen}} = \dot{S}_{\mathrm{system}} - \sum \dot{Q}_j/T_j - \dot{S}_{\mathrm{mass,in}} + \dot{S}_{\mathrm{mass,out}}$$

Use boundary temperatures at which heat actually crosses. All flows, storage and units must refer to the same control volume.

DERIVATION AND APPLICATIONS

Count the entropy generated when two ideal gases mix

Two separated ideal gases can have the same temperature and pressure, yet their mixing increases entropy. Consider distinct nonreacting species in insulated rigid compartments. Remove the separator and let them equilibrate in the common volume. Initially both sides have temperature T and pressure p, and no shaft work is performed. The total internal energy of ideal gases depends only on temperature, so the final equilibrium temperature stays T. Denote mole numbers by n_i and total volume by V.

For each species, the entropy change can be calculated along a reversible isothermal expansion into its accessible final volume, even though the actual mixing is irreversible. Integrating $\delta Q_{\mathrm{rev}}/T = p_i\,dV/T = n_i R\,dV/V$ gives

$$\Delta S_i = n_i R \ln(V/V_i).$$

The initial volumes were proportional to mole numbers because p and T were common. Thus $V_i/V = n_i/\sum_j n_j = y_i$, the final mole fraction. Add the species changes and use zero boundary entropy transfer:

$$\Delta S_{\mathrm{mix}} = S_{\mathrm{gen}} = -R\sum_i n_i \ln y_i.$$

For two gases with y and $1-y$, the generation per mole of mixture is maximal when both are equally abundant. The derivative and its zero are

$$\frac{d}{dy}\frac{\Delta S_{\mathrm{mix}}}{n} = R\ln[(1-y)/y] = 0 \quad\Longrightarrow\quad y = \tfrac{1}{2}.$$

The entropy change tends to zero when either fraction tends to one. If the compartments contain the same gas and the initial temperatures and pressures already agree, removing the partition does not create distinguishable species; assigning artificial labels would predict a spurious mixing term.

WORKED EXAMPLE

Estimate the lost work of mixing two samples

Change the inputs and test this calculation.

Mix 1.00 mol of one ideal gas and 3.00 mol of another at 300 K. The mole fractions are 0.250 and 0.750. The generated entropy is

$$-8.314[\ln 0.25 + 3\ln 0.75] = 18.7\,\mathrm{J/K}.$$

Relative to an environment at 300 K, the destroyed work potential is $300(18.7) = 5.61$ kJ. No energy has disappeared from the insulated vessel.

SECOND APPLICATION

Compare equal and strongly unequal amounts

Change the inputs and test this calculation.

For a fixed total of 4.00 mol, an equal mixture generates $4R\ln 2 = 23.1$ J/K. With only 0.0400 mol of one species the fractions become 0.0100 and 0.990. The generation falls to

$$4R[-0.01\ln 0.01 - 0.99\ln 0.99] = 1.86\,\mathrm{J/K}.$$

Each rare molecule gains access to much more volume, but there are few such molecules; both contributions must be counted before comparing samples.

Try it · Mixing

Set the inputs to the values in the worked example, then vary one quantity while holding the others fixed. Compare the curve and numerical balance with the second application. The calculation uses the assumptions derived above; its readout identifies bounds such as choking, phase-region changes and inconsistent measurements.

Open this live calculation

Try the calculation

Total amount (mol)	4
First-species mole fraction	0.25
Environment temperature (K)	300

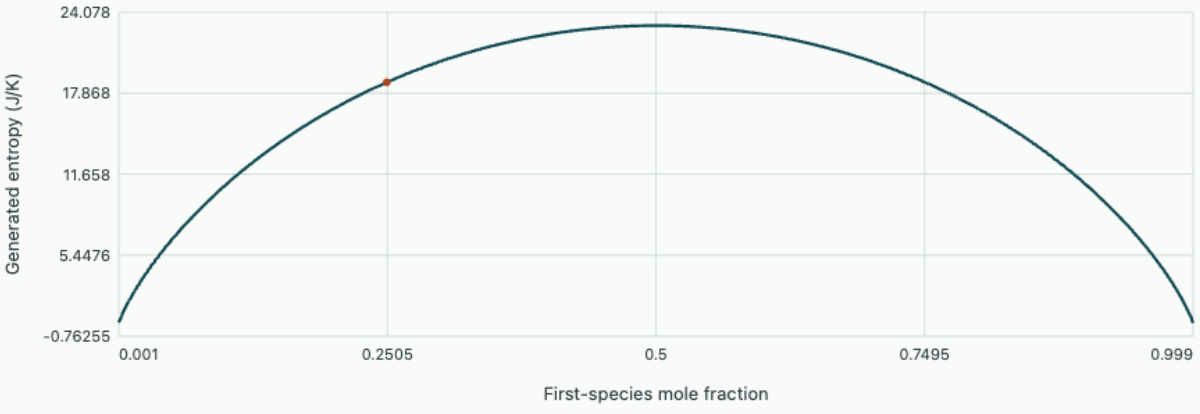

Two distinct ideal species Selected mixture

Entropy generation 18.7021 J/K; destroyed work potential T0 Sgen = 5.61062 kJ. Equal initial temperature and pressure; the same-species case has no mixing term.

Energy bookkeeping is not sufficient

In an insulated two-body experiment, heat lost by one body is gained by the other. That energy balance alone would allow heat to move spontaneously from a colder body to a hotter one. Entropy supplies the missing direction constraint. It is a property of a state. For a simple equilibrium system a reversible differential heat transfer defines $dS = \delta Q_{\mathrm{rev}}/T$. An irreversible process between the same endpoints has the same entropy difference; the actual heat transfer divided by the body temperature need not equal that difference.

For a closed system, the integrated entropy balance reads $\Delta S = \int \delta Q/T_b + S_{\mathrm{gen}}$, where T_b is the absolute temperature at the boundary where the heat crosses. Entropy generation satisfies $S_{\mathrm{gen}} \geq 0$. A system can lose entropy by exporting it in heat or in a mass stream. A negative change in the entropy of a single hot body therefore need not violate the second law.

Background · Why does a small temperature difference produce a second-order loss at fixed conductance?

Let a link of constant conductance G join reservoirs held at $T + \Delta T/2$ and $T - \Delta T/2$. In steady operation the heat current is $G\Delta T$. Subtract the entropy leaving the hot reservoir from that entering the cold reservoir.

$$\dot{S}_{\mathrm{gen}} = G\Delta T\left(\frac{1}{T - \Delta T/2} - \frac{1}{T + \Delta T/2}\right) = \frac{G(\Delta T)^2}{T^2 - (\Delta T/2)^2}.$$

At fixed conductance, halving a small difference approximately quarters the entropy production per second, while it also halves the heat current. If a design must keep the heat rate fixed, halving the difference requires nearly twice the conductance and only halves the production. This distinction matters when comparing a bigger exchanger with the original one.

Calculate a property change with a convenient path

For a calorically perfect ideal gas, combine $T\,ds = du + p\,dv$ with $du = c_v\,dT$ and $p = RT/v$. This gives

$$ds = c_v\,dT/T + R\,dv/v.$$

Replacing v using the ideal-gas law yields the alternative $ds = c_p\,dT/T - R\,dp/p$. Integrate between the initial and final equilibrium states. These are property relationships, even when the actual path between states includes friction, throttling or free expansion.

An ideal gas passing through an insulated valve with equal inlet and outlet speeds has $h_2 = h_1$, so its temperatures are equal in the constant-c_p approximation. Dropping air from 400 to 100 kPa gives $s_2 - s_1 = -0.287\ln(100/400) = 0.398$ kJ/(kg K). With no heat transfer or other streams, the positive entropy change per kilogram is entirely generated. A reversible insulated expansion turbine would generally produce a lower outlet temperature and shaft work.

Trace entropy across two reservoirs

An insulated conductor transfers 600 W from a reservoir at 400 K to another at 300 K. Its stored energy and entropy are steady. The hot reservoir loses entropy at $600/400 = 1.50$ W/K, while the cold one gains $600/300 = 2.00$ W/K. The total generation rate is 0.50 W/K. Heat rate is the same at both ends; entropy transfer in that heat differs because the temperatures differ.

$$\dot{S}_{\mathrm{gen}} = \dot{Q}\left(\frac{1}{T_c} - \frac{1}{T_h}\right).$$

If the surroundings defining the available-work reference are at $T_0 = 300$ K, the corresponding destruction of available work is $T_0 \dot{S}_{\mathrm{gen}} = 150$ W. For these particular reservoirs a reversible engine supplied with 600 W at 400 K could have produced $600(1 - 300/400) = 150$ W before rejecting the remainder. The conductor instead rejects all 600 W to the cold sink. Energy did not disappear; an opportunity for work was lost.

Locate sources and choose a boundary

A heat exchanger can have little external heat leak yet generate appreciable entropy internally through finite temperature differences and pressure losses. Write $\dot{S}_{\mathrm{gen}} = \sum_{\mathrm{out}} \dot{m} s - \sum_{\mathrm{in}} \dot{m} s - \sum_j \dot{Q}_j / T_{b,j}$ for steady operation. Both streams belong inside the same control volume when testing internal heat exchange. If the boundary encloses only one stream, heat crossing the wall carries entropy with it and must appear explicitly. The largest reduction in entropy generation per extra unit of area or pumping effort is a design question. The thermodynamic balance identifies losses without alone determining the economically preferred geometry.

Compare the conditions on a valve and a turbine

Air passes through a steady insulated throttling valve from 400 to 100 kPa with negligible changes of speed. With no shaft work the first law gives equal inlet and outlet enthalpy. For ideal air this implies equal temperature. Its entropy change follows from $\Delta s = c_p \ln(T_2/T_1) - R \ln(p_2/p_1) = 0.287 \ln 4 = 0.398\,\mathrm{kJ/(kg\,K)}$. At 0.10 kg/s the entropy generation is 0.0398 kW/K. An ideal adiabatic turbine undergoing the same pressure drop has constant entropy and delivers work as its temperature falls. Pressure ratio alone therefore cannot determine the temperature change across an unfamiliar device.

Obtain a work bound for two finite bodies

The reversible efficiency $1 - T_c/T_h$ describes operation between reservoirs held at fixed temperatures. If the two bodies are finite, each increment of transferred heat changes the temperatures and the instantaneous efficiency. Consider two incompressible blocks in a common insulated enclosure. Their heat capacities C_h, C_c are taken as constant. A reversible engine operates between them and exports mechanical work until no temperature difference remains. The working fluid returns to its starting state after each incremental cycle, so conservation of entropy for this reversible composite system determines the final temperature.

$$C_h \ln(T_f/T_h) + C_c \ln(T_f/T_c) = 0,$$
$$T_f = T_h^{C_h/(C_h+C_c)} T_c^{C_c/(C_h+C_c)}.$$

Energy conservation gives maximum work $W_{\max} = C_h T_h + C_c T_c - (C_h + C_c) T_f$ for consistent internal-energy references. The reversible final temperature is a weighted geometric mean. If the two blocks are simply put in thermal contact without work output, first-law conservation instead gives the weighted arithmetic mean. The arithmetic mean is larger for unequal initial temperatures; additional energy remains as sensible heat after irreversible contact. These different final temperatures matter when comparing two ways of using the same finite resources.

In a real engine running at finite heat-transfer rates there are temperature differences between each block and the working fluid. Total entropy then increases. For a specified common final temperature the entropy change of the blocks is fixed by state properties; measuring it indicates the departure from the ideal reversible construction without needing the working-fluid cycle.

Compare direct contact and a reversible engine quantitatively

Take a hot assembly of heat capacity 2.0 kJ/K at 450 K and a cooler assembly of 1.0 kJ/K at 300 K. Solving the log balance gives

$$T_f = (450^2 \, 300)^{1/3} = 393.1 \,\mathrm{K}.$$

The hot assembly releases about $2(450 - 393.1) = 113.8 \,\mathrm{kJ}$ and the cold one receives 93.1 kJ. Their difference gives approximately 20.7 kJ of maximum work. Acting as though the initial temperatures stayed at 450 and 300 K throughout would assign an incorrectly high efficiency to later parts of the transfer.

By direct contact the common temperature would be $(2 \times 450 + 300)/3 = 400 \,\mathrm{K}$. The entropy generated is

$$2 \ln(400/450) + \ln(400/300) = 0.0521 \,\mathrm{kJ/K}.$$

If instead a device delivers a measured 12 kJ of work before the blocks reach a common temperature, energy fixes that temperature as $(1200 - 12)/3 = 396 \,\mathrm{K}$. The measured process then generates $2 \ln(396/450) + \ln(396/300) = 0.0220 \,\mathrm{kJ/K}$, a positive value between the two limiting cases. Entropy generation multiplied by an environmental temperature gives destroyed exergy relative to that environment; comparing the finite-body work outputs also requires acknowledging their different final stored energies.

The geometric-mean result has an informative limiting case. If the cold assembly is so large that C_c/C_h tends to infinity, its temperature remains effectively fixed, and the engine work reduces to $C_h[(T_h - T_c) - T_c \ln(T_h/T_c)]$. This is the integral of the reversible work recovered as the finite hot block approaches a constant sink. It is smaller than applying the initial Carnot efficiency to all of the hot block’s released heat.

Which mixture generates the most entropy?

Mix streams of the same incompressible liquid at 370 K and 290 K, at equal pressure, with constant $c_p = 4180$ J/(kg K). The energy balance fixes the outlet temperature. Evaluate the outlet entropy minus the

mass-weighted inlet entropy using $c_p \ln T$.

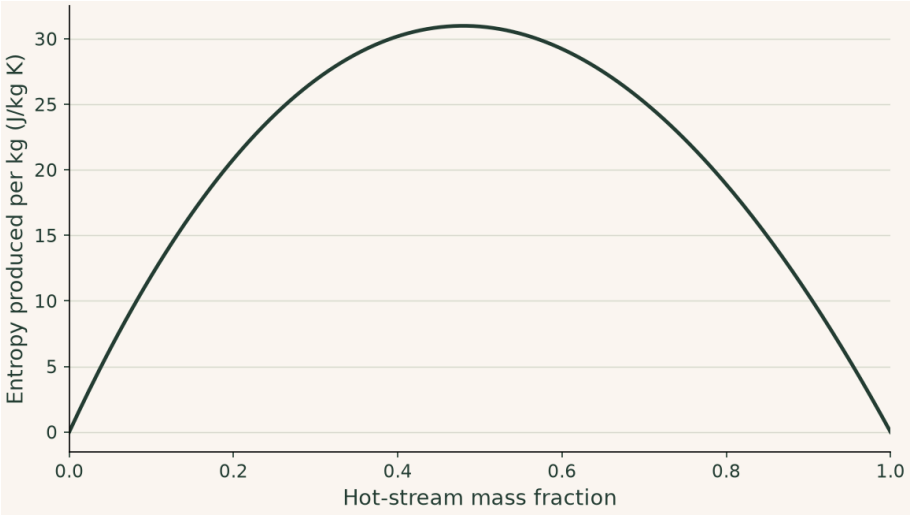

Either pure stream gives zero entropy production. An equal-mass mixture settles at 330 K and produces about 30.9 J/(kg K). The maximum is near, but not exactly at, equal fractions because temperature appears inside a logarithm.

▶ **Read values from the plot**

Interactive calculation · change the inputs of this example

Move the two inlet temperatures toward each other and watch entropy production vanish throughout the mixture range.

The curves below are recalculated from this section's equations in your browser. The fixed reference plot above remains available for comparison.

Open this live calculation

Try the calculation

Hot stream (K)	**370**
Cold stream (K)	**290**
Liquid heat capacity (J/kg K)	**4180**

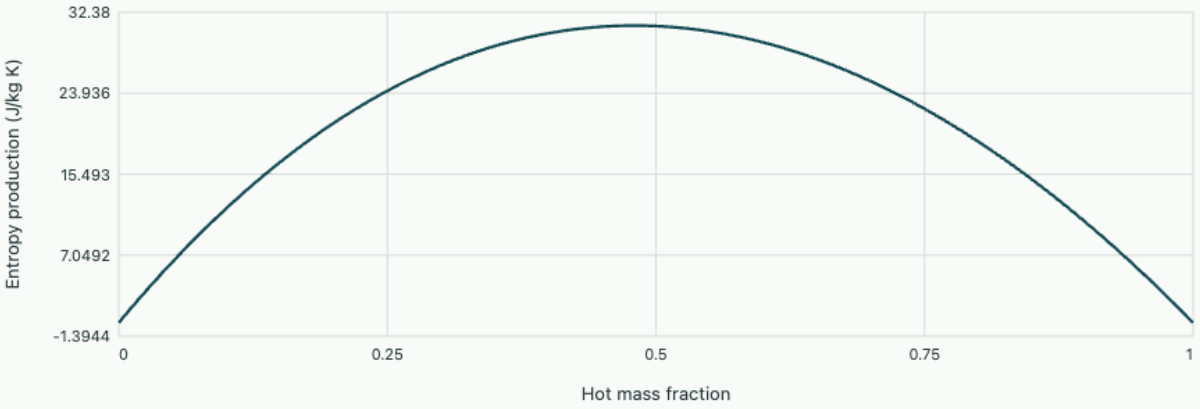

Adiabatic mixture

Equal mass fractions reach 330.000 K and generate 30.935 J/(kg K). Endpoints are unmixed and generate zero.

Phase equilibrium and steam quality

Engineering thermodynamics and heat transfer

Locate a two-phase state from two independent measurements and use mass fractions to find its volume and energy.

FROM FIRST PRINCIPLES

Derive the slope of a pure-substance coexistence line

At liquid–vapour equilibrium, transferring a small amount of matter between phases cannot lower the total Gibbs energy. For a pure substance this requires equality of the specific Gibbs energies $g_f = g_g$ at common temperature and pressure. To find how pressure must change when temperature changes while maintaining coexistence, differentiate this equality along the saturation curve. The fundamental differential $dg = -s\,dT + v\,dp$ gives

$$-s_f\,dT + v_f\,dp = -s_g\,dT + v_g\,dp.$$

Rearranging expresses the measured slope as the entropy change of vaporization divided by its specific-volume change. Reversible phase change at constant temperature absorbs specific enthalpy $h_{fg} = T(s_g - s_f)$, so

$$\left(\frac{dp}{dT}\right)_{\text{sat}} = \frac{s_g - s_f}{v_g - v_f} = \frac{h_{fg}}{T(v_g - v_f)}.$$

The units are Pa/K. No ideal-gas approximation has been made. If the vapour is dilute and the liquid volume negligible, further substitution can simplify the slope, but this fails near the critical point. The slope describes neighbouring coexistence states; determining a system's vapour mass fraction additionally requires an extensive property such as specific enthalpy or volume.

Background · Vapour quality

$$x = m_v/(m_v + m_l)$$

This mass fraction applies in a two-phase equilibrium region. Mole fraction, mass fraction and fraction of volume occupied by vapour can be very different.

DERIVATION AND APPLICATIONS

Trace an isentropic pressure change through the saturation dome

A vapour fraction inferred from one property must reproduce the others at the same pressure. Consider a pure substance in local liquid–vapour equilibrium. Let x be the vapour mass fraction and use subscripts f, g for saturated liquid and vapour at the current pressure. Neglect interfacial contributions and assume each

phase has the tabulated saturation state. Specific extensive properties then obey the lever rule because their totals add over the masses of both phases:

$$s = (1-x)s_f + x s_g = s_f + x s_{fg},$$
$$h = h_f + x h_{fg}.$$

In a steady adiabatic reversible expansion the entropy balance gives

$$s_2 = s_1.$$

At the specified exit pressure solve the entropy relation before using the enthalpy relation:

$$x_{2s} = (s_1 - s_{f,2})/s_{fg,2},$$
$$h_{2s} = h_{f,2} + x_{2s} h_{fg,2}.$$

This interpolation is physical only when $0 \le x_{2s} \le 1$. For an actual adiabatic turbine with negligible kinetic-energy change, use isentropic efficiency $\eta_t = (h_1 - h_2)/(h_1 - h_{2s})$ and check the actual phase state separately:

$$h_2 = h_1 - \eta_t (h_1 - h_{2s}),$$
$$x_2 = (h_2 - h_{f,2})/h_{fg,2}$$

if the actual outlet is still a mixture. Irreversibility reduces extracted work, which can leave the exhaust drier at the same pressure. A drier exit therefore does not by itself demonstrate higher efficiency.

WORKED EXAMPLE

Infer equilibrium quality for two expansion histories

Change the inputs and test this calculation.

At a low-pressure outlet take rounded steam properties $h_f = 192$, $h_{fg} = 2393$ kJ/kg, $s_f = 0.649$ and $s_{fg} = 7.50$ kJ/(kg K). Steam enters with $h_1 = 3200$ kJ/kg and $s_1 = 6.65$ kJ/(kg K). The reversible quality is $(6.65 - 0.649)/7.50 = 0.800$, giving $h_{2s} = 2107$ kJ/kg. For $\eta_t = 0.800$, the actual exit enthalpy is about 2326 kJ/kg and the quality 0.892. The actual specific work is about 874 kJ/kg.

SECOND APPLICATION

Detect when a reported quality has no physical meaning

Change the inputs and test this calculation.

At the same outlet pressure an independent calculation gives $h_2 = 2630$ kJ/kg. Interpolating would give

$$x = (2630 - 192)/2393 = 1.019.$$

Since the saturated vapour enthalpy is only 2585 kJ/kg, the enthalpy lies above the saturation range. The outlet is superheated; finding its temperature requires superheated-vapour properties. Reporting a vapour

fraction above unity would hide a change of thermodynamic region.

Try it · Turbine quality

Set the inputs to the values in the worked example, then vary one quantity while holding the others fixed. Compare the curve and numerical balance with the second application. The calculation uses the assumptions derived above; its readout identifies bounds such as choking, phase-region changes and inconsistent measurements.

Open this live calculation

Try the calculation

Turbine isentropic efficiency **0.8**

Inlet entropy (kJ/kg/K) **6.65**

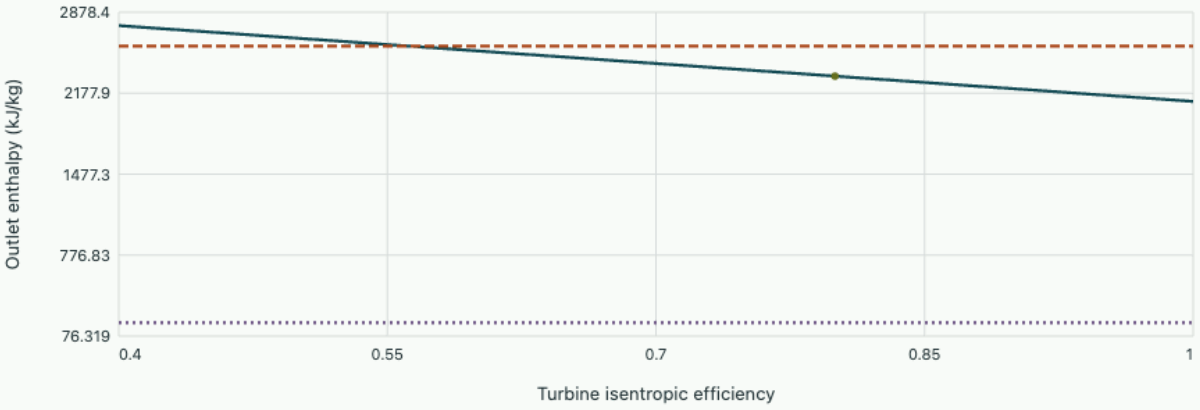

Actual outlet enthalpy · Saturated-vapour bound · Saturated-liquid bound · Selected efficiency

Isentropic outlet quality 0.800133; selected work 874.625 kJ/kg. Actual mixture quality 0.891507. Rounded outlet saturation properties are fixed as in the derivation.

Temperature can stay fixed while heat enters

In an open pan at a fixed ambient pressure, liquid water begins to boil at a particular temperature. Additional heat can then convert liquid to vapour while the temperature remains close to that value. Specifying both saturation temperature and saturation pressure of a pure substance does not locate a unique mixture state. Along the coexistence curve they are linked, and the fraction of material in each phase supplies the missing information. Outside that curve pressure and temperature can usually be used as independent coordinates for a simple compressible substance.

Define quality $x = m_g/(m_f + m_g)$ as the vapour mass fraction, using subscripts f for saturated liquid and g for saturated vapour. A property such as volume or enthalpy is additive over the two masses when interfacial

contributions are negligible. Writing $V = m_f v_f + m_g v_g$ and dividing by total mass gives the mixing relation. The same derivation applies to specific internal energy, enthalpy and entropy.

$$v = v_f + x(v_g - v_f),$$

$$h = h_f + x h_{fg},$$

$$h_{fg} = h_g - h_f.$$

Quality lies between zero and one in a saturated mixture. A calculated value beyond either endpoint signals a single-phase state or inconsistent measurements. Temperature does not interpolate: both phases in equilibrium have the same temperature.

Distinguish vapour mass from space occupied

A closed vessel contains 2.00 kg of saturated water and steam at approximately 100 kPa in 0.500 m³. Use the supplied rounded values $v_f = 0.00104,\ v_g = 1.694\,\mathrm{m^3/kg}$,

$$h_f = 417.5,\ h_{fg} = 2258\,\mathrm{kJ/kg}.$$

The measured specific volume is 0.250 m³/kg. Quality follows as $(0.250 - 0.00104)/(1.694 - 0.00104) = 0.1471$. Thus vapour mass is 0.294 kg while 1.706 kg remains liquid.

The vapour occupies $0.2941(1.694) = 0.4982\,\mathrm{m^3}$, almost the entire vessel. The liquid occupies roughly 0.0018 m³. Specific enthalpy is

$$417.5 + 0.1471(2258) = 749.6\,\mathrm{kJ/kg}.$$

Specific internal energy follows from $u = h - pv = 749.6 - 100(0.250) = 724.6\,\mathrm{kJ/kg}$ when using kPa and m³/kg. A level indicator measures occupied volume; using its reading directly as a mass fraction would give a large error.

Connect latent heat to the coexistence curve

For two pure phases in equilibrium their specific Gibbs free energies are equal. Along the coexistence line this equality persists, so $dg_f = dg_g$. From $dg = v\,dp - s\,dT$ one obtains

$$dp_{\mathrm{sat}}/dT = (s_g - s_f)/(v_g - v_f) = h_{fg}/[T(v_g - v_f)].$$

At about 373 K using the same values, the slope is

$$2258/[373(1.693)] = 3.58\,\mathrm{kPa/K}.$$

A small increase of boiling pressure by 3.6 kPa therefore raises saturation temperature by about one kelvin locally. Larger changes require integrating the changing properties or reading a saturation table rather than treating this local slope as a constant.

Predict flash after throttling

A saturated liquid with inlet enthalpy 500 kJ/kg passes through an insulated valve to the pressure of the drum above. With negligible work and velocity change, enthalpy remains 500 kJ/kg. The downstream quality is therefore

$$(500 - 417.5)/2258 = 0.0365.$$

Only about 3.7% of the water flashes by mass, but the large specific volume of the vapour can make its occupied volume dominate the downstream pipe.

Trace a saturated mixture along a constant-volume path

The drum example determines one state. If a sealed rigid vessel is heated, its pressure and temperature change together along saturation while its specific volume remains fixed. Provided equilibrium is maintained and no other phase appears, the quality at every pressure is

$$x = (v - v_f)/(v_g - v_f).$$

The energy equation for the vessel is $Q = \Delta U$ when it does no boundary or shaft work. Replacing internal-energy change by enthalpy change would add the spurious $V\Delta p$ term of a rigid container.

Quality does not change solely because latent heat has entered. Differentiate $v = (1 - x)v_f + xv_g$ along the saturation curve at fixed v. This gives

$$\frac{dx}{dT} = -\frac{(1 - x)\,dv_f/dT + x\,dv_g/dT}{v_g - v_f}.$$

Saturated liquid usually expands as temperature rises while saturated vapour specific volume decreases over the ordinary water range. The sign therefore depends on the initial mixture and local property slopes. A vessel containing almost all liquid by volume can, under some conditions, follow the saturated-liquid boundary on heating rather than drying out. Determine the actual phase region from the property data after each change. When v lies outside the interval between v_f and v_g, a saturated quality formula is no longer applicable.

Find how much energy changes phase in a sealed vessel

Take 1.00 kg of water in a rigid 0.300 m³ vessel. Use rounded saturation data at 100 kPa:

$$v_f = 0.00104,\ v_g = 1.694\,\mathrm{m^3/kg},\ u_f = 417.4,\ u_g = 2506\,\mathrm{kJ/kg}.$$

At 200 kPa use $v_f = 0.00106,\ v_g = 0.8857,\ u_f = 504.5,\ u_g = 2529\,\mathrm{kJ/kg}$. The volume bounds contain 0.300 m³/kg at both pressures. Initial quality is $(0.300 - 0.00104)/(1.694 - 0.00104) = 0.1766$; final quality is

$$(0.300 - 0.00106)/(0.8857 - 0.00106) = 0.3379.$$

Thus about 0.161 kg additionally evaporates.

Mixture internal energies are $417.4 + 0.1766(2506 - 417.4) = 786.3\,\mathrm{kJ/kg}$ and $504.5 + 0.3379(2529 - 504.5) = 1188.6\,\mathrm{kJ/kg}$, giving heat input about 402 kJ. Merely multiplying evaporated mass by one latent heat would overlook the sensible energy changes in each phase and the variation of saturation properties. As an independent check, $\Delta H = \Delta U + V\Delta p$ exceeds the required heat by $0.300(200 - 100) = 30.0\,\mathrm{kJ}$. If additional heating brings saturated vapour specific volume down to 0.300 m³/kg, the last liquid disappears. Beyond that state the fixed-volume path enters the superheated region and one must determine temperature using a single-phase equation of state or superheated-vapour data.

The vapour volume fraction in the second state is $mxv_g/V = 0.3379(0.8857)/0.300 = 0.998$ approximately. Although about one third of the mass is vapour, almost every part of the vessel volume is vapour-filled. A level probe that detects a very small liquid volume is consistent with the presence of substantial residual liquid mass. This has a direct consequence for rounding: replacing v_f by zero only gives a modest error in calculated quality here, but using that same approximation to predict the remaining liquid depth gives a nonsensical zero. Retain the finite liquid volume whenever the question concerns an interface height or the mass recovered through a bottom drain.

The units of the property table matter just as much as the phase labels when predicting a small liquid inventory.

Can an equation of state predict energy changes during expansion?

A caloric model $du = c_v\,dT$ works for an ideal gas. To discover the extra term for a real fluid, let v, s and u denote molar volume, entropy and energy for a single homogeneous phase. At a fixed composition the fundamental relation is

$$du = T\,ds - p\,dv.$$

Introduce the Helmholtz energy $f = u - Ts$. Then $df = -s\,dT - p\,dv$. Equality of the two mixed partial derivatives of the smooth function $f(T, v)$ gives the Maxwell relation $(\partial s/\partial v)_T = (\partial p/\partial T)_v$.

Now expand ds in T, v. The term $T(\partial s/\partial T)_v$ equals c_v, obtained directly by holding v fixed in the fundamental relation. Substituting both coefficients gives

$$du = c_v\,dT + \left[T\left(\frac{\partial p}{\partial T}\right)_v - p\right]dv.$$

The bracket has units of pressure. For $p = RT/v$ it vanishes. For a van der Waals model $p = RT/(v - b) - a/v^2$, the repulsive volume correction cancels between the two terms and the attraction term leaves a/v^2. Assuming constant c_v, integration yields

$$\Delta u = c_v(T_2 - T_1) + a(1/v_1 - 1/v_2).$$

An insulated, rigid container contains a gas on one side of a removable internal partition and vacuum on the other. Removing the partition causes no external boundary work and exchanges no heat. States before and after settling therefore have equal internal energy. Let $T_1 = 300\,\mathrm{K}$, $v_1 = 0.00200\,\mathrm{m^3\,mol^{-1}}$, $v_2 = 0.00400\,\mathrm{m^3\,mol^{-1}}$, $a = 0.150\,\mathrm{Pa\,m^6\,mol^{-2}}$, $b = 4.0 \times 10^{-5}\,\mathrm{m^3\,mol^{-1}}$ and $c_v = 24.0\,\mathrm{J\,mol^{-1}K^{-1}}$. The volume contribution is

$$0.150(500-250) = 37.5\,\mathrm{J\,mol^{-1}}.$$

To keep energy constant the temperature falls by $37.5/24.0 = 1.56\,\mathrm{K}$, to $298.44\,\mathrm{K}$. The differential identity links the equilibrium endpoints even though the actual expansion was irreversible. The value of b matters for their pressures and cancels from this energy change.

Gas mixtures, humidity and condensation

Engineering thermodynamics and heat transfer

Translate relative humidity into a water mass balance and determine how much condensate a cooling coil removes.

FROM FIRST PRINCIPLES

Obtain humidity ratio from partial pressures and dry-air mass

Treat moist air as an ideal mixture of dry air and water vapour occupying the same volume V at temperature T. Dalton's law gives total pressure $p = p_{da} + p_v$. For each component $m_i = p_i V/(R_i T)$, so dividing vapour mass by dry-air mass cancels the shared volume and temperature:

$$\omega = \frac{m_v}{m_{da}} = \frac{R_{da}}{R_v}\,\frac{p_v}{p - p_v} \simeq 0.622\,\frac{p_v}{p - p_v}.$$

Relative humidity is a different observable. It compares actual vapour partial pressure with the saturation pressure at the same temperature. Cooling without water removal keeps ω constant, and approximately constant total pressure then keeps p_v constant until the saturation curve reaches it:

$$\varphi = p_v / p_{\mathrm{sat}}(T),$$
$$p_{\mathrm{sat}}(T_{\mathrm{dew}}) = p_v.$$

Below the dew point some water can leave the gas as liquid, while dry-air mass is conserved. A steady coil receiving humidity ratio ω_i and delivering ω_o therefore removes condensate at rate $\dot{m}_{da}(\omega_i - \omega_o)$. Inferring the same removal from relative humidity alone would ignore its dependence on temperature.

DERIVATION AND APPLICATIONS

Measure the bypass fraction of a cooling and dehumidifying coil

Humid air crossing a coil may leave warmer and wetter than air in equilibrium with the coldest effective surface. A common reduced model treats a fraction b of the dry-air stream as bypassing the surface while the remainder reaches a saturated state at the apparatus dew-point temperature. Consider a steady coil at nearly constant total pressure, neglect fan work and heat gains, and use dry-air mass as the conserved basis. Denote inlet, saturated-surface and mixed outlet states by i, s, o, with humidity ratio ω and moist-air enthalpy h per kilogram of dry air.

Water-vapour and mixed-stream enthalpy balances give two independent ways to estimate the same bypass fraction:

$$\omega_o = b\omega_i + (1 - b)\omega_s,$$
$$h_o = bh_i + (1 - b)h_s.$$

Solving either for b exposes a consistency check on measured data and the assumed single surface state:

$$b = \frac{\omega_o - \omega_s}{\omega_i - \omega_s} = \frac{h_o - h_s}{h_i - h_s}.$$

The condensed water remains a separate outlet from the coil. With condensate leaving at liquid enthalpy h_ℓ in the same reference convention, an overall balance gives positive heat removal

$$\dot{m}_\ell = \dot{m}_{da}(\omega_i - \omega_o),$$

$$\dot{Q}_{\text{out}} = \dot{m}_{da}(h_i - h_o) - \dot{m}_\ell h_\ell.$$

The mixed-state model requires

$$0 \le b \le 1.$$

It also requires incoming air wet enough for condensation on the assumed surface. A dry surface with solely sensible cooling cannot be represented by forcing a saturated leaving branch.

WORKED EXAMPLE

Evaluate air and water leaving a coil

Change the inputs and test this calculation.

Use $\omega_i = 0.0160$, $\omega_s = 0.00760$ kg/kg dry air, $h_i = 71.0$, $h_s = 29.3$ kJ/kg dry air and $b = 0.250$. Then $\omega_o = 0.00970$ and $h_o = 39.7$ kJ/kg dry air. With 1.20 kg/s dry air, the condensate rate is 0.00756 kg/s. For $h_\ell = 42.0$ kJ/kg, heat removal is

$$1.20(71.0 - 39.725) - 0.00756(42.0) = 37.2\,\text{kW}.$$

SECOND APPLICATION

Test whether three measurements share one apparatus dew point

Change the inputs and test this calculation.

On another run with the same inlet and assumed surface, measured outlet values are $\omega_o = 0.01096$ and $h_o = 41.8$ kJ/kg dry air. The humidity ratio implies $b = 0.400$, while enthalpy gives $b = 0.300$. At $b = 0.400$ the model would predict $h_o = 46.0$ kJ/kg. The 4.2 kJ/kg discrepancy indicates that measurement uncertainty, a more distributed surface temperature or another heat transfer must be considered before inferring a unique bypass fraction.

Try it · Coil bypass

Set the inputs to the values in the worked example, then vary one quantity while holding the others fixed. Compare the curve and numerical balance with the second application. The calculation uses the assumptions derived above; its readout identifies bounds such as choking, phase-region changes and inconsistent measurements.

Open this live calculation

Try the calculation

Bypass fraction inferred from humidity	**0.25**
Measured outlet enthalpy (kJ/kg dry air)	**39.725**

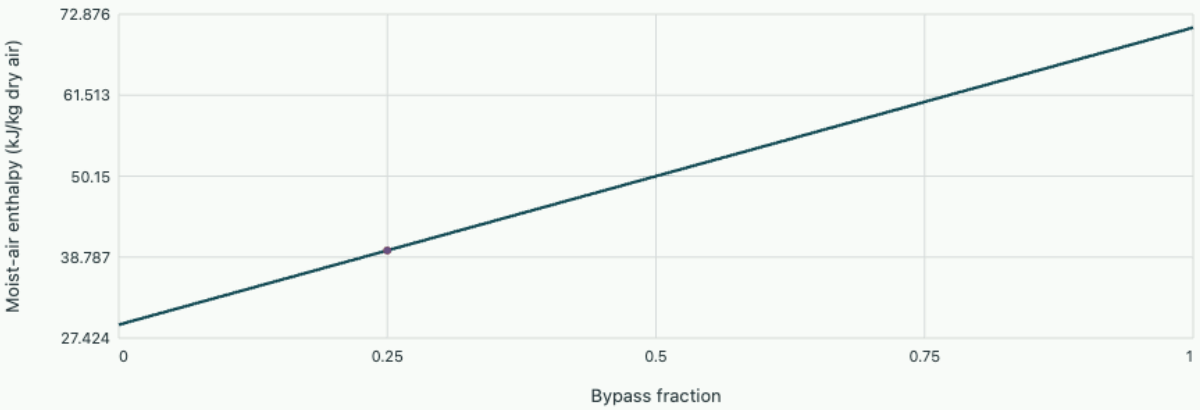

Humidity predicts ωout=0.0097 kg/kg and hout=39.725 kJ/kg; enthalpy measurement implies b=0.25. Discrepancy 0 kJ/kg. At the model state and 1.2 kg/s dry air: condensate 0.00756 kg/s; heat removal 37.2125 kW including liquid enthalpy 42 kJ/kg.

Choose a quantity that survives condensation

A ventilation duct can carry an unchanged mass flow of dry air while water evaporates into it or condenses out. It is therefore convenient to express water content and mixture energy per kilogram of dry air. Model the dry air and water vapour as ideal gases sharing a common temperature and total volume. The partial pressure of each gas is the pressure it would exert if it occupied this volume alone. Dalton's law gives

$$p = p_a + p_v.$$

Relative humidity $\phi = p_v / p_{\mathrm{sat}}(T)$ describes proximity to saturation at the current temperature. It is not the mass fraction of water. For humidity ratio $w = m_v / m_a$, divide the two ideal-gas mass equations. The

common temperature and volume cancel, leaving a ratio of molar masses multiplied by the partial-pressure ratio.

$$w = \frac{M_v}{M_a}\frac{p_v}{p - p_v} \simeq 0.622\frac{p_v}{p - p_v},$$

$$w_{\mathrm{sat}}(T) = 0.622\frac{p_{\mathrm{sat}}(T)}{p - p_{\mathrm{sat}}(T)}.$$

Before any water condenses during cooling at constant total pressure, the humidity ratio and vapour partial pressure remain fixed. Relative humidity rises because saturation pressure falls. The dew point is the temperature at which saturation pressure equals that fixed vapour pressure. Once that point is crossed, further cooling can remove water.

Find water removal and the associated cooling duty

Air enters a coil at 30 °C, 100 kPa and 60% relative humidity, and leaves saturated at 12 °C. Use $p_{\mathrm{sat}}(30) = 4.24$ kPa and $p_{\mathrm{sat}}(12) = 1.40$ kPa. The inlet vapour partial pressure is 2.544 kPa, giving $w_1 = 0.01624$ kg water per kg dry air. The outlet ratio is $w_2 = 0.00883$. For dry-air flow 0.50 kg/s the condensate flow is $0.50(0.01624 - 0.00883) = 0.00371\,\mathrm{kg/s}$, about 13.3 kg/h.

Using a 0 °C liquid-water reference, approximate moist-air enthalpy per kilogram dry air by $h = 1.005T + w(2501 + 1.86T)$ kJ/kg dry air with T in °C. The duct streams carry about 71.7 and 34.3 kJ/kg dry air. Condensate leaving at 12 °C carries $h_l \simeq 4.18(12) = 50.2\,\mathrm{kJ/kg}$. Energy conservation gives heat removed per kilogram dry air as $h_1 - h_2 - (w_1 - w_2)h_l$, approximately 37.0 kJ/kg dry air, or 18.5 kW.

Separate humidity ratio from relative humidity

Suppose the outlet stream is reheated to 20 °C without adding water. Its vapour pressure remains 1.40 kPa at the same total pressure. Since $p_{\mathrm{sat}}(20) = 2.34\,\mathrm{kPa}$, the new relative humidity is 59.8%. Neither evaporating water nor drying the air was required during this reheat. A sensor reporting 60% at both the inlet and reheated outlet therefore describes two air streams containing rather different masses of water. The coil water balance, written on the dry-air basis, explains the apparent coincidence.

Bracket an unknown dew point

At the original inlet the vapour pressure was 2.544 kPa. Suppose a saturation table gives 2.487 kPa at 21 °C and 2.644 kPa at 22 °C. Linear interpolation over this short interval puts the dew point at $21 + (2.544 - 2.487)/(2.644 - 2.487) = 21.36^{\circ}\mathrm{C}$. A coil surface warmer than this value cannot start condensing the entering air locally under the stated conditions.

Balance an adiabatic spray chamber

A direct evaporative cooler behaves differently from a refrigerated cooling coil. Water is added to an airstream, the dry-bulb temperature falls, and the humidity ratio increases. Assume steady operation,

negligible shaft work and heat leakage, unchanged total pressure, and no liquid droplets carried out. Express enthalpy per kilogram dry air. The dry-air flow is conserved; added liquid flow per kilogram dry air equals $w_2 - w_1$. The resulting energy balance is

$$h_1 + (w_2 - w_1) h_{\ell,\mathrm{in}} = h_2 .$$

Use the same enthalpy reference for liquid and moist air. If the added water has low liquid enthalpy on that reference, the moist-air enthalpy is approximately constant. Often the recirculated liquid is near the leaving-air temperature; the more accurate energy equation above can be used without that approximation. An ideal infinitely long saturated contactor has $\phi_2 = 1$, but a finite cooler generally leaves unsaturated air. An outlet temperature or a contact effectiveness is needed in addition to the balances unless saturation is specified.

With constant heat capacities one may write $h_2 = c_{pa} T_2 + w_2 (h_{g0} + c_{pv} T_2)$. Substituting this form into the balance and grouping the unknown w_2 shows explicitly which measured quantities are needed. The process cannot cool below the adiabatic saturation limit for the given inlet state and water supply merely by increasing the spray rate.

Check outlet humidity and water consumption

A cooler receives air at 35 °C and 100 kPa with inlet humidity ratio 0.0080 kg/kg dry air and dry-air flow 0.60 kg/s. Its measured outlet temperature is 23 °C and liquid makeup enters at 20 °C. Use $h = 1.005T + w(2501 + 1.86T)$ kJ/kg dry air and $h_\ell = 4.18T$ kJ/kg. The inlet carries $1.005(35) + 0.008[2501 + 1.86(35)] = 55.70\,\mathrm{kJ/kg}$ dry air; liquid enters with $h_\ell = 83.6\,\mathrm{kJ/kg}$.

The outlet equation is $55.70 + (w_2 - 0.0080)83.6 = 1.005(23) + w_2[2501 + 1.86(23)]$, giving $w_2 = 0.01297\,\mathrm{kg/kg}$ dry air. Thus liquid consumption is $0.60(0.01297 - 0.0080) = 0.00298\,\mathrm{kg/s}$, or about 10.7 kg/h. The outlet vapour pressure is

$$100(0.01297)/(0.622 + 0.01297) = 2.04\,\mathrm{kPa}.$$

At a supplied saturation pressure of 2.81 kPa for 23 °C, this corresponds to about 73% relative humidity, consistent with a finite contact process.

Treating air as a dry sensible stream would suggest a cooling duty of roughly $0.60(1.005)(35 - 23) = 7.24\,\mathrm{kW}$ removed from the cabinet. In this adiabatic device most of the sensible decrease supplies evaporation inside the boundary; no corresponding cooling-coil duty exists. A downstream process sensitive to water content may therefore react very differently to evaporative cooling and cooling plus dehumidification, even when both report the same outlet dry-bulb temperature.

An independent check can be made by weighing the makeup reservoir over a stable ten-minute test. At the predicted consumption it should lose about $10.7/6 = 1.78$ kg, provided no liquid is stored in the spray circuit during that interval. A greater measured loss together with the same outlet humidity suggests entrained droplets, drainage or leakage rather than necessarily more evaporation. A humidity probe downstream of an effective droplet separator and a separate liquid-drain measurement distinguish those paths and make the water balance close.

Identify condensation from a humidity-ratio chart

At a total pressure of 101.325 kPa,

$$w = 0.62198\, p_v / (p - p_v),$$

$$p_v = \phi\, p_{\mathrm{sat}}(T).$$

The plot uses the Buck saturation-pressure correlation over liquid water. Follow a sealed parcel cooled at constant pressure with no condensation until it reaches saturation.

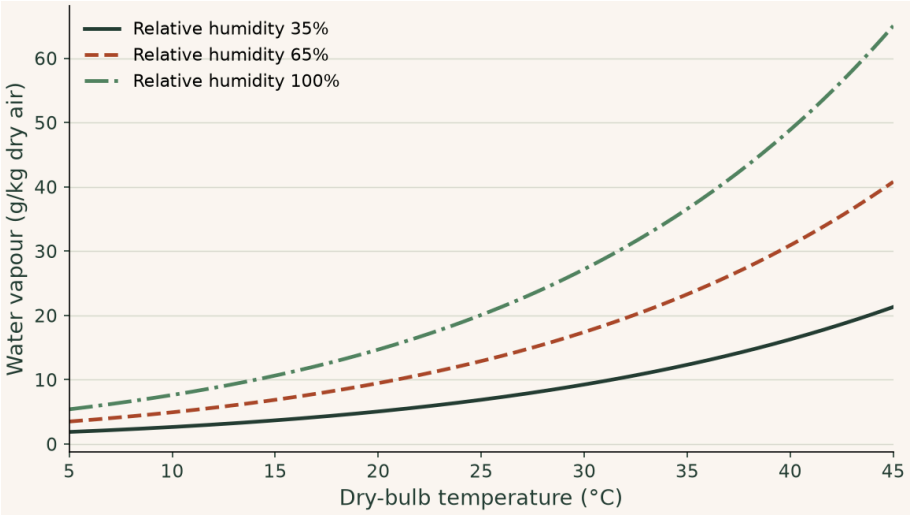

Air at 30°C and 65% relative humidity contains about 17.4 g/kg dry air. Following that horizontal humidity ratio to the saturation curve gives a dew point near 22.7°C. Continuing cooling at constant total water content then produces liquid condensate.

▶ **Read values from the plot**

Interactive calculation · change the inputs of this example

Choose an initial temperature and relative humidity, then follow the horizontal parcel-water line to saturation.

The curves below are recalculated from this section's equations in your browser. The fixed reference plot above remains available for comparison.

Open this live calculation

Try the calculation

Relative humidity	0.65
Total pressure (kPa)	101
Initial dry bulb (°C)	30

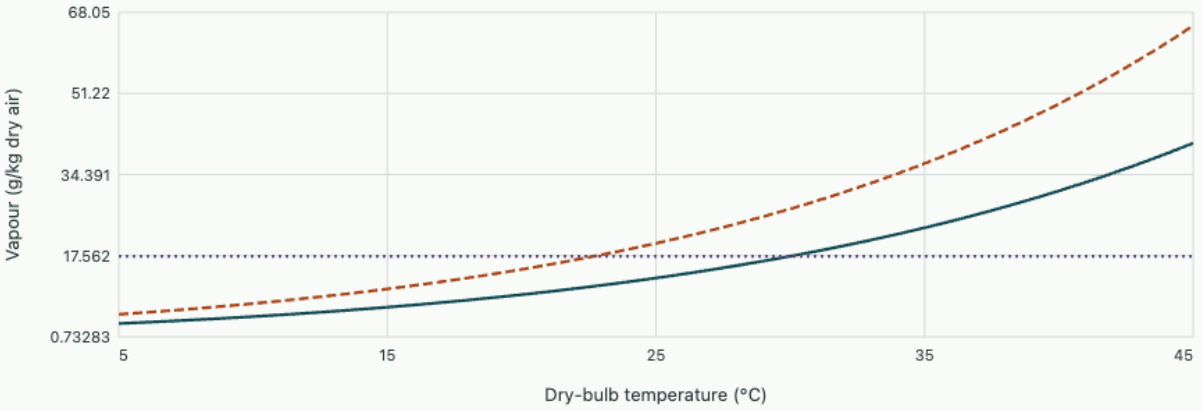

Initial humidity 17.470 g/kg dry air; dew point 22.701 °C. Below that point cooling at fixed total water produces liquid.

Useful work and exergy destruction

Engineering thermodynamics and heat transfer

Evaluate the useful-work potential of a steady stream and distinguish energy conservation from the destruction of useful-work potential.

FROM FIRST PRINCIPLES

Derive maximum useful work relative to an environment

For a closed body changing toward equilibrium with surroundings at T_0, p_0, consider the reversible path that gives the largest possible useful work. Neglect bulk kinetic and gravitational energy. With heat Q_0 positive into the body and total boundary plus shaft work W positive out, the first law gives

$$W = U_i - U_f + Q_0.$$

Reversible heat exchange with this environment requires

$$Q_0 = T_0(S_f - S_i).$$

Part of total work simply displaces ambient pressure, so

$$W_{\text{useful,rev}} = U_i - U_f + T_0(S_f - S_i) - p_0(V_f - V_i).$$

Choose the final state as the dead state, denoted by subscript zero. Grouping the available initial differences defines nonflow physical exergy:

$$B = (U - U_0) + p_0(V - V_0) - T_0(S - S_0).$$

If the process generates entropy, the same first and second law combination lowers the useful work by $T_0 S_{\text{gen}}$. The expression depends on the specified surrounding state and assumes no other heat reservoirs or chemical changes. A stream crossing a control volume carries its pressure flow work as well, giving a distinct enthalpy-based expression for flow exergy.

Background · Exergy destruction

$$\dot{X}_{\text{dest}} = T_0 \dot{S}_{\text{gen}}$$

The environmental reference temperature enters in kelvin. Irreversibility destroys useful-work potential even when the total energy balance closes.

DERIVATION AND APPLICATIONS

Derive the work potential of a compressed flowing gas

A stream can deliver useful work by cooling or by expanding, even when no fuel or chemical reaction is involved. Choose a steady reference environment at T_0, p_0 with the same gas composition. Consider one kilogram taken reversibly from its inlet state to the environmental state. Neglect kinetic and potential energy. Let q_0 be heat supplied from the environment, positive into the device, and w shaft work delivered. The steady energy balance includes flow work in enthalpy:

$$w = h - h_0 + q_0.$$

The only external heat exchange occurs with the environment. Reversibility then gives $s_0 - s = q_0/T_0$, eliminating the unmeasured heat transfer. The maximum work is the specific physical flow exergy

$$b = (h - h_0) - T_0(s - s_0).$$

For an ideal gas with constant c_p and R, insert $dh = c_p\,dT$ and $ds = c_p\,dT/T - R\,dp/p$. Integrate relative to T_0, p_0:

$$b = c_p[(T - T_0) - T_0 \ln(T/T_0)] + RT_0 \ln(p/p_0).$$

The pressure term survives at ambient temperature. Near $T = T_0$ the thermal term is quadratic in the temperature difference because energy and the ambient heat contribution cancel to first order. The formula describes work relative to a specified reference and assumes the inlet and outlet composition is the same. An actual dissipative path destroys part of the reversible work potential.

WORKED EXAMPLE

Estimate the work available from warm compressed air

Change the inputs and test this calculation.

Take air at $T = 360$ K, $p = 300$ kPa entering an environment at $T_0 = 300$ K and $p_0 = 100$ kPa. With $c_p = 1.005$ and $R = 0.287$ kJ/(kg K), the thermal contribution is $1.005[60 - 300 \ln 1.2] = 5.33$ kJ/kg. The pressure contribution is $0.287(300) \ln 3 = 94.6$ kJ/kg. Total available work is approximately 99.9 kJ/kg, mostly associated with pressure.

SECOND APPLICATION

Account for loss in an ambient-temperature throttle

Change the inputs and test this calculation.

Let 0.100 kg/s of ideal air throttle adiabatically from 300 to 100 kPa, both at 300 K. Enthalpy stays constant, no shaft work is extracted, and the exergy destruction is $0.100(94.6) = 9.46$ kW. Independent entropy accounting gives $\dot{S}_{\mathrm{gen}} = \dot{m}R \ln 3 = 0.0315$ kW/K. Its product with the ambient 300 K recovers the same 9.46 kW. A device balancing energy perfectly can therefore expend substantial work potential.

Try it · Flow exergy

Set the inputs to the values in the worked example, then vary one quantity while holding the others fixed. Compare the curve and numerical balance with the second application. The calculation uses the assumptions derived above; its readout identifies bounds such as choking, phase-region changes and inconsistent measurements.

Open this live calculation

Try the calculation

Inlet temperature (K) **360**

Selected pressure / ambient pressure **3**

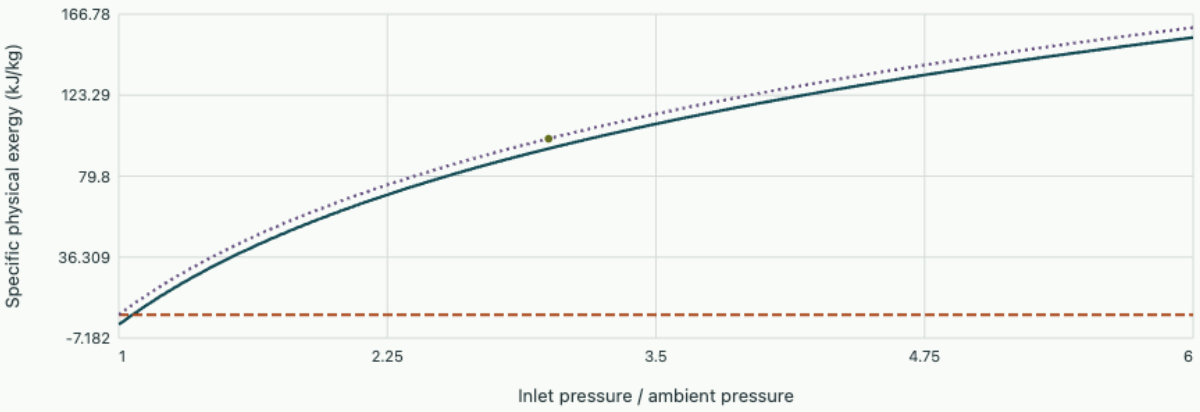

Environment 300 K, 100 kPa. Thermal 5.33005, pressure 94.5905, total 99.9206 kJ/kg. At ambient temperature an ideal-gas throttle from this pressure to ambient destroys 9.45905 kW at 0.1 kg/s without changing enthalpy.

How much work is lost when two finite bodies touch?

Two insulated metal blocks start at different temperatures. A small cyclic heat engine could transfer heat from one to the other and deliver shaft work until the temperatures equalise. Simply bringing the blocks into thermal contact also gives a common final temperature, but it conserves a different quantity. Calculate the two end states before estimating how much work the initial temperature difference made available.

Assume constant positive heat capacities C_h, C_c, no phase change, no external heat exchange and no mechanical work by either block. In direct passive contact no outside work is produced. The total energy change is therefore zero, giving $T_{\mathrm{mix}} = (C_h T_h + C_c T_c)/(C_h + C_c)$. If the intervening engine instead operates reversibly, each small entropy transfer leaving the hot block equals the entropy entering the cold one. At the final common temperature T_f,

$$C_h \ln(T_f/T_h) + C_c \ln(T_f/T_c) = 0,$$

$$T_f = \exp\left[\frac{C_h \ln T_h + C_c \ln T_c}{C_h + C_c}\right].$$

Here all temperatures are absolute. More formally the logarithms can be taken of ratios to the same reference temperature; that reference cancels. The weighted geometric mean lies below the weighted arithmetic mean. The reversible final state thus contains less internal energy than the passive final state. The difference has emerged as work,

$$W_{\max} = (C_h + C_c)(T_{\mathrm{mix}} - T_f).$$

An irreversible engine generates entropy and ends at a higher common temperature with smaller available work.

Take two blocks with capacities of 1.2 kJ/K, initially at 360 K and 280 K. Passive contact ends at 320.0 K; the reversible engine ends at $\sqrt{360 \times 280} = 317.49$ K and delivers 6.02 kJ. Passive contact generates 0.01890 kJ/K of entropy. To visualise the reversible path, let s be the entropy transferred from one block into the other. The instantaneous temperatures are $T_h(s) = T_h(0)e^{-s/C_h}$ and $T_c(s) = T_c(0)e^{s/C_c}$. The next increment of work is $(T_h - T_c)ds$, so the signed area between the two temperature curves reproduces the total work.

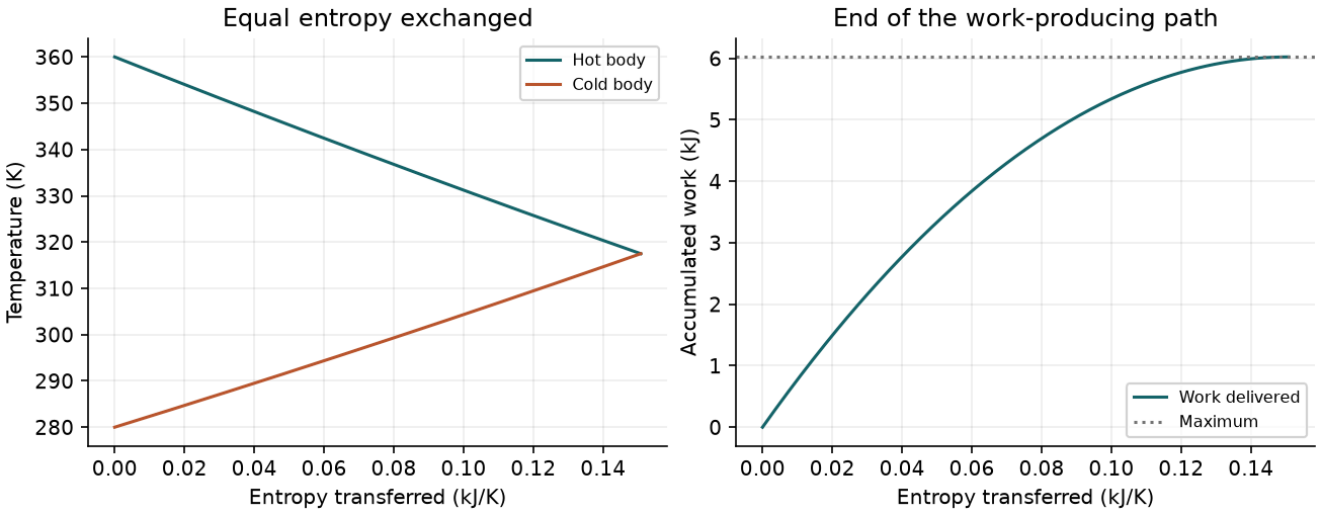

The slope of the work curve is the temperature gap. It vanishes at the reversible end state even though the engine delivered finite work earlier.

The scaling becomes particularly clear for equal capacities C when the starting temperatures are $T_m + \Delta$ and $T_m - \Delta$. The mixed temperature is T_m, while the reversible final temperature is $T_m\sqrt{1 - (\Delta/T_m)^2}$. Expanding the square root for small temperature gaps gives

$$W_{\max} \simeq C\Delta^2/T_m.$$

For $T_m = 320$ K and $\Delta = 40$ K this is 6.0 kJ, close to the exact 6.02 kJ. Halving the temperature gap at the same mean temperature gives roughly a quarter of the work. The energy initially contained in the warmer block above the mean temperature scales only linearly with the gap, so that energy alone is a poor measure of usable work.

If the surroundings never exchange heat with either block, should their temperature appear in this maximum-work calculation?

▼ **Work through the solution**

The reversible engine can use only the two specified bodies. Their initial states and conservation of their total entropy determine how far the work extraction goes. A different problem allows ultimate equilibration with an external reservoir at T_0. In that enlarged problem the standard exergy destruction associated with the irreversible contact is $T_0 S_{\mathrm{gen}}$. These work comparisons have different permitted end states. The notebook keeps the enclosure insulated and checks the work using the energy balance on both bodies throughout the exchange.

Interactive calculation · Finite source temperatures

Vary the notebook inputs and compare both plotted observables. The same equations generate the browser curves and the notebook reference; fixed material, geometry and model assumptions are stated in the readout.

Open this live calculation

Try the calculation

Hot initial temperature (K)	360
Cold initial temperature (K)	280
Hot / cold heat capacity	1

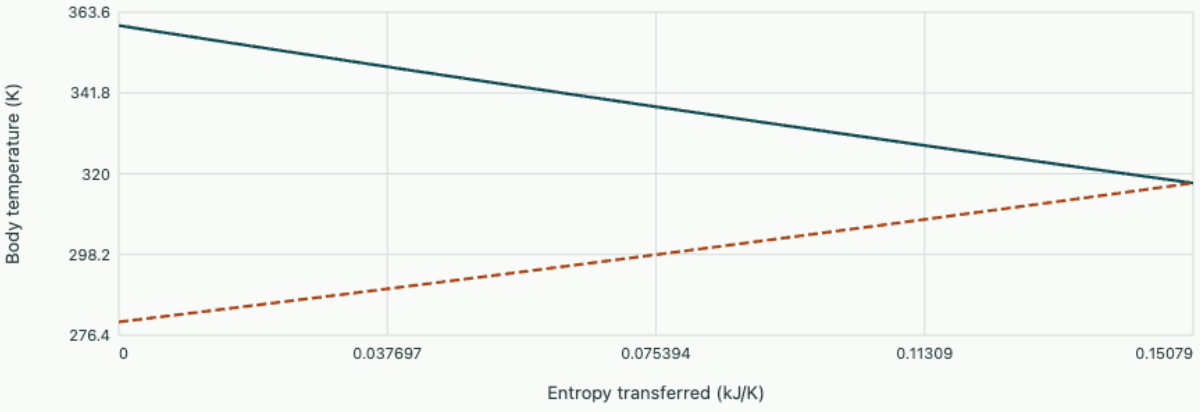

Common reversible temperature 317.490 K; maximum work 6.024 kJ. Cold heat capacity1.2 kJ/K. Direct contact reaches a warmer arithmetic-weighted mean and generates entropy.

Download the editable notebook for this example

Interactive calculation · Finite source work

Vary the notebook inputs and compare both plotted observables. The same equations generate the browser curves and the notebook reference; fixed material, geometry and model assumptions are stated in the readout.

Open this live calculation

Try the calculation

Hot initial temperature (K)	360
Cold initial temperature (K)	280
Hot / cold heat capacity	1

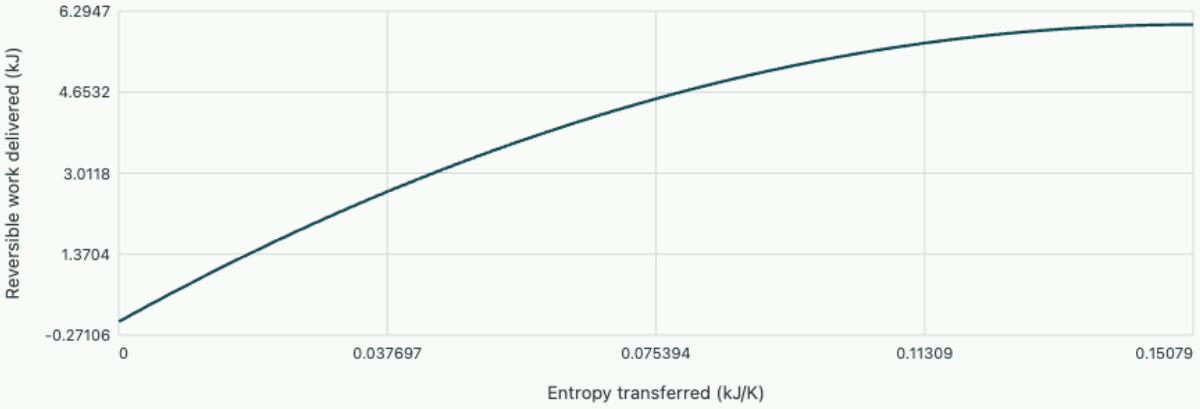

Accumulated reversible work

Common reversible temperature 317.490 K; maximum work 6.024 kJ. Cold heat capacity1.2 kJ/K. Direct contact reaches a warmer arithmetic-weighted mean and generates entropy.

Download the editable notebook for this example

Open editable notebook Download .ipynb

Specify the surroundings before rating a resource

A warm stream contains energy relative to a chosen reference. The amount that could become useful work depends on the surroundings available to receive heat. If stream and environment reach mutual thermal and mechanical equilibrium, there is no further physical work potential associated with their difference. This equilibrium is the physical dead state at T_0, p_0. Exclude chemical reactions here and consider a steady flow with constant composition.

Imagine a reversible device that takes a stream from its current state to the dead state while interacting thermally only with the environment. The steady first law relates delivered work to changes in stream enthalpy, kinetic and potential energy, and heat. The reversible entropy balance gives heat received from the environment $q_0 = T_0(s_0 - s)$. Eliminating that heat gives the specific flow exergy

$$\psi = (h - h_0) - T_0(s - s_0) + V^2/2 + g(z - z_0).$$

Use consistent units for the kinetic and potential terms. The entropy contribution reduces the fraction of enthalpy that a reversible device can turn into useful work. At the dead state, with the same velocity and elevation as the reference, that work potential vanishes.

Cool a liquid to the environment

Take a stream of liquid water at approximately ambient pressure, 360 K, with negligible motion and elevation change. It flows at 0.20 kg/s towards a large environment at 300 K. Assume $c_p = 4.18\,\mathrm{kJ/(kg\,K)}$ constant. Integration gives

$$h - h_0 = c_p(T - T_0),$$
$$s - s_0 = c_p \ln(T/T_0).$$

The stream's specific physical exergy is therefore

$$4.18[60 - 300\ln(360/300)] = 22.2\,\mathrm{kJ/kg}.$$

Its enthalpy difference is much larger, 250.8 kJ/kg.

Cooling all the way to 300 K makes at most $0.20(22.2) = 4.44\,\mathrm{kW}$ of useful work available, even though the stream releases 50.16 kW of heat. Direct heat exchange with the environment delivers no shaft work. The environment entropy gain is $50.16/300 = 0.1672\,\mathrm{kW/K}$, while the stream loses $0.20(4.18)\ln(360/300) = 0.1524\,\mathrm{kW/K}$. The difference, about 0.0148 kW/K, is the entropy production rate. Multiplication by T_0 gives the lost 4.44 kW of work potential.

Include heat crossing another temperature boundary

Heat Q supplied reversibly by a reservoir at a uniform boundary temperature T_b carries exergy $Q(1 - T_0/T_b)$. This follows by comparing its energy with the environmental heat required to discharge the associated entropy Q/T_b. A heat input of 10 kW from a 600 K reservoir in a 300 K environment carries 5 kW of exergy. A transfer from a finite cooling stream has changing source temperature and requires integration, as in the water example. Irreversibility in the device destroys exergy at rate $T_0\dot{S}_{\mathrm{gen}}$; directing heat to a useful secondary process can retain part of the exergy that direct cooling would destroy.

Include the work potential of pressure

For an ideal gas already at T_0, an elevated pressure still gives physical flow exergy $RT_0\ln(p/p_0)$. Air at 300 K and twice environmental pressure has about $0.287(300)\ln 2 = 59.7\,\mathrm{kJ/kg}$ of pressure-related work potential although its ideal-gas enthalpy equals that at the dead state.

Derive the least work needed to heat a finite stream

Heat pumps are commonly compared by heat delivered per unit of electrical work. A second-law comparison additionally accounts for the temperature at which heat is delivered. Consider a steady stream heated from T_i to T_o at nearly environmental pressure, with constant heat capacity rate $C = \dot{m}c_p$. A reversible heat pump extracts heat from surroundings at T_0 and delivers it through a temperature-matched sequence of infinitesimal condensers. For one increment $dQ = C\,dT$ at stream temperature T, the reversed Carnot relation gives

$$dW_{\min} = dQ(1 - T_0/T).$$

Integrating over the stream yields

$$W_{\min} = C[(T_o - T_i) - T_0 \ln(T_o/T_i)].$$

This is the increase of its physical flow exergy, and is smaller than evaluating the Carnot work factor at the final high temperature for the whole heat duty. A device that condenses at a single temperature above the warmest outlet generates entropy during delivery to the colder portions. The result assumes that the source of heat is the environment at T_0; a colder source would introduce an additional resource requirement.

For steady equipment receiving only electricity and environmental heat, with negligible parasitic flows, a meaningful second-law efficiency is the stream’s exergy increase divided by actual electrical input. Electric power is entirely work potential, whereas a kilowatt of heat near room temperature carries only a small fraction of a kilowatt of exergy.

Compare a heat pump with a resistance heater

Water flows at 0.10 kg/s from 293 K to 333 K in a 293 K environment. With $c_p = 4.18\,\mathrm{kJ/(kg\,K)}$, heat duty is $0.418(40) = 16.72\,\mathrm{kW}$. The exergy increase is $0.418[40 - 293\ln(333/293)] = 1.05\,\mathrm{kW}$ approximately. This relatively small figure reflects the modest average temperature lift. A well-insulated heat pump drawing 4.20 kW from the electrical supply has heating COP $16.72/4.20 = 3.98$ and exergy efficiency about 25%. Its surroundings provide 12.52 kW of heat under the stated energy balance.

A resistance heater supplying the identical water service draws 16.72 kW when its losses to the room are negligible. It is nearly 100% efficient on the simple electrical-energy-to-heat basis, yet its exergy efficiency is only about 6.3%. For the heat pump the implied destruction of work potential is $4.20 - 1.05 = 3.15\,\mathrm{kW}$; for resistance heating it is about 15.67 kW. Electricity used by a circulating pump, external heat leaks or a source below environmental temperature belongs in the appropriate device boundary before drawing a practical comparison. The heat-duty equality ensures that the machines are judged against the same delivered temperature service.

For the resistance case the entropy balance can be checked without using the exergy formula. Electrical work carries no entropy across the control surface, and assume no heat leakage. The stream increases its entropy at $0.418\ln(333/293) = 0.0535\,\mathrm{kW/K}$. Multiplying by the 293 K environment gives 15.7 kW destroyed, consistent with the difference between input work and stream exergy gain. This agreement depends on keeping the heater and heated stream in the same chosen boundary.

Why a real gas cools or warms across a throttle

Engineering thermodynamics and heat transfer

A gas passes through an insulated porous plug. Its pressure drops and its measured temperature changes. Start with the energy flowing through the plug and obtain a prediction from the equation of state.

Establish which property travels through the plug

A steady control volume with one inlet, one outlet and no shaft work has $h_1 + u_1^2/2 + g z_1 = h_2 + u_2^2/2 + g z_2$ when the casing is insulated. Choose the measurement stations in large, level pipes where changes in mean kinetic energy are negligible. Then $h_1 = h_2$. The pressure gradient is sustained by irreversible friction in the plug. Along the family of equilibrium outlet states, a small pressure change produces $dT = \mu_{JT} dp$, with $\mu_{JT} = (\partial T/\partial p)_h$. This family lets us compute endpoints without resolving the interior pore flow.

Recover the pressure dependence of enthalpy

Use molar properties. For a simple compressible pure substance, $dh = T\,ds + v\,dp$. View entropy as a function of T, p. At fixed pressure its temperature derivative is c_p/T. To obtain its pressure derivative, differentiate $g = h - Ts$. Its differential is $dg = -s\,dT + v\,dp$. Equality of mixed derivatives gives

$$(\partial s/\partial p)_T = -(\partial v/\partial T)_p.$$

Substitution yields

$$dh = c_p\,dT + \left[v - T\left(\frac{\partial v}{\partial T}\right)_p\right] dp,$$

$$\mu_{JT} = \frac{T(\partial v/\partial T)_p - v}{c_p}.$$

For an ideal gas $v = RT/p$ the two numerator terms cancel at all temperatures. Real gas attractions and finite molecular volume affect this near cancellation. Since $dp < 0$ for the pressure drop, positive μ_{JT} indicates cooling. Check units by dividing cubic metres per mole by joules per mole per kelvin. The result is kelvins per pascal.

An instructive gas with an inversion temperature

Represent a dilute model gas by $v = RT/p + B(T)$, with $B = b - a/(RT)$. Here a, b are constants and the omitted equation-of-state terms become important at higher density. The ideal part cancels identically. Differentiating the correction gives

$$\mu_{JT} \simeq \frac{T B'(T) - B(T)}{c_p} = \frac{2a/(RT) - b}{c_p},$$

$$T_{\mathrm{inv}} = \frac{2a}{Rb}.$$

With $b = 4.0 \times 10^{-5}$ m³/mol, $a = 0.080$ Pa m⁶/mol², $R = 8.314$ J/(mol K) and $c_p = 30$ J/(mol K), the inversion lies near 481 K. At 300 K, $\mu_{JT} = 0.0805$ K/bar. A modest drop from 3 bar to 1 bar cools this gas by about 0.161 K. At 600 K the derivative is about −0.0264 K/bar. The same pressure drop then warms the gas by 0.0528 K. Treat this as a local calculation. Over a large temperature range one should integrate the property derivative using a validated equation of state.

Alter the pressure drop while holding the attractive constant fixed. Locate the zero crossing at two pressure drops. Explain why it remains at the same temperature in this first-order calculation and why the cooling magnitudes change.

Open this experiment online

Try the calculation

Attractive constant a (Pa m⁶/mol²) **0.08**

Pressure drop (bar) **2**

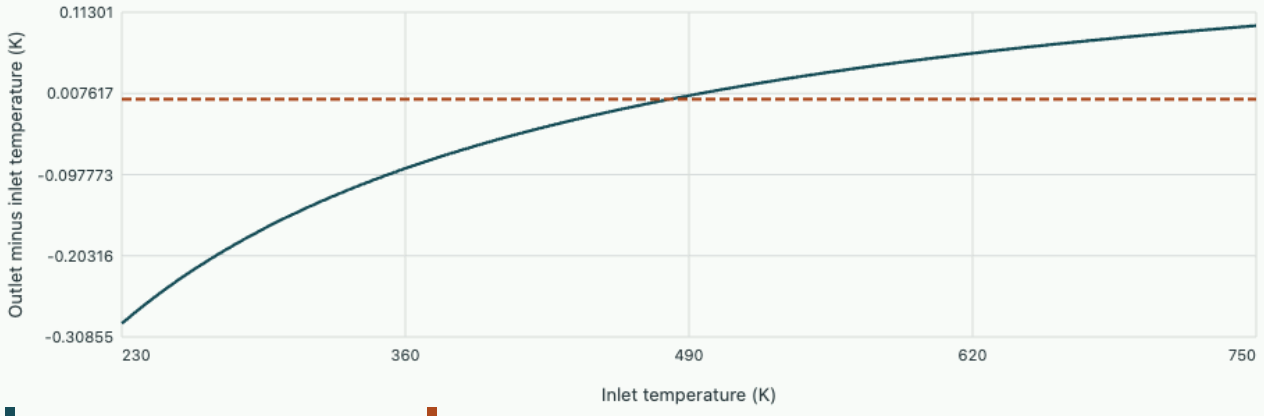

The inversion temperature is 481.1 K. At 300 K the local pressure-drop estimate gives an outlet temperature change of -0.161 K.

Tutorial · infer the attractive correction

A low-density gas at 320 K gives a positive slope of 0.060 K/bar in a throttle experiment. Assume the same b, c_p as above. Estimate a and the inversion temperature. How would an unaccounted heat leak from a warm room bias the inference if the gas cools?

▼ **Show the inference and sign check**

Convert the slope to 6.0×10^{-7} K/Pa. Rearranging gives $a = (RT/2)(c_p \mu_{JT} + b) = 0.0772$ Pa m⁶/mol², whence $T_{\mathrm{inv}} = 464$ K. A warm leak partially cancels cooling, making the observed outlet temperature drop too small. Because the pressure drop is also negative, the apparent positive slope is smaller. In this parameter fit the estimated a is therefore too low.

Further reading · MIT 2.43 · property relations and the Joule–Thomson coefficient

Finding the vapour fraction of an equilibrium flash

Engineering thermodynamics and heat transfer

A mixture enters a separator at given overall composition, pressure and temperature. Some components prefer the vapour more strongly than others. Derive a scalar equation for the fraction that evaporates and test for a two-phase solution before iterating.

Connect composition to the equilibrium ratios

Let $F = L + V$ be the total molar feed and liquid and vapour flow rates. Overall component fractions are z_i, liquid fractions x_i and vapour fractions y_i. At equilibrium the chemical potential of each species agrees across phases. In an ideal gas over an ideal liquid this reduces to Raoult's law $y_i p = x_i p_i^{\mathrm{sat}}(T)$, so $K_i = y_i / x_i = p_i^{\mathrm{sat}}(T)/p$. A larger K enriches that species in the vapour. For a real mixture K can depend on phase composition and require an outer property iteration. Here use fixed positive K values to isolate the material balance.

Eliminate the component flows

Define $\beta = V/F$. Component conservation gives

$$z_i = (1-\beta)x_i + \beta y_i = x_i[1 + \beta(K_i - 1)].$$

Thus $x_i = z_i / [1 + \beta(K_i - 1)]$. For both compositions to sum to one, enforce $\sum_i (y_i - x_i) = 0$, giving the Rachford–Rice equation:

$$R(\beta) = \sum_i \frac{z_i (K_i - 1)}{1 + \beta(K_i - 1)} = 0,$$

$$R'(\beta) = -\sum_i \frac{z_i (K_i - 1)^2}{[1 + \beta(K_i - 1)]^2}.$$

The denominators are positive for $0 \le \beta \le 1$. With at least one non-unity K of positive feed fraction, R decreases strictly, so at most one physical root exists. A positive R(0) and a negative R(1) bracket two phases. If $R(0) \le 0$ the stable fixed-K split is all liquid; if $R(1) \ge 0$ it is all vapour, with equality corresponding to incipient phase appearance. When every K is one this balance cannot specify β. A bracketed numerical method preserves the meaningful interval even near a phase boundary.

Perform a binary flash and change the feed composition

Choose $K_A = 2,\ K_B = 0.5,\ z_A = z_B = 0.5$. Then

$$R(0) = 0.25,\ R(1) = -0.25.$$

Multiplying by the two denominators gives $0.25 - 0.50\beta = 0$, hence $\beta = 0.50$. Liquid fractions are $x_A = 1/3,\ x_B = 2/3$, and vapour fractions are

$$y_A = 2/3,\ y_B = 1/3.$$

For 10 mol/s feed each outlet carries 5 mol/s. The A balance reads $5(1/3) + 5(2/3) = 5$ mol/s.

Increase z_A to 0.60 at the same T, p and K. The physical root becomes $\beta = 0.80$. The phase compositions remain at the same two endpoints; their proportions change to accommodate the new feed. Indeed the overall A fraction is

$$0.20(1/3) + 0.80(2/3) = 0.60.$$

At $z_A = 0.75$ the feed lies beyond the vapour endpoint. Then R(1) is positive, and searching outside the interval would produce a vapour fraction above unity. Report a single vapour phase instead. An adiabatic flash at unknown temperature would require a further enthalpy balance.

Start from equal feed fractions and increase the fraction of A until the root leaves the physical interval. Predict which pure-phase endpoint is reached. Compare residual signs at both endpoints before interpreting the reported phase.

Open this experiment online

Try the calculation

Feed fraction of A **0.5**

Equilibrium ratio K of A **2**

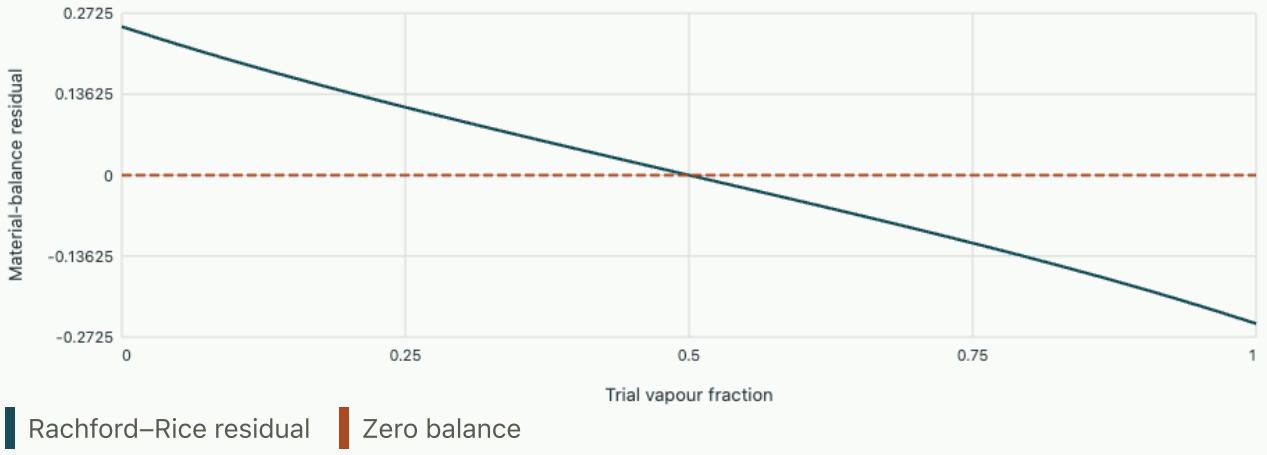

Bracketed balance gives vapour fraction 0.5000. A fractions in the liquid and vapour are 0.3333 and 0.6667.

Tutorial · determine whether an iteration is warranted

At an operating point take $K_A = 3,\ K_B = 0.6$. Compute the liquid and vapour A fractions when both phases coexist. For $z_A = 0.30$, calculate β. Repeat the phase test for $z_A = 0.10$ and explain why normalized hypothetical phase compositions can hide a failed balance.

▼ **Solve phase closure and the lever rule**

Vapour closure gives $3x_A + 0.6(1 - x_A) = 1$, hence $x_A = 1/6$, $y_A = 1/2$. For overall fraction 0.30,

$$\beta = (0.30 - 1/6)/(1/2 - 1/6) = 0.40.$$

At $z_A = 0.10$,

$$R(0) = 0.10(2) + 0.90(-0.4) = -0.16.$$

The feed remains liquid. Renormalizing an unphysical trial vapour composition could make it sum to unity while violating the original component and equilibrium equations. Check the physical endpoint signs first.

Further reading · Penn State, PNG 520 · elementary vapour–liquid equilibrium, flash objective function and numerical solution

When does a polar crystal change its preferred state?

Fields and the response of matter

Derive stable and unstable dielectric responses from a free-energy landscape. Separate a reversible small-signal capacitance from a switching experiment.

Why one voltage can support several measured polarisations

Some insulating crystals develop a spontaneous electric polarisation as temperature is lowered. The high-temperature phase may be symmetric under inversion so that states with polarisation $+P$ and $-P$ are equivalent in zero field. Near a continuous transition, treat one component of polarisation as an order parameter. A local free-energy density compatible with the symmetry is $f(P) = a(T - T_c)P^2/2 + bP^4/4$ with $a, b > 0$. The omitted terms limit how far in temperature or field this truncated expansion is reliable. The positive quartic term keeps the energy bounded when the quadratic coefficient is negative.

For a device whose applied electric field E is controlled, the appropriate potential in this simple model is $g(P; E) = f(P) - EP$. The linear term records external work; changing what is held fixed changes the energy to minimise. Differentiation gives the constitutive equation $E = a(T - T_c)P + bP^3$. An algebraic root is an equilibrium candidate. A stable local response also requires positive second derivative $a(T - T_c) + 3bP^2$. A very accurate rootfinder can return an unstable stationary point if this second test is omitted.

Find the equilibrium branches and their response

Introduce scaled polarisation p, field e and temperature parameter s such that $g = p^4/4 + sp^2/2 - ep$. Then $e = p^3 + sp$. For $s > 0$, the slope $3p^2 + s$ stays positive. Each field has one stable solution. At $e = 0$, $p = 0$ and the small-signal susceptibility $dp/de = 1/s$ increases when approaching the ideal critical point from above.

For $s = -1,\ e = 0$, stationary values are $-1, 0, +1$. The middle root has curvature –1 and is a local maximum. Each outer root has curvature 2 and is a minimum. The susceptibility about a selected outer state is $1/2$ in the scaled units. The fact that the two domains can carry opposite zero-field polarisation does not make the harmonic response negative. The mere existence of a negative-slope segment in the cubic constitutive curve is insufficient to infer a stable isolated small-signal device with that differential response.

Derive the limit of a metastable homogeneous state

Imagine preparing the negative-polarisation state at $s = -1$, then increasing the applied field slowly. A positive field lowers the energy of the positive state. The negative state can persist locally behind a barrier even after it ceases to be the global minimum. A model of perfectly homogeneous evolution without early nucleation follows this metastable branch until the curvature vanishes. Solve $3p^2 - 1 = 0$, choose $p = -1/\sqrt{3}$, and insert it into $e = p^3 - p$. The positive limiting field is

$$e = 2/(3\sqrt{3}) \simeq 0.3849.$$

At the opposite spinodal the signs reverse.

This is a limit of the particular energy model. Real switching can occur earlier through domain walls, nucleation, defects or thermally activated barrier crossing. Observing a smaller coercive field in a specimen is not evidence that the cubic equation was solved inaccurately. To compare an actual hysteresis loop with a calculation, state the sweep rate, boundary circuit, sample preparation and a dynamical or nucleation model in addition to equilibrium free energy.

Read a landscape rather than tracing an unlabelled loop

Open the live experiment

Try the calculation

Temperature coefficient s **-1**

Applied field e **0.18**

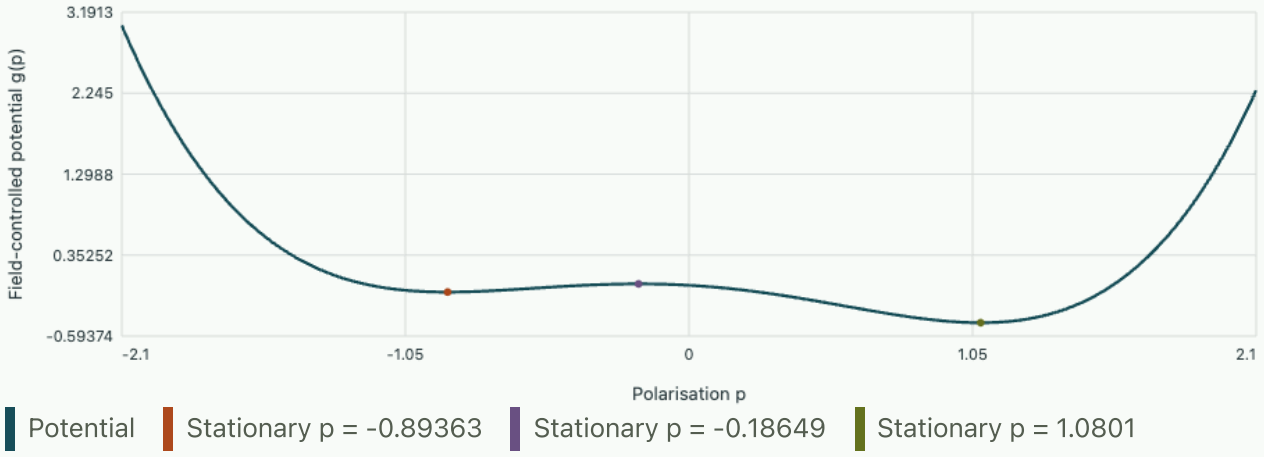

Stationary polarisations: -0.89363 (minimum), -0.18649 (maximum), 1.0801 (minimum).

Set $s = -1$. At zero applied field, locate the two equal wells. Increase positive e gradually and watch the negative well become shallower. The browser locates real stationary points of the cubic, classifies them by curvature and reports their positions. Its plotted potential shows the energy barrier separating basins. Choose $s > 0$ and repeat to see why fitting the same S-shaped hysteresis loop above a continuous transition would be inappropriate.

A below-transition sample can also form spatially varying domains. Adding a positive gradient term $K(\nabla P)^2/2$ penalises interfaces, while depolarising fields can favour a mixture with reduced net surface charge. This competition sets new length scales absent from the homogeneous model. An electrode measuring total charge can return near-zero mean polarisation from oppositely oriented domains even though each domain is strongly polarised. A local probe and a bulk capacitance measurement therefore constrain different parts of the model.

Turn a capacitance measurement into a consistency test

For a small harmonic voltage imposed about a prepared local minimum, the differential response is $\delta P/\delta E = 1/[a(T - T_c) + 3bP^2]$, provided the driving frequency is slow enough for this local variable to respond and the amplitude is small enough to stay in the same well. Add the vacuum and fast polarisation terms when converting to total displacement D. In a parallel-plate geometry capacitance per area then follows from $\delta D/\delta E$ divided by thickness.

A large-signal charge–voltage measurement crosses different branches and may dissipate work around a loop. The energy per volume over a closed polarisation cycle is $\oint E\,dP$ with the sign convention for electrical work. A single equilibrium graph of $f(P)$ returns to its initial state with no prescribed loop area. Measured hysteresis thus carries dynamical information that a polynomial fit to one reversible branch cannot uniquely reconstruct.

Tutorial

For $s = -4$ in the scaled model, find zero-field minima and the positive homogeneous spinodal field.

▼ **Work through the solution**

Zero-field minima lie at $p = \pm 2$. The negative spinodal is

$$p = -\sqrt{4/3} = -2/\sqrt{3}.$$

Substitution gives positive field $16/(3\sqrt{3}) \simeq 3.08$. The $|s|^{3/2}$ dependence follows from the cubic.

A Newton solver initialised at $p = 0$ reports exact convergence for $s = -1, e = 0$. What additional observation rejects that result as the prepared stable state?

▼ **Work through the solution**

Its residual is exactly zero but the curvature $3p^2 + s = -1$ is negative. A small perturbation lowers the potential. Check a measured weak-field slope or the relaxation after a perturbation, and compare the energy of the two minima.

Why can the homogeneous coercive field be much larger than the measured coercive field?

▼ **Work through the solution**

The calculation removes a local barrier only when a spatially uniform state ceases to exist. A real material can switch by creating a small favourable domain and moving a wall; electrodes and defects can lower the needed activation energy. The relevant material and time scales differ.

Compare material descriptions

Use MIT's ferroelectrics and polarisation material to distinguish a microscopic source of displacement from the macroscopic free-energy model used here.

Moving constraints, boundary work and a changing energy

Analytical mechanics from constraints to invariants

A pendulum with a reeled string can gain energy while its support force does no virtual work. A bead on a spinning guide can accelerate without gravity. Follow the work done by the apparatus, derive the relation between a reduced energy and laboratory energy, and identify which quantity a measurement tests.

Parameterise the motion of a guide as well as the bead

Let a prescribed moving guide define positions $\mathbf{r}(q,t)$. Its relative velocity along the guide is $\mathbf{r}_q\dot q$ and its imposed local velocity is $\mathbf{r}_t$, where t is differentiated at fixed coordinate. The kinetic energy $T = m|\mathbf{r}_q\dot q + \mathbf{r}_t|^2/2$ includes a quadratic term in $\dot q$, a mixed term linear in $\dot q$, and a term independent of $\dot q$. Dropping the last two merely because the guide motion is prescribed generally changes the bead equation. A term independent of $\dot q$ may still depend on its position q.

For $L = T - V$ define $p = \partial_{\dot q} L$ and $\mathcal{E} = p\dot q - L$. Differentiating along a solution of the Euler–Lagrange equation gives $d\mathcal{E}/dt = -\partial_t L$ in the absence of extra nonconservative generalised forces. This identity involves the partial time derivative at fixed $q, \dot q$. The laboratory energy is

$$E_{\mathrm{lab}} = m|\dot{\mathbf{r}}|^2/2 + V.$$

Direct substitution gives $\mathcal{E} = E_{\mathrm{lab}} - m\dot{\mathbf{r}}\cdot\mathbf{r}_t$ for a single particle and ordinary position potential. Thus a time-independent reduced Lagrangian can conserve a quantity different from the laboratory mechanical energy. Which observer supplies a force, and which independent variables are held fixed, determine the energy being tracked.

Work supplied by a steadily rotating straight slot

A smooth horizontal guide turns about its inner end with prescribed angle $\varphi = \Omega t$. Let r be the bead's distance from the pivot. The velocity is $\dot r\hat r + r\Omega\hat\varphi$, and the reduced Lagrangian is $m(\dot r^2 + \Omega^2 r^2)/2$. Its Euler–Lagrange equation is

$$\ddot r = \Omega^2 r.$$

With $r(0) = r_0, \dot r(0) = 0$ the solution is $r_0\cosh\Omega t$ until the bead leaves the slot. In the laboratory no force acts along the radial tangent of the slot, yet radial acceleration in polar coordinates is $\ddot r - r\Omega^2$, precisely zero. The tangential acceleration $2\Omega\dot r$ is provided by the slot wall.

The wall reaction $R_\varphi = 2m\Omega\dot r$ exerts a torque rR_φ on the bead. The driving motor balances the opposite reaction on the slot and supplies power $\Omega r R_\varphi = 2m\Omega^2 r\dot r$. Differentiating $E_{\mathrm{lab}} = m(\dot r^2 + \Omega^2 r^2)/2$ and using the radial equation gives the same power. The reduced energy $m(\dot r^2 - \Omega^2 r^2)/2$ is constant. Its difference from E_{lab} equals $-\Omega L_z$, with $L_z = mr^2\Omega$. This is a consequence of uniform imposed rotation, not a negative physical rotational kinetic energy.

For $m = 0.10\,\mathrm{kg}$, $r_0 = 0.050\,\mathrm{m}$ and $\Omega = 4.0\,\mathrm{rad/s}$, consider $t = 0.25\,\mathrm{s}$. The radial position is $77.2\,\mathrm{mm}$, radial speed $0.235\,\mathrm{m/s}$, and wall reaction $0.188\,\mathrm{N}$. The motor supplies about $0.0580\,\mathrm{W}$ at that instant. Laboratory energy has risen from $0.00200\,\mathrm{J}$ to approximately $0.00752\,\mathrm{J}$. Measuring only the angular motor

speed would miss the incoming mechanical work. The torque on the drive shaft or the changing moment of inertia supplies the other observable required for an energy budget.

A pendulum with prescribed length

Let a point pendulum have length $l(t) > 0$ and angle θ from downward vertical. Using $x = l\sin\theta$, $y = -l\cos\theta$ gives $T = m(\dot{l}^2 + l^2\dot{\theta}^2)/2$ and $V = -mgl\cos\theta$ for a fixed gravitational datum. The angular Euler–Lagrange equation becomes

$$\frac{d}{dt}(ml^2\dot{\theta}) = -mgl\sin\theta,$$

$$\ddot{\theta} + 2\frac{\dot{l}}{l}\dot{\theta} + \frac{g}{l}\sin\theta = 0.$$

The term proportional to $\dot{l}\dot{\theta}$ comes directly from differentiating angular momentum about the pivot. With increasing length the angular speed at fixed angular momentum decreases; reeling in has the opposite effect. It can look like damping or amplification in the angular equation even with no friction. A mechanical energy balance distinguishes the reversible exchange with a reel from energy lost as heat.

The radial Newton equation, taking outward as positive, is $m(\ddot{l} - l\dot{\theta}^2) = mg\cos\theta - T_{\mathrm{string}}$. Hence $T_{\mathrm{string}} = m(g\cos\theta + l\dot{\theta}^2 - \ddot{l})$. The tension power on the bob is $-T_{\mathrm{string}}\dot{l}$. This has the correct sign when a heavy nearly vertical bob is lowered slowly: $\dot{l} > 0$ and the gravitational potential decreases as the string takes energy out. At the bottom of a swing the tension is larger, so a reel shortening the string there does more positive work per unit length than it would at a turning point. Timing a cyclic length change can pump an oscillation while the guide returns to precisely its initial length.

With $m = 0.50\,\mathrm{kg}$, $l = 1.0\,\mathrm{m}$, $\dot{\theta} = 2.0\,\mathrm{rad/s}$, and the bob at the bottom, a short interval of constant-rate shortening at $\dot{l} = -0.030\,\mathrm{m/s}$ has tension $6.91\,\mathrm{N}$ and delivers $0.207\,\mathrm{W}$. When the bob is momentarily at rest at 30°, the tension for the same length and zero radial acceleration is $4.25\,\mathrm{N}$. Reversing the reel rate there would remove only $0.127\,\mathrm{W}$. A complete finite cycle requires integrating the time-varying tensions and checking the acceleration during each reel reversal.

Translate a pivot and separate real forcing from a coordinate choice

For a pivot translating horizontally at prescribed $X(t)$ and constant string length l, write $x = X + l\sin\theta$, $y = -l\cos\theta$. The kinetic energy includes $ml\dot{X}\dot{\theta}\cos\theta$. In the angular Euler–Lagrange equation the terms proportional to $\dot{X}\dot{\theta}\sin\theta$ cancel, leaving $l\ddot{\theta} = -g\sin\theta - \ddot{X}\cos\theta$. A constant pivot velocity therefore does not force the relative angle. Replacing the pivot velocity by its magnitude as a driving force would wrongly violate Galilean invariance. A horizontal pivot acceleration does force the angle; it appears as the opposite inertial acceleration in the translating frame.

For $|\theta| \ll 1$ and $X = X_0\cos\omega t$, the equation becomes

$$\ddot{\theta} + \omega_0^2\theta = (X_0\omega^2/l)\cos\omega t,$$

$$\omega_0^2 = g/l.$$

At drive frequencies well below ω_0,

$$\theta \simeq -\ddot{X}/g.$$

The bob nearly aligns with the effective gravity in the moving frame. Far above ω_0, the particular solution is approximately $-X/l$; then the bob's absolute horizontal coordinate $X + l\theta$ changes little even while the pivot moves. Near resonance the undamped steady particular amplitude is inappropriate, and the initial transient or a physical damping model matters.

Calculate work done by a moving wall in an elastic collision

A particle of mass m moves along a line and collides elastically with a wall whose prescribed laboratory velocity is U during the short contact. In the instantaneous wall frame the incoming relative velocity $v^- - U$ reverses, so

$$v^+ = 2U - v^-.$$

In the laboratory the impulse is $J = m(v^+ - v^-) = 2m(U - v^-)$ and the kinetic energy change is $\Delta K = m[(v^+)^2 - (v^-)^2]/2 = UJ$. The energy gained by the particle equals wall velocity times the impulse delivered by the wall. A perfectly elastic contact still changes laboratory kinetic energy when the wall moves, since elasticity was defined for the relative normal velocity.

Suppose a particle of mass $0.20\,\mathrm{kg}$ moves right at $3.0\,\mathrm{m/s}$ into a right wall moving left at $0.50\,\mathrm{m/s}$. Its final velocity is $-4.0\,\mathrm{m/s}$, and the impulse on the particle is $-1.40\,\mathrm{N\,s}$. The particle's kinetic energy rises from 0.90 to $1.60\,\mathrm{J}$, exactly $UJ = 0.70\,\mathrm{J}$. The actuator keeping the wall on schedule supplies the energy. When the same wall recedes to the right at $0.50\,\mathrm{m/s}$, the outgoing velocity is $-2.0\,\mathrm{m/s}$ and kinetic energy falls by $0.50\,\mathrm{J}$. The difference is not symmetric under reversal of wall speed at fixed incoming particle velocity, because that reversal also changes the relative approach speed.

Now trap a particle between a fixed wall and a length $l(t)$ varying slowly compared with the collision period. To leading order a right-wall collision changes the particle's speed by $-2\dot{l}$, and such collisions occur about once in $2l/v$. Hence $dv/dt \simeq -(v/l)\dot{l}$, so vl is approximately constant. The one-dimensional action around a full bounce cycle is

$$J_{\mathrm{act}} = (2\pi)^{-1}(2mvl) = mvl/\pi.$$

This elementary wall argument explains the adiabatic action without requiring the smooth-force approximation at the instantaneous collision. If the wall speed varies quickly within a round trip, the collision phases matter and the averaged invariant can fail.

Starting with $l_0 = 0.40\,\mathrm{m}$ and speed $4.0\,\mathrm{m/s}$, slow compression to $0.30\,\mathrm{m}$ predicts $v \simeq 5.33\,\mathrm{m/s}$. Its kinetic energy increases by a factor $16/9$, supplied through the moving boundary. A pressure sensor at the wall measures impulse per time, $2mv$ times the collision frequency $v/(2l)$, giving average force mv^2/l at leading order. Multiplying this force by the inward speed gives the power required by the predicted energy law. An instantaneous compression while the particle is elsewhere would instead change no kinetic energy until a later collision. The final length alone does not determine the energy unless the rate protocol is specified.

Follow a string through a moving boundary in configuration and time

For a vertical moving pivot with height $Y(t)$, coordinates $x = l\sin\theta$, $y = Y - l\cos\theta$ give kinetic energy $m[\dot{Y}^2 + 2l\dot{Y}\dot{\theta}\sin\theta + l^2\dot{\theta}^2]/2$. Derive the angular equation by differentiating $p_\theta = ml^2\dot{\theta} + ml\dot{Y}\sin\theta$. The terms containing $\dot{Y}\dot{\theta}\cos\theta$ cancel as in horizontal translation; the remaining equation is

$$l\ddot{\theta} = -(g + \ddot{Y})\sin\theta.$$

A vertical acceleration changes the restoring coefficient multiplicatively, unlike a horizontal displacement whose small-angle effect is an additive forcing term. The distinction is visible even when the two pivot motions have identical amplitudes and frequencies.

For $Y = A\cos\Omega t$, the effective gravity oscillates as $g - A\Omega^2\cos\Omega t$. The perfectly vertical solution remains an exact solution at any drive amplitude. A small deviation need not remain small, because its linear equation has a periodic coefficient. An observed growth of swing from a tiny initial disturbance near twice the natural frequency is a parametric effect. A camera viewing the pivot and bob together can separate relative angle from large imposed vertical displacement. Estimating angular amplitude from absolute bob height alone introduces a component at the drive frequency unrelated to the swing.

At modest drive $A = 1.0\,\mathrm{mm}$, $\Omega = 20\,\mathrm{rad/s}$, the pivot acceleration amplitude is $0.40\,\mathrm{m/s^2}$, about four percent of gravity. A pendulum of length $0.10\,\mathrm{m}$ has small-angle frequency $9.90\,\mathrm{rad/s}$, putting this drive near twice its natural frequency despite the small displacement. For an exactly stationary vertical bob relative to the pivot, the string tension is $m(g + \ddot{Y})$ as long as that value is positive. A force measurement thus contains a direct drive-frequency signal even before transverse oscillation develops. Combining an angle sensor with a load cell helps distinguish contact loss from growth of parametric oscillation.

A fast vertical drive can also modify the slow stability of the inverted position. A quantitative treatment requires separating a fast small angular ripple from the slow tilt and averaging its feedback, with the support acceleration and frequency satisfying the relevant ordering. The instantaneous sign of $g + \ddot{Y}$ alone is insufficient to establish the stability of the inverted solution. An upright motion that briefly samples a favourable sign each cycle may still drift away over many cycles. Compare initial perturbations at several drive phases and fit the envelope over complete periods before drawing a stability conclusion.

Treat the drive itself as a degree of freedom

A prescribed rotating slot conceals the energy source behind a fixed pattern speed. Replace the motor by a free rotating turntable with moment of inertia I_0, and let its angle φ evolve dynamically. In the horizontal frictionless problem the total kinetic energy is $m\dot{r}^2/2 + (I_0 + mr^2)\dot{\varphi}^2/2$. The angle is cyclic, so the total angular momentum $P = (I_0 + mr^2)\dot{\varphi}$ is constant. As the bead moves out, the turntable slows. The radial equation remains $\ddot{r} = r\dot{\varphi}^2$, but its coefficient now decreases according to the total angular momentum instead of remaining externally fixed.

Eliminate $\dot{\varphi}$ from the total energy to find

$$E = m\dot{r}^2/2 + P^2/[2(I_0 + mr^2)].$$

The reduced radial potential decreases outwards and feeds radial kinetic energy, but the total energy available is finite. Differentiating the potential gives a radial acceleration $P^2 r/(I_0 + mr^2)^2$, consistent with

the original radial equation. One should not simply insert $\dot{\varphi} = P/(I_0 + mr^2)$ into the original kinetic Lagrangian before varying r; the correct fixed-momentum Routh reduction subtracts $P\dot{\varphi}$, reversing the sign of the eliminated rotational term. The energy and Lagrangian methods agree once their controlled quantities are distinguished.

With $I_0 = 0.020\,\mathrm{kg\,m^2}$, $m = 0.10\,\mathrm{kg}$, $r_0 = 0.10\,\mathrm{m}$, initial angular speed $4.0\,\mathrm{rad/s}$ and zero radial speed, total momentum is $0.084\,\mathrm{kg\,m^2/s}$ and energy $0.168\,\mathrm{J}$. At $r = 0.30\,\mathrm{m}$ the rotation rate has fallen to $0.084/0.029 = 2.90\,\mathrm{rad/s}$. The remaining radial kinetic energy is $0.168 - 0.084^2/(2 \cdot 0.029) = 0.0463\,\mathrm{J}$, giving radial speed $0.962\,\mathrm{m/s}$. Under a prescribed constant angular speed of four, the radial first integral instead gives $\dot{r} = 4\sqrt{r^2 - r_0^2} = 1.13\,\mathrm{m/s}$ there. The two runs started with identical positions and velocities but have different power connections.

The drive can also be modelled between those limits. Let a servo apply torque $\tau_{\mathrm{drive}} = K_{\mathrm{s}}(\Omega_{\mathrm{command}} - \dot{\varphi})$. Then $d[(I_0 + mr^2)\dot{\varphi}]/dt = \tau_{\mathrm{drive}}$, and the rate of increase of total mechanical energy is $\tau_{\mathrm{drive}}\dot{\varphi}$. A high control gain with finite response time can approximate the prescribed-speed case over part of a trajectory, yet a rapid radial change can cause a measurable speed sag. Recording shaft torque and speed together tests the drive rather than attributing its lag to a fictitious radial friction on the bead.

Infer drive work from a closed loop of length and measured tension

The reel-pendulum experiment produces synchronised records of string force and length. Over a period in which the length returns to its original value, the work on the bob from the reel is $-\oint T\,dl$. A graph of tension against length can enclose a signed area although the string force at each instant is perpendicular to the bob's allowed virtual angular displacement. That loop area measures energy transferred by the imposed control. If the bob returns to its full mechanical state and there are no other energy sources, any net positive reel work must go into losses. If only the string length returns, the bob's energy can increase cycle after cycle.

Suppose the reel shortens by $20\,\mathrm{mm}$ during a high-tension interval averaging $7.0\,\mathrm{N}$, then pays out the same length while the mean tension is $4.5\,\mathrm{N}$. The net mechanical work on the bob is $0.020(7.0 - 4.5) = 0.050\,\mathrm{J}$ for that cycle. A torque sensor on a reel of changing winding radius must convert torque to tension with the actual instantaneous radius. A force signal delayed relative to the length encoder can create a spurious enclosed area, especially when tension changes rapidly near a turning point. Repeating the measurement slowly with a known hanging calibration weight provides a direct check of polarity and timing.

Derive rapid-pivot stabilisation rather than replacing gravity by its mean

The vertical-pivot equation can be written $\ddot{\theta} + [g/l - (A\Omega^2/l)\cos\Omega t]\sin\theta = 0$. The time average of the oscillatory coefficient is zero. Nevertheless a high-frequency pivot can alter slow stability because the small fast ripple in θ correlates with the fast acceleration. Let $\theta = \Theta + \xi$, where Θ varies little in a drive period and ξ is small. Keeping the leading rapidly oscillating term gives $\ddot{\xi} \simeq (A\Omega^2/l)\sin\Theta\cos\Omega t$, hence $\xi \simeq -(A/l)\sin\Theta\cos\Omega t$. Its amplitude scales with displacement ratio A/l, so large acceleration and small angular ripple can coexist.

Expand $\sin(\Theta+\xi)\simeq\sin\Theta+\xi\cos\Theta$, insert the fast solution, and average. Terms linear in the cosine average to zero but the product of ripple and imposed acceleration gives a mean cosine square of one half. To leading order the slow equation is

$$\ddot{\Theta}+\frac{g}{l}\sin\Theta+\frac{A^2\Omega^2}{2l^2}\sin\Theta\cos\Theta=0.$$

It derives from an effective potential per angular inertia $U_{\mathrm{eff}}/(ml^2)=-(g/l)\cos\Theta+[A^2\Omega^2/(4l^2)]\sin^2\Theta$, up to a constant. Linearising near $\Theta=\pi+\eta$ gives

$$\ddot{\eta}+[A^2\Omega^2/(2l^2)-g/l]\eta=0.$$

The inverted state is stable within the averaged approximation if $A^2\Omega^2>2gl$. The restoring contribution depends on squared pivot speed amplitude. Replacing effective gravity by its simple time average would miss it entirely.

Choose a rigid pendulum arm with $l=0.20\,\mathrm{m}$, drive amplitude $A=0.0040\,\mathrm{m}$, and $\Omega=600\,\mathrm{s}^{-1}$. The small displacement ratio is 0.020, the pivot speed amplitude is $2.4\,\mathrm{m/s}$, and $A^2\Omega^2=5.76\,\mathrm{m}^2/\mathrm{s}^2$, above $2gl=3.92\,\mathrm{m}^2/\mathrm{s}^2$. The predicted slow angular frequency about the inverted state is $\sqrt{5.76/(2\cdot0.20^2)-9.81/0.20}=4.79\,\mathrm{s}^{-1}$, far below the drive frequency. A slow ripple of tilt $0.10\,\mathrm{rad}$ has a leading drive-frequency micromotion of roughly $0.0020\,\mathrm{rad}$. A camera sampling at an unfortunate phase can hide this micromotion or alias it into apparent slow drift.

The rigid arm in this example is essential. An ordinary string cannot push its bob when an averaged upright configuration requires alternating compression and tension throughout a fast cycle. Check the actual instantaneous radial reaction and the physical actuator's ability to deliver the large acceleration. Also check drive-period, arm elasticity and support resonance before identifying an observed stable wobble with the averaged rigid-arm prediction. The leading formula assumes small ripple and a clear separation between Ω and all slow angular frequencies; near threshold the slow restoring coefficient is a small difference and higher-order effects or friction become relatively important.

Separate support work from energy in a moving coordinate

Prescribe an accelerating support. Compare the relative response with what a laboratory observer sees and identify how the drive supplies energy.

Open this experiment online

Try the calculation

Support acceleration (m/s²) 1.6

Spring stiffness (N/m) 8

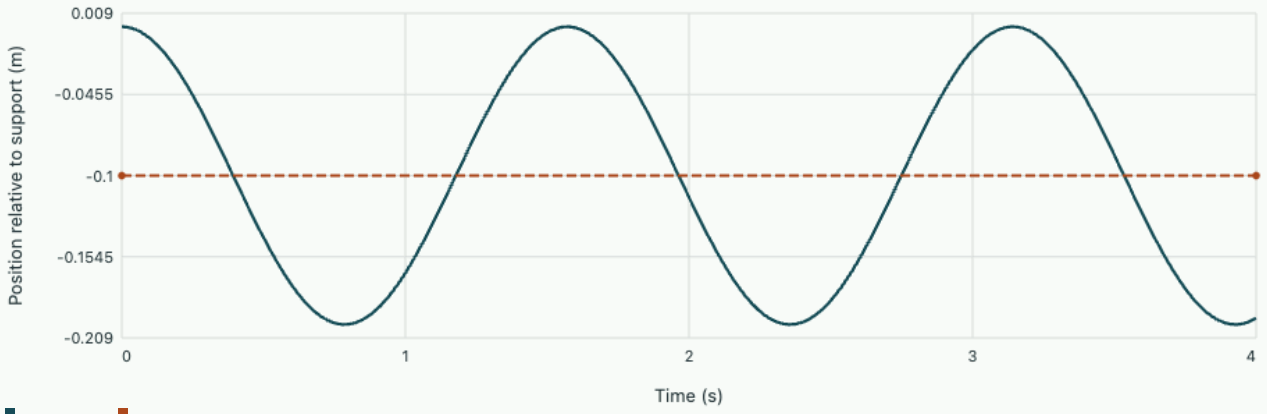

Bead Accelerated equilibrium

The equilibrium is at -0.1 m. The first turning point is -0.2 m at 0.7854 s.

Distinguish two ways of moving an endpoint in Hamilton's principle

When a constraint itself is explicitly time dependent, the identity $\delta g = \nabla g \cdot \delta x = 0$ describes a virtual change at fixed time on the moving surface. A variational problem may instead allow the endpoint time to change. If a terminal point moves along a prescribed spacetime curve $q_b(t)$, the endpoint variation includes $dq_b = \dot{q}_b\, dt$; it is not an arbitrary tangent displacement at fixed laboratory clock. The boundary term in the on-shell action is $p\, dq - H\, dt$. Reading these two terms separately tells which physical terminal conditions a derivation has assumed.

For a free particle of mass m moving in one dimension from the origin at time zero to q at time T, the extremal action is $S = mq^2/(2T)$. Its derivatives are $S_q = mq/T = p$ and $S_T = -mq^2/(2T^2) = -H$. If the terminal marker moves as $q_b(T) = uT$, the action of the connecting extremal is $mu^2T/2$. Its derivative along that endpoint line is $pu - H = mu^2/2$, exactly the direct derivative. Holding the endpoint position fixed while differentiating instead gives $-mu^2/2$. Both derivative experiments are legitimate, but the different signs demonstrate why one must state what is fixed.

Can a numerical action fit identify momentum and energy from a single sequence of terminal markers all moving at speed u?

▼ **Work through the solution**

Along that one curve it returns the combination $pu - H$. Separating momentum and energy requires independent endpoint variations in position and time, or another dynamical relation. For the free-particle model the known quadratic Hamiltonian supplies such a relation. For an unknown effective system a single endpoint track can leave different combinations indistinguishable. A grid of nearby fixed-time endpoints estimates p, and a separate grid of terminal times at fixed place estimates $-H$. This measurement

interpretation is the finite-dimensional counterpart of separating virtual reaction work from actual boundary power.

Infer whether a moving wall may be treated as prescribed

In an actual impact apparatus the wall has finite effective mass M. For a one-dimensional elastic impact with an initially moving wall, conservation of total momentum and reversal of relative velocity give

$$v^+ = [(m - M)v^- + 2MU]/(m + M),$$
$$U^+ = [2mv^- + (M - m)U]/(m + M).$$

Only the limit $M/m \to \infty$, or active feedback able to enforce the schedule on the collision time scale, gives the prescribed-wall expression used earlier. A motor that restores its target speed a second later may still behave as a finite inertia during the collision itself.

Take $m = 0.20\,\mathrm{kg}$, $M = 2.0\,\mathrm{kg}$, $v^- = 3.0\,\mathrm{m/s}$ and $U = -0.50\,\mathrm{m/s}$. The finite wall sends the particle back at approximately $-3.36\,\mathrm{m/s}$ and recoils to $+0.136\,\mathrm{m/s}$, compared with the prescribed-wall prediction of $-4.0\,\mathrm{m/s}$. Check that the initial and final total momenta both equal about $-0.40\,\mathrm{kg\,m/s}$. Measuring wall speed immediately before and after contact separates an imperfectly enforced boundary from an inelastic impact; both can reduce the magnitude of particle rebound relative to an ideal prescribed boundary.

Tutorial on a prescribed moving apparatus

The rotating slot is held at fixed Ω, but the bead is initially given inward speed $-\Omega r_0$. Predict its subsequent motion and whether the motor delivers or receives power.

▼ Work through the solution

The general solution is a combination of growing and decaying exponentials. These initial conditions select $r = r_0 e^{-\Omega t}$. The wall reaction $2m\Omega\dot{r}$ is negative, and its power $2m\Omega^2 r\dot{r}$ is also negative. The motor receives energy from the bead while maintaining its speed. The same geometric constraint can therefore inject or absorb energy depending on the actual trajectory. The origin is approached only asymptotically in the ideal solution.

At $l = 0.80\,\mathrm{m}$ and $\theta = 0$, a bob has zero angular speed and the prescribed reel tries to accelerate the length outward at $12\,\mathrm{m/s^2}$. Can a string enforce this motion?

▼ Work through the solution

The radial-force equation gives $T = m(g - 12)$, a negative tension. A string cannot push the bob outward. The reel pays out slack and the bob initially accelerates under gravity, so treating $l(t)$ as its exact radius becomes inadmissible. A rigid telescopic arm capable of compression would impose a different physical constraint and could enforce the same geometrical schedule.

A pivot is displaced sinusoidally by $2.0\,\mathrm{mm}$ at $1.0\,\mathrm{Hz}$ with $l = 0.25\,\mathrm{m}$. Is the low-frequency tilt formula suitable?

▼ **Work through the solution**

The natural angular frequency is $\sqrt{9.81/0.25} = 6.26\,\mathrm{rad/s}$, almost equal to the drive angular frequency $2\pi = 6.28\,\mathrm{rad/s}$. The low-frequency expansion requires a frequency much smaller than the natural one and fails here. Measure damping and initial conditions, or model a finite-time drive. A steady frictionless amplitude calculated from the small difference of squared frequencies is extremely sensitive to uncertainties in the length and drive frequency.

Compare a second rotating constraint and identify what the motor holds fixed.

David Tong's Lagrangian dynamics chapter develops explicitly time-dependent coordinate changes.

Heat stored in quantum degrees of freedom

Build the single-particle spectrum before taking the thermal average. Compare acoustic branches, magnetic ordering and paired electrons, and trace how a radiometer samples a photon reservoir.

1. Vibrations of monoatomic and diatomic chains
2. Filling electronic states and measuring the Fermi surface
3. Magnetic moments and susceptibility
4. Exchange, spontaneous magnetisation and critical response
5. Pairing, excitation gaps and what a spectrum measures
6. How the same electrons carry heat and charge
7. Why a polar crystal has two optical phonon frequencies
8. Exchange and two-particle energies
9. Density matrices and loss of phase
10. Thermal radiance, étendue and detector calibration
11. Why an oscillator can move without its packet spreading
12. From a Coulomb orbit to the fine-structure multiplet
13. When a magnetic field changes the quantum numbers of an atom
14. Two electrons, exchange splitting and the cost of localisation
15. From a finite pulse to a continuum transition rate
16. Measure tunnelling from a four-site spectrum

Vibrations of monoatomic and diatomic chains

Condensed matter and band theory

Derive normal modes from a local force law and identify the physical motion behind the acoustic and optical branches.

FROM FIRST PRINCIPLES

Turn a local spring law into a phonon dispersion relation

Consider identical atoms of mass m spaced by a along a chain. Write u_j for the displacement from the equilibrium position of atom j. For small changes of bond length, expand the bond energy to second order about its minimum. Its first derivative vanishes and its positive second derivative K is the measured bond stiffness. The bond on the right pulls with force $K(u_{j+1}-u_j)$; the left bond contributes $K(u_{j-1}-u_j)$. Newton's law gives

$$m\ddot{u}_j = K(u_{j+1} - 2u_j + u_{j-1}).$$

Translation of a normal mode by one site can change its phase but not its shape. Insert $u_j = Ue^{i(kja-\omega t)}$, divide by the common amplitude, and simplify:

$$-m\omega^2 = K(e^{ika} + e^{-ika} - 2) = -4K\sin^2(ka/2),$$

$$\omega = 2\sqrt{K/m}\,|\sin(ka/2)|.$$

For wavelengths much larger than a,

$$\sin(ka/2) \simeq ka/2.$$

The slope gives the sound speed $a\sqrt{K/m}$. At $k=0$ all atoms move together and no bond stretches, which explains the zero frequency. More than one atom per primitive cell permits internal relative motion even at $k=0$, the ingredient needed for an optical branch.

DERIVATION AND APPLICATIONS

Count the modes that crowd against a chain's band edge

Why can a narrow frequency bin near the top of a vibration band contain a substantial share of all modes? Take a monoatomic chain of N equal masses m with one longitudinal displacement each. Impose periodic ends and use identical harmonic nearest-neighbour springs of stiffness K with spacing a. The frequency on the positive half of the first zone rises monotonically to $\omega_m = 2\sqrt{K/m}$. Inverting that branch gives the wavevector bounding the states below a trial frequency,

$$k(\omega) = \frac{2}{a}\arcsin\!\left(\frac{\omega}{\omega_m}\right), \qquad 0 \le \omega \le \omega_m.$$

Allowed wavevectors are spaced by $2\pi/(Na)$. Counting the interval of width $2k(\omega)$ includes both propagation directions. In a large chain the fractional count and its derivative per atom are therefore

$$F(\omega) = \frac{2k(\omega)}{2\pi/(Na)}\frac{1}{N} = \frac{2}{\pi}\arcsin\left(\frac{\omega}{\omega_m}\right),$$

$$g(\omega) = \frac{dF}{d\omega} = \frac{2}{\pi\omega_m\sqrt{1-(\omega/\omega_m)^2}}.$$

Here $g\,d\omega$ is a fraction of the modes. The normalization follows from $F(\omega_m) = 1$, as required by the count of independent displacements. Near the edge the slope of the dispersion tends to zero, so evenly spaced wavevectors map into progressively smaller frequency intervals. This causes an integrable singularity in the continuum density. A finite spectrometer bin always contains a finite count. A measured spectrum can also weight modes differently through a probe's coupling to atomic motion.

WORKED EXAMPLE

Count modes within the upper tenth of the band

Change the inputs and investigate this derivation.

For $N = 10{,}000$, the interval between $0.9\omega_m$ and the upper edge contains $N[1-(2/\pi)\arcsin(0.9)] \simeq 2871$ modes, allowing a mode or two for endpoints. A uniform distribution in frequency would give 1000. A nearly stationary wave packet near the upper edge and a large number of modes per unit frequency have the same origin in the small group velocity.

SECOND APPLICATION

Use the spectral median to infer a spring stiffness

Change the inputs and investigate this derivation.

An ideal chain with atomic mass $5.00 \times 10^{-26}\,\mathrm{kg}$ has half of its longitudinal modes below an ordinary frequency of $4.00\,\mathrm{THz}$. Setting $F = 1/2$ gives

$$\omega_{1/2}/\omega_m = \sin(\pi/4).$$

The upper frequency is $5.66\,\mathrm{THz}$, and $K = m(2\pi f_m)^2/4 = 15.8\,\mathrm{N/m}$. The frequency ratio is unchanged by conversion to angular frequency, but the final stiffness requires that conversion.

Try it · Chain count

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 fraction of band-edge frequency **0.9**

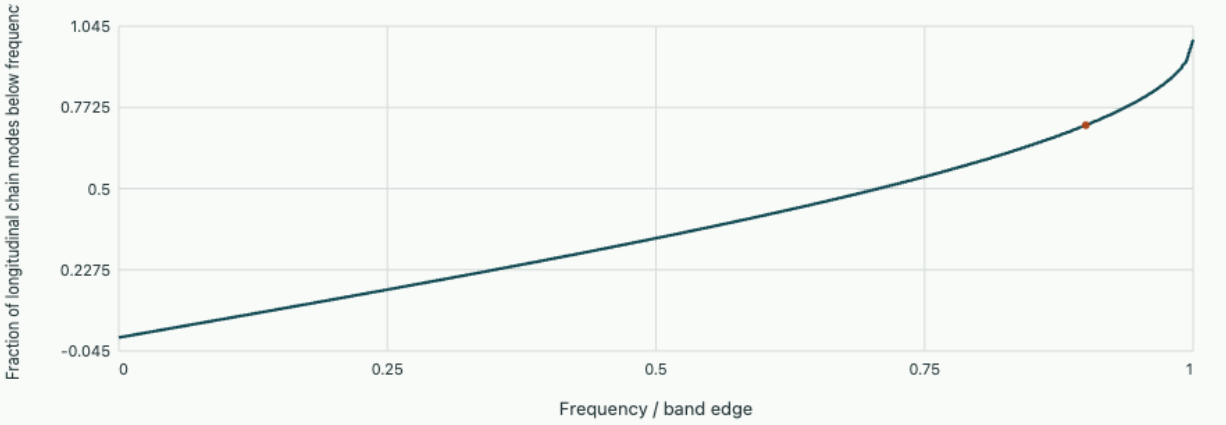

Uniform wavevector count Selected cutoff

Fraction above selected cutoff 0.28713. The spectral median is ωmax/√2; mode density diverges integrably at the band edge.

From bond extension to a travelling displacement

Let atoms of mass m on a line be coupled by springs of stiffness K, spaced a apart. The spring to the right pulls site n by $K(u_{n+1}-u_n)$; the left contributes $K(u_{n-1}-u_n)$. A uniform shift produces no force. Translation symmetry suggests $u_n = u e^{i(kna-\omega t)}$. Substitution gives

$$\omega^2(k) = \frac{2K}{m}(1-\cos ka) = \frac{4K}{m}\sin^2(ka/2).$$

The sound speed at long wavelength is $a\sqrt{K/m}$, consistent with a continuum rod whose one-dimensional stiffness is Ka and whose mass per length is m/a. At $k=\pi/a$ neighbouring atoms move in opposition and the slope of the dispersion vanishes. A wave packet made of nearby k's moves with the group velocity $d\omega/dk$, which can differ greatly from the phase velocity ω/k near the boundary.

Background · Why must one vibration have exactly zero frequency?

Give every atom the same displacement. If the potential energy depends only on separations, no bond changes length and no atom experiences a restoring force. Differentiate the force with respect to this uniform translation. For one displacement direction the resulting row-sum condition on the force constants is

$$\sum_j K_{ij} = 0,$$

$$D\,(\sqrt{m_1}, \ldots, \sqrt{m_N})^T = 0,$$

$$D = M^{-1/2} K M^{-1/2}.$$

With unequal masses the translational eigenvector is uniform in physical displacement, while its mass-weighted components vary. A nonzero frequency at zero wavevector may reveal missing neighbour interactions in a numerical model. A real spring to a substrate also breaks translation symmetry and produces a physical gap. The boundary conditions decide which explanation is plausible.

Introduce an internal coordinate

Place masses m, M alternately, with spacing $a/2$ between neighbours and cell period a. Let their displacements in cell n be Ue^{ikna} and Ve^{ikna} times the same time factor. The two force balances contain sums of neighbouring amplitudes. For example the A mass sees B in cell n and in cell $n-1$, giving $m\omega^2 U = 2KU - K(1 + e^{-ika})V$. Exchanging the masses and conjugating the phase gives the second row. Requiring a nonzero solution yields

$$mM\omega^4 - 2K(m+M)\omega^2 + 4K^2 \sin^2(ka/2) = 0.$$

There are two nonnegative solutions for ω^2 per k. With N cells under periodic boundary conditions this gives $2N$ longitudinal modes, exactly matching the number of independent displacements. At $k = 0$ the lower branch has $U = V$ and zero frequency. For the upper branch the centre of mass stays fixed, $mU + MV = 0$, with $\omega^2 = 2K(1/m + 1/M)$. An electric field can excite some optical modes in polar crystals; the branch name alone does not guarantee optical activity.

A gap and a sound speed answer different questions

At the zone boundary $1 + e^{ika} = 0$, so the sublattices decouple. If $M = 4m$, the two edge frequencies are $\sqrt{K/(2m)}$ and $\sqrt{2K/m}$, a factor of two apart. When the masses become identical the two folded branches meet, recovering the original monoatomic chain in a doubled cell. An apparent branch boundary created by plotting in a smaller zone need not be a physical frequency gap.

Expand the determinant for small k on the lower branch. Since ω^2 scales as k^2, discard the term of order k^4. This gives

$$v_s = a\sqrt{K/[2(m+M)]}.$$

The two springs in a cell act in series for long waves, while both masses contribute to inertia. The sound speed alone determines a stiffness-to-mass ratio; measuring the upper branch supplies independent information about the internal motion.

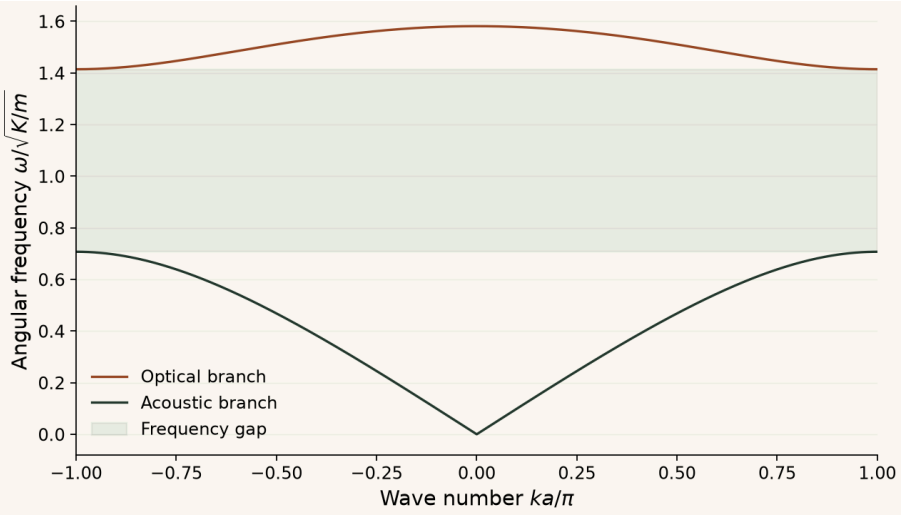

For a four-to-one mass ratio the zone-edge frequencies differ by a factor of two. The shaded interval contains no longitudinal normal mode of the infinite chain.

Try it · Diatomic

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

Heavy / light mass 4

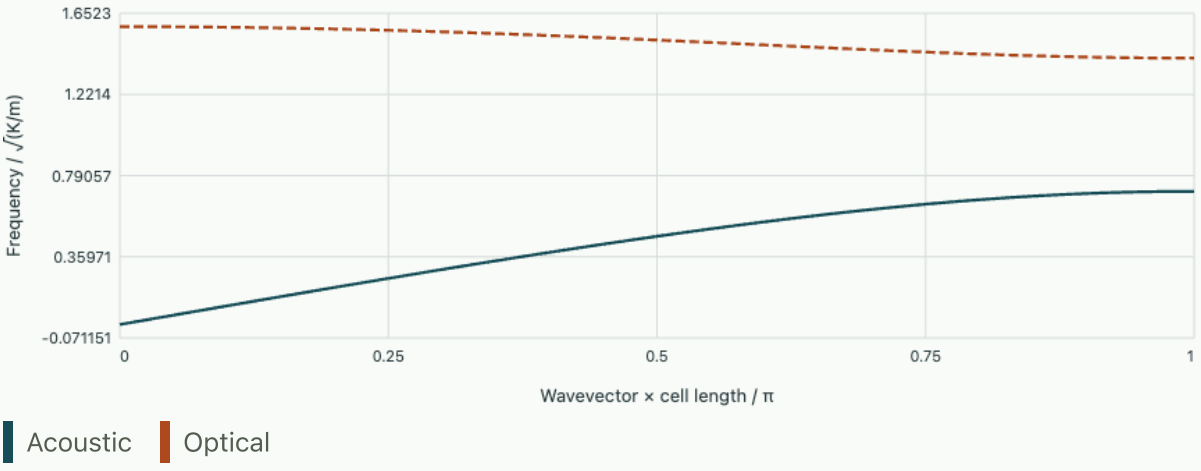

Zone-edge frequencies 0.70711 and 1.4142. The gap closes at equal masses.

Convert a mode count into thermal predictions

Quantisation assigns a harmonic mode levels $\hbar\omega(j+1/2)$. The thermal part of its mean energy is $\hbar\omega/[\exp(\hbar\omega/k_BT)-1]$. Low-temperature excitations thus select long-wavelength acoustic modes while high-frequency optical modes remain scarcely populated. Counting states in a three-dimensional sphere gives a low-frequency density proportional to ω^2 for linear branches. Rescaling the energy integral by $\hbar\omega/k_BT$ gives thermal energy proportional to T^4 and heat capacity proportional to T^3. Dimensionality and linear dispersion are essential assumptions behind that power.

Count thermally generated vacancies

A mode calculation for a perfect crystal starts with one atom on every site. Heating a real sample can leave some atoms at a surface or reservoir and some lattice sites empty. How large should that defect population be even at thermal equilibrium? Consider N equivalent sites and n_v independent vacancies. At fixed temperature and pressure let g_f be the Gibbs cost of making one vacancy and transferring its atom to the

reservoir. It may contain local vibrational contributions. We count the separate entropy of choosing the missing sites explicitly.

There are $W = N!/[n_v!(N-n_v)!]$ such arrangements. With $x = n_v/N$ and macroscopic occupations, Stirling's formula gives

$$\frac{G - G_0}{N} = x g_f + k_B T\,[x \ln x + (1-x)\ln(1-x)].$$

Minimizing with respect to x gives

$$g_f + k_B T \ln[x/(1-x)] = 0.$$

The second derivative is positive. Thus $x = 1/(1 + e^{g_f/(k_B T)})$. For positive $g_f \gg k_B T$ there are few vacancies and $x \simeq e^{-g_f/(k_B T)}$. Each vacancy costs energy, but at tiny x the number of possible positions grows fast enough to favour a nonzero population. A finite microscopic specimen can still contain zero vacancies at a given instant.

Take $g_f = 0.800\,\mathrm{eV}$ as constant over the chosen range and $k_B = 8.6173 \times 10^{-5}\,\mathrm{eV/K}$. At $900\,\mathrm{K}$, $k_B T = 0.07756\,\mathrm{eV}$ and $x = 3.31 \times 10^{-5}$. A region with 10^{15} sites contains an expected 3.31×10^{10} empty sites. At $600\,\mathrm{K}$ the equilibrium fraction drops to 1.91×10^{-7}. Rapid cooling may leave an excess because vacancies must migrate to sinks before the concentration relaxes. The formation free energy in this calculation sets an equilibrium population. A migration barrier would enter the time to attain it.

A light substitution can localize a vibration

An infinite monatomic chain has a continuous vibration band, while a single altered atom can support a vibration spatially concentrated near one site. Let identical springs of stiffness K connect masses m on a chain of spacing a, except the central mass at site zero is M. Seek a symmetric localized solution with alternating neighbouring amplitudes $u_n = A(-q)^{|n|}e^{-i\omega t}$ for $0 < q < 1$. Alternation is motivated by the top of the perfect-chain band where adjacent atoms oscillate out of phase. Exponential decay supplies the localization.

At any ordinary site away from the impurity, the equation $-m\omega^2 u_n = K(u_{n+1} + u_{n-1} - 2u_n)$ gives

$$m\omega^2 = K(2 + q + q^{-1}).$$

Since $q + q^{-1} > 2$, this frequency lies above the perfect-chain band edge $\omega_{\max} = 2\sqrt{K/m}$. At site zero both neighbours have displacement $-qu_0$, giving $M\omega^2 = 2K(1+q)$. Eliminating ω and cancelling $1+q$ yields

$$q = \frac{M}{2m - M},$$

$$\omega^2 = \frac{K}{m}(2 + q + q^{-1}).$$

The physical condition $0 < q < 1$ requires a lighter central mass $0 < M < m$. A heavy isotope can still scatter travelling waves or produce resonant response within the band. This ansatz predicts no exponentially

localized symmetric mode above the band for the heavier substitution. Its absence follows from the admissible decay factor rather than an assumption that a heavy impurity cannot move.

Predict a localized spectral peak

Take host atomic mass $m = 4.00 \times 10^{-26}$ kg, spring stiffness $K = 40.0$ N/m, lattice spacing $a = 0.300$ nm and impurity $M = m/2$. The host upper angular frequency is

$$2\sqrt{40.0/(4.00 \times 10^{-26})} = 6.32 \times 10^{13}\,\mathrm{rad/s}.$$

For the impurity $q = 1/3$, so $\omega^2 = (16/3)K/m$, giving $\omega = 7.30 \times 10^{13}$ rad/s or $f = 11.6$ THz. The ordinary chain ends near 10.1 THz.

The neighbour amplitude is minus one third of the central amplitude, and the second neighbour has plus one ninth. Expressing the absolute envelope as $e^{-|n|a/\xi}$ gives localization length $\xi = a/\ln 3 = 0.273$ nm. As M approaches m from below, q approaches one, localization length diverges and the extra frequency merges with the band edge. A finite-chain calculation should therefore vary chain length before claiming that a very weak impurity mode is spatially localized: a box shorter than its decay length cannot separate that broad tail from a standing host mode.

Spatially resolved displacement data can provide two independent estimates of impurity mass. Invert the decay-factor relation to obtain $M/m = 2q/(1+q)$ from neighbouring displacement amplitudes. Separately insert the measured localized frequency and the same host force constant in the central equation of motion. Agreement tests the assumption that the impurity changes only inertia. A substituted atom may also change the two neighbouring springs: in that case the observed frequency and decay factor can no longer be fitted by altering mass alone. This inversion distinguishes an isotope-like substitution from a stronger local bonding perturbation.

A weak second impurity far away perturbs the localized mode only through its exponentially small tail. Bringing two such impurities within a few decay lengths allows the modes to mix and split in frequency, offering another test of their spatial localization.

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

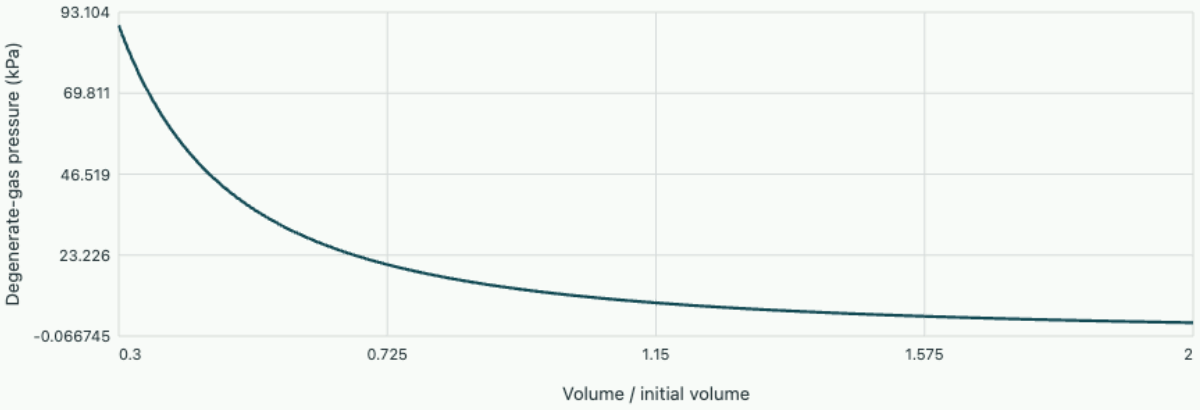

Zero-temperature 3D ideal Fermi gas

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)
```

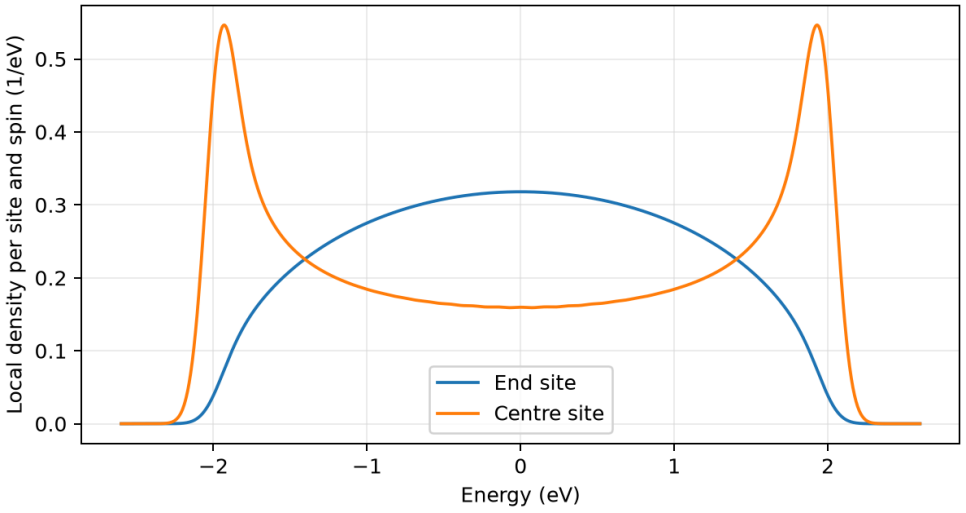

The integrated areas agree. Near zero energy the end site carries about 0.318 states/eV per spin and the centre carries 0.160. A change in probe location redistributes spectral weight.

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**

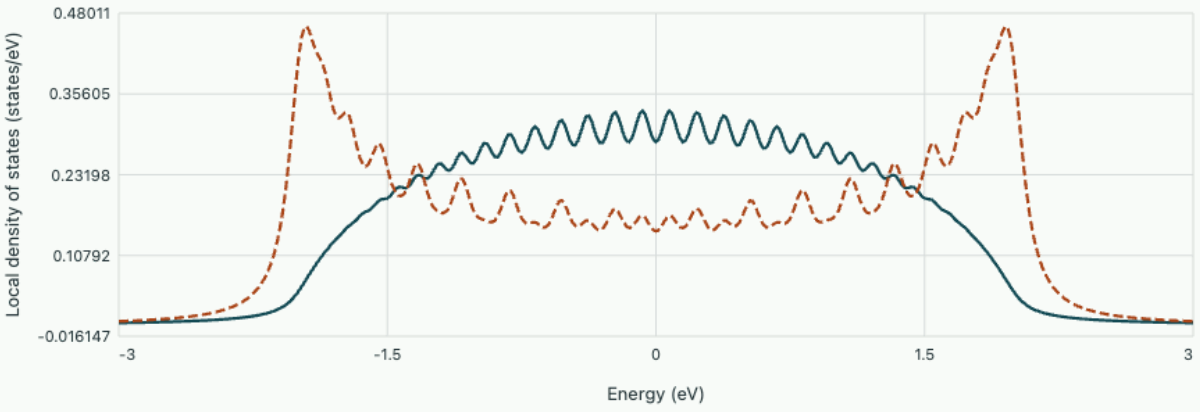

Selected site Central site

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.

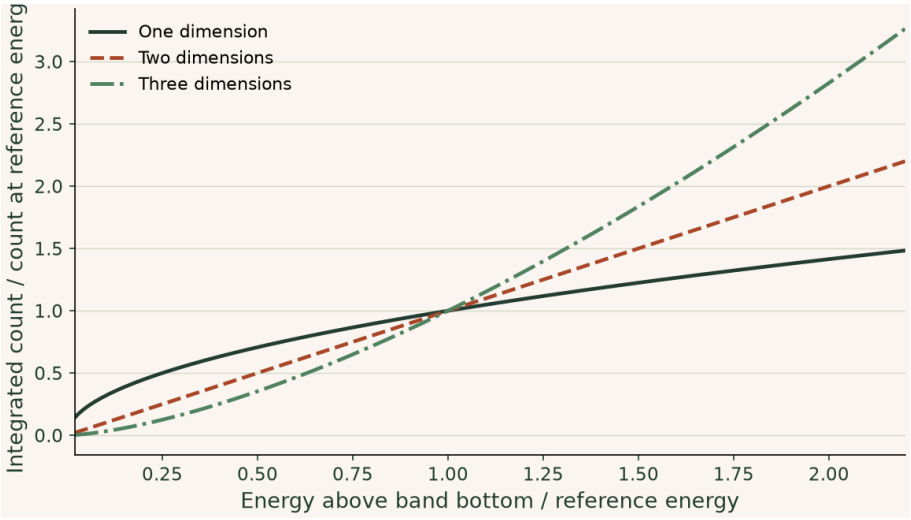

Doubling the Fermi energy raises the count by √2, 2 or 2√2. The local slope gives the density of states. Its behaviour at the band edge is quite different even when the three integrated counts agree at the reference energy.

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

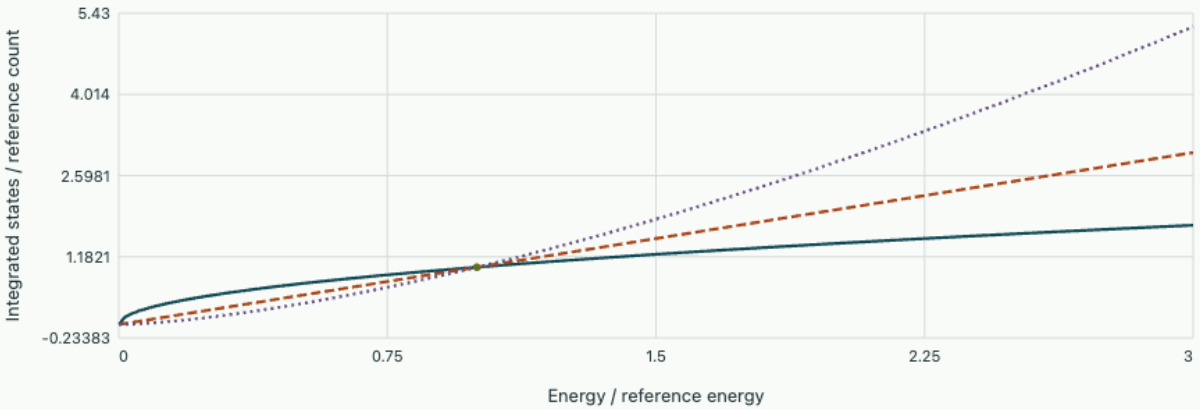

1 spatial dimensions | 2 spatial dimensions | 3 spatial dimensions | Selected 3D count

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**

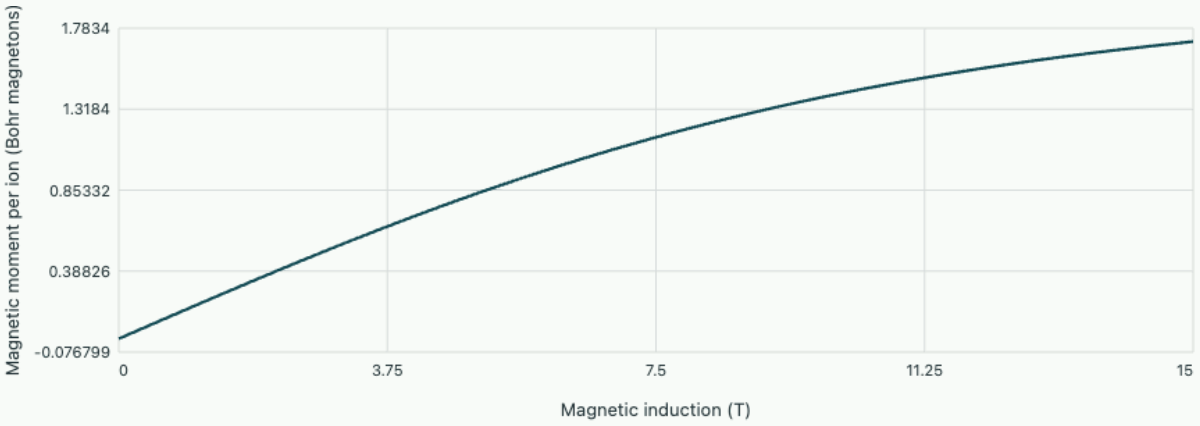

Exact spin partition sum, g=2

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.

Exchange, spontaneous magnetisation and critical response

Condensed matter and band theory

Use a mean-field model and a free-energy expansion to explain how a macroscopic moment can arise in zero applied field.

FROM FIRST PRINCIPLES

How an interaction permits spontaneous magnetization

Put an Ising spin $s_i = \pm 1$ on each lattice site and use a ferromagnetic bond energy $-Js_is_j$ with $J > 0$. If every site has z neighbours, replacing their actual spins by the common mean $m = \langle s \rangle$ assigns a trial spin energy $-s(zJm + \mu B)$. Thermal averaging this two-level problem must reproduce the same mean used to construct the local field:

$$m = \frac{\sum_{s=\pm 1} s\, e^{s(zJm+\mu B)/(k_BT)}}{\sum_{s=\pm 1} e^{s(zJm+\mu B)/(k_BT)}} = \tanh\!\left(\frac{zJm + \mu B}{k_BT}\right).$$

At zero applied field the right-hand side is odd, so $m = 0$ always solves the equation. To decide when two nearby nonzero solutions appear, compare the slope at the origin with one, then retain the cubic correction. Write $a = zJ/(k_BT)$:

$$m = am - \tfrac{1}{3}a^3m^3 + \cdots,$$
$$T_c^{\mathrm{MF}} = zJ/k_B, \qquad m^2 \simeq 3(a-1)/a^3.$$

Just below this mean-field critical temperature either sign of m is possible even in zero applied field. The approximation replaces correlated neighbours by an average. In short-range one-dimensional systems that omission can even change whether a finite-temperature transition exists.

DERIVATION AND APPLICATIONS

Calculate a correlation length without a mean-field transition

Mean-field theory replaces the neighbours of each spin by a common average. A chain gives a case in which the energy of a local reversal can instead be counted exactly. Take an open chain of Ising spins $s_i = \pm 1$, zero external field, and ferromagnetic nearest-neighbour interaction $J > 0$. The energy is $-J\sum_i s_is_{i+1}$. Specify the first spin and all bond variables $b_i = s_is_{i+1}$. This is a one-to-one description of every configuration. Since the energy is a sum over these independent binary bonds, each bond has weights e^K and e^{-K}, where $K = J/(k_BT)$.

$$p_{\mathrm{wall}} = P(b_i = -1) = \frac{1}{1 + e^{2K}},$$
$$\langle b_i \rangle = 1 - 2p_{\mathrm{wall}} = \tanh K.$$

A wall costs $2J$ compared with an aligned bond. To compare two spins separated by r bonds, multiply the intervening b_i. Every interior spin then occurs twice and drops out. Bond independence makes the average of this product particularly simple,

$$\langle s_i s_{i+r} \rangle = \left\langle \prod_{j=i}^{i+r-1} b_j \right\rangle = (\tanh K)^r .$$

Writing the result as an exponential in the physical separation ra, where a is the site spacing, gives

$$\xi = -\frac{a}{\ln(\tanh K)},$$

$$\xi \simeq \frac{a}{2} e^{2K} \quad (K \gg 1).$$

At low temperature the last expression follows by expanding $\tanh K \simeq 1 - 2e^{-2K}$. At every positive temperature the probability of a wall remains positive. Thus widely separated spins lose correlation in this short-range one-dimensional model. A finite sample shorter than ξ can still look ordered during a measurement.

WORKED EXAMPLE

Estimate decorrelation over a microscope field

Change the inputs and investigate this derivation.

Let $J/k_B = 20.0\,\mathrm{K}$, $T = 10.0\,\mathrm{K}$ and $a = 0.400\,\mathrm{nm}$. The wall probability is $1/(1+e^{4}) = 0.0180$, giving about 18 walls per 1000 bonds on average. The exact correlation length is $10.9\,\mathrm{nm}$. Across $20.0\,\mathrm{nm}$, or 50 bonds, the measured spin correlation should be approximately $(\tanh 2)^{50} = 0.160$.

SECOND APPLICATION

Infer exchange from a measured spatial decay

Change the inputs and investigate this derivation.

In a different chain the correlation drops to one half over eight bonds at $T = 30.0\,\mathrm{K}$. Each step has correlation $2^{-1/8} = 0.9170$. Inverting $K = \tfrac{1}{2}\ln[(1+\langle b\rangle)/(1-\langle b\rangle)]$ gives

$$J/k_B = 47.1\,\mathrm{K}.$$

A local image counting opposite neighbours should then find a fraction $(1 - 0.9170)/2 = 0.0415$, an independent test using adjacent spins.

Try it · Ph ising ring

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 $k_\beta T/J$ **1**

Sites around the ring **32**

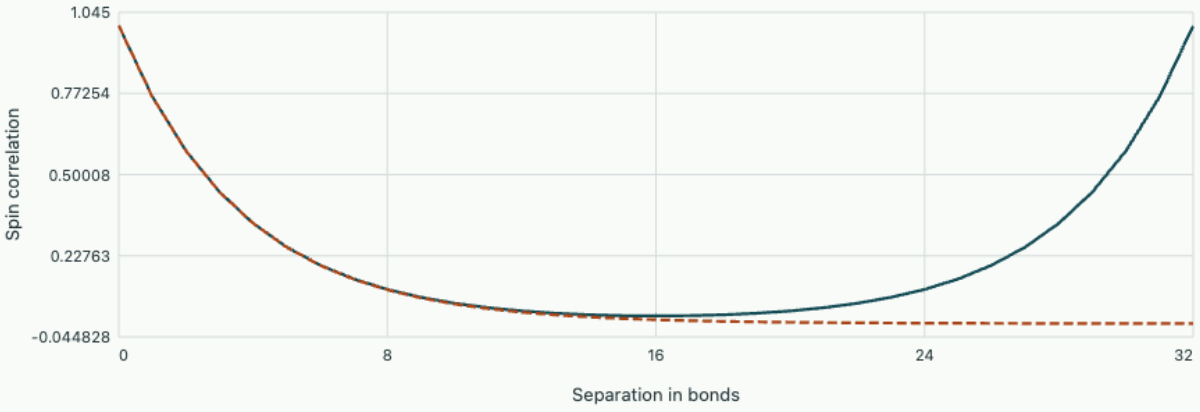

Closed finite ring Infinite chain

Infinite-chain correlation length 3.672 lattice spacings. The opposite site on the 32-site ring has correlation 0.02562.

Replace neighbouring moments by their average

An independent moment has zero mean at zero field for any positive temperature. In a ferromagnet aligned neighbours can lower the exchange energy. Consider Ising variables $s_i = \pm 1$ with energy $-J\sum_{\langle ij\rangle} s_i s_j - \mu B \sum_i s_i$, $J > 0$. Each distinct nearest-neighbour bond is counted once. In a uniform mean-field approximation replace a neighbour's spin by $m = \langle s \rangle$. For coordination number z, an individual spin then feels an effective energy bias $zJm + \mu B$. Applying the two-state result to that spin requires self-consistency:

$$m = \tanh[(zJm + \mu B)/(k_B T)].$$

At $B = 0$, $m = 0$ always solves the equation. The slope of the right-hand side at the origin is $zJ/(k_B T)$. Above the mean-field transition temperature $T_c = zJ/k_B$ only the disordered solution exists. Below it two symmetry-related nonzero stable solutions appear. The sign selected in a macroscopic region can depend on a small field and on its history.

Approach the transition from either side

For small field above T_c, expand $\tanh x \simeq x$. This gives $m \simeq \mu B/[k_B(T - T_c)]$, a Curie–Weiss susceptibility. A positive intercept of an inverse-susceptibility plot can indicate dominant ferromagnetic exchange over a suitable high-temperature interval. It is not an exact measurement of the true ordering temperature, particularly when fluctuations, geometry or competing interactions matter.

Just below T_c, at zero field, expand through the cubic term. Writing $t = T_c/T$, the nonzero solution has $m^2 \simeq 3(t - 1)/t^3$, approaching $3(T_c - T)/T_c$ sufficiently near the transition. The square-root onset is a mean-field prediction. A finite-temperature transition of this type does not exist for a one-dimensional short-range Ising chain in the thermodynamic limit, despite the same simple self-consistency calculation predicting one. Neglecting the cost and entropy of fluctuations can qualitatively change a prediction.

Interpret the curvature of a free energy

A local free energy even under $m \mapsto -m$ at zero field may be expanded as $f = \frac{1}{2}a(T - T_c)m^2 + \frac{1}{4}bm^4 - hm$, $a, b > 0$. At $h = 0$, the minima sit at 0 above the transition and at $m = \pm\sqrt{a(T_c - T)/b}$ below it. A positive second derivative indicates local stability, and the differential response $dm/dh = 1/f''$ grows as the curvature weakens. For $a(T - T_c) = -1, b = 4$ in scaled units, the minima are at ± 0.5, have curvature 2 and are lower than the intervening maximum by 0.0625 in free-energy density units.

From a local moment to a macroscopic measurement

Different regions can select opposite stable magnetisations. Domain walls cost exchange and anisotropy energy, while domain formation can reduce the magnetic field energy outside a specimen. A small measured zero-field net moment can therefore coexist with substantial local order. A static uniform free-energy graph alone cannot predict coercivity: wall pinning, nucleation, defects and measurement times affect a hysteresis loop. Temperature-dependent susceptibility, local probes and field-cycling measurements answer complementary questions about the same ordered sample.

Derive the profile and energy of a domain wall

Two domains magnetized in opposite easy-axis directions can coexist, but rotating the magnetization in between costs exchange and anisotropy energy. Consider a uniaxial ferromagnet with magnetization magnitude fixed and angle $\theta(x)$ relative to the easy axis varying along one spatial coordinate. Assume a planar 180-degree wall in a geometry where additional magnetostatic costs can be neglected for this profile. With exchange stiffness A in J/m and anisotropy density K in J/m³, the wall energy per unit area is

$$\sigma_w = \int_{-\infty}^{\infty} [A(\partial_x \theta)^2 + K \sin^2 \theta]\, dx.$$

A narrow wall saves anisotropy energy but makes the spatial gradient steep. A broad wall saves exchange energy but places more material away from the preferred axis. Variation gives

$$2A\theta'' = 2K \sin\theta \cos\theta.$$

Multiply by θ' and integrate once. Since both gradient and anisotropy term vanish deep inside either domain, the integration constant is zero, yielding $A(\theta')^2 = K\sin^2\theta$. Exchange and anisotropy costs are locally equal on the minimum-energy profile.

Take increasing θ between zero and π and define $\Delta = \sqrt{A/K}$. Separation gives

$$d\theta/\sin\theta = dx/\Delta.$$

With the wall centred at zero,

$$\tan(\theta/2) = \exp(x/\Delta).$$

Inserting the first integral back into the wall energy and changing variable from x to θ gives

$$\sigma_w = 2K\Delta\int_0^{\pi}\sin\theta\,d\theta = 4\sqrt{AK}.$$

The profile has long tails, so a quoted wall width must state its convention; $\pi\Delta$ is a frequently used slope-based width rather than a sharp boundary.

Compare two anisotropies at fixed exchange stiffness

Use $A = 1.00\times10^{-11}$ J/m and $K = 1.00\times10^{5}$ J/m³. The scale Δ is 10.0 nm and conventional width $\pi\Delta$ is 31.4 nm. Wall energy is 4.00 mJ/m². A planar patch of area 1.00 µm² costs 4.00×10^{-15} J, even though the two far domains have the same zero-field bulk energy.

If a second specimen at the same temperature has four times the anisotropy and the same exchange stiffness, the predicted width halves while the cost per area doubles. The combination of width and energy could therefore independently estimate A and K, whereas either measurement by itself fixes only a ratio or a product. In the limit $K\to0$ the preferred axis loses meaning and the wall spreads indefinitely in this model. Finite sample size, magnetostatic fields, boundaries and defects then set spatial scales omitted from the one-dimensional calculation.

One can also express the measured magnetization components explicitly. With the local moment in a fixed rotation plane, the easy-axis component is proportional to $\cos\theta = -\tanh(x/\Delta)$ and the transverse component to $\sin\theta = 1/\cosh(x/\Delta)$. The easy-axis profile passes smoothly through zero even though the moment magnitude remains fixed throughout the wall. A probe sensitive only to that component might incorrectly interpret the central zero as suppression of ferromagnetism. Measuring a transverse component or a quantity sensitive to the full moment tests the distinction between a rotating order parameter and a vanishing one.

Distinguish a soft magnetic response from a loss of stability

In zero applied field the lowest stabilising even expansion is $f = am^2/2 + m^4/4$ in these units. Vary the quadratic coefficient through zero. Equilibrium requires both stationarity and positive curvature at a minimum.

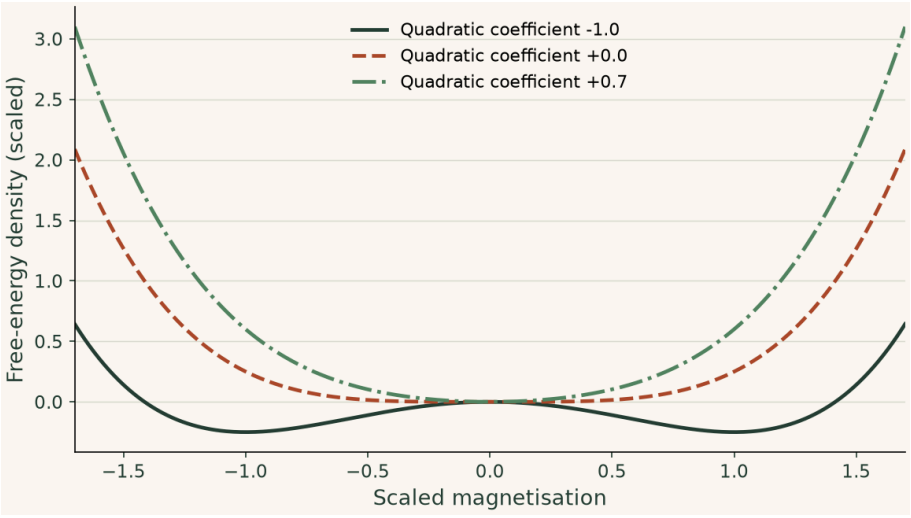

For $a = -1$ the stable minima lie at ± 1 with energy density -0.25. At $a = 0$ the quartic term stabilises a flat minimum. This bulk potential cannot by itself predict the field at which an inhomogeneous sample will reverse its magnetisation.

Try it · Landau potential

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

Quadratic coefficient a **-1**

Applied dimensionless field **0**

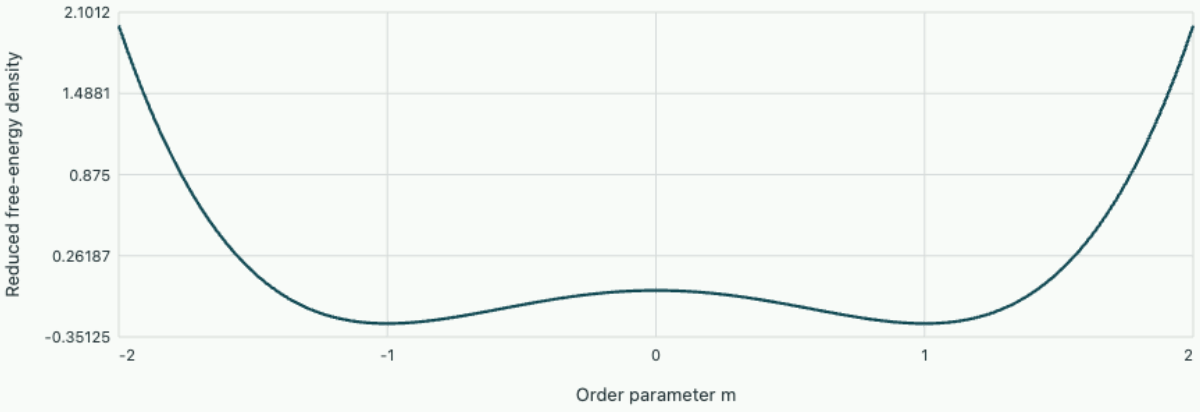

Landau potential

The quartic coefficient is one. At zero field, minima lie at ±1. A bulk potential does not predict a real specimen’s coercive field.

▶ **Read values from the plot**

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

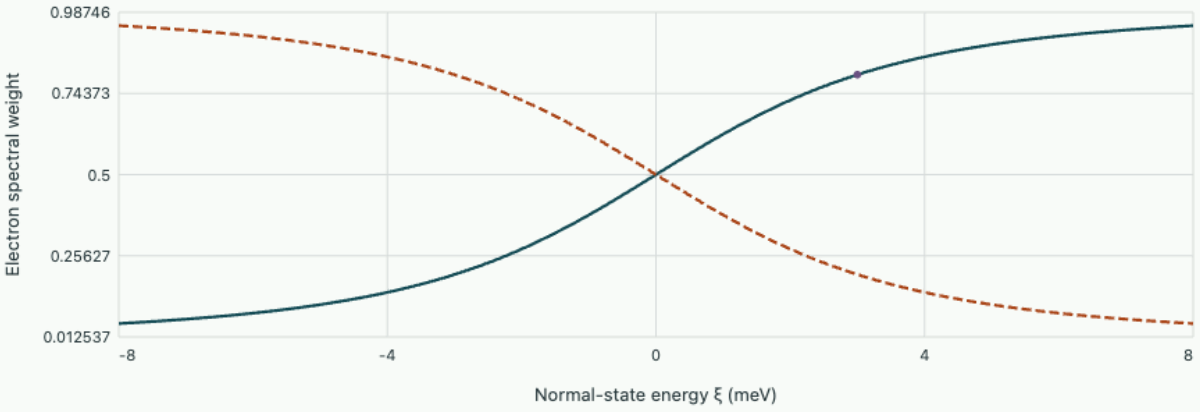

Positive-energy electron weight u² Negative-energy electron weight v²

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.

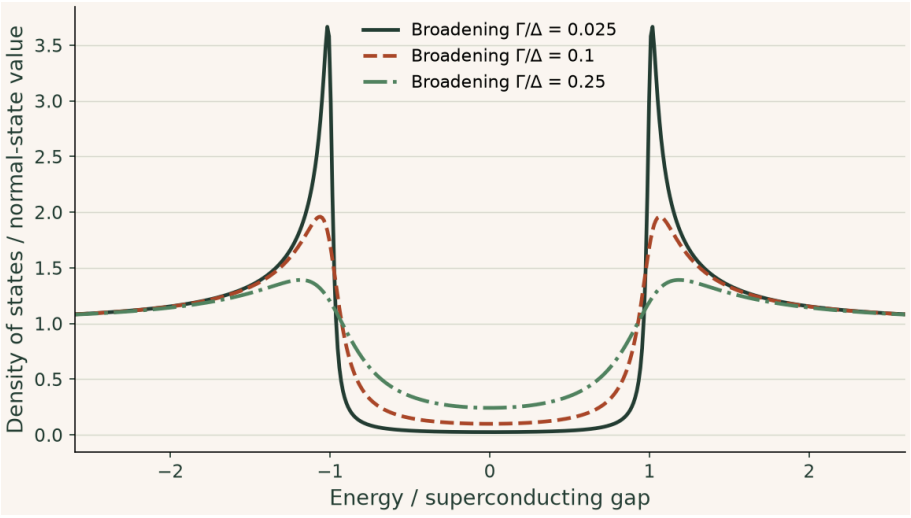

A broader peak and finite central density are possible even with the same gap parameter. At $E = 0$ the ratio is $\gamma/\sqrt{(1+\gamma^2)}$, where $\gamma = \Gamma/\Delta$. Inferring the gap from a threshold in conductance also requires the thermal and instrumental resolution.

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**

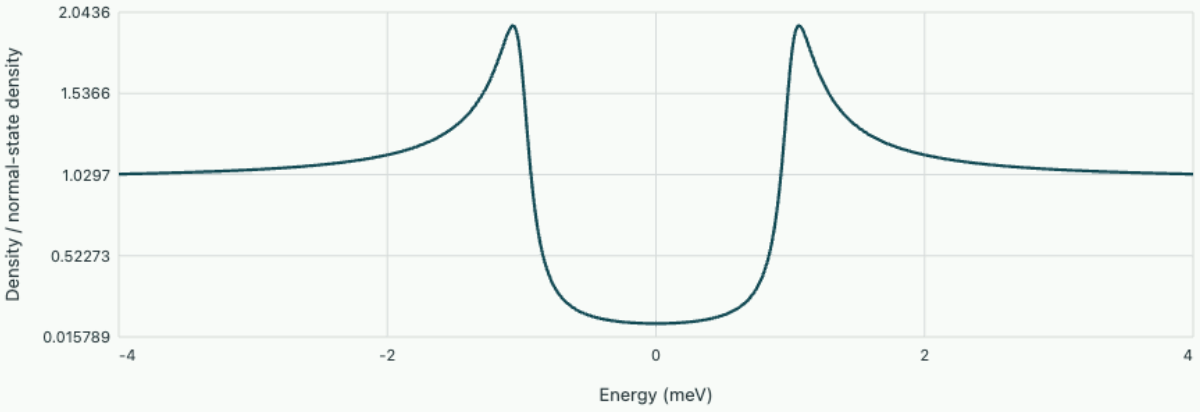

Broadened superconducting spectrum

Zero-energy density ratio 0.099504. This density-of-states model precedes thermal and instrument convolution.

▶ **Read values from the plot**

How the same electrons carry heat and charge

Condensed matter and band theory

A metal with high electrical conductivity often also conducts heat well. The quantitative relation involves an open-circuit measurement, the thermal width of the Fermi surface, and the energy dependence of the states able to carry current.

Weigh the states that can carry a current

Use an isotropic band of quasiparticles with energy E, velocity magnitude $v(E)$, density of states per volume $D(E)$ and elastic relaxation time $\tau(E)$. In a weak electric field or temperature gradient the departure from local equilibrium is proportional to the relaxation time times the applied driving term. For a given transport direction the angular average gives $\langle v_x^2 \rangle = v^2/3$ in three dimensions. Define $\Sigma(E) = D(E)v(E)^2\tau(E)/3$. Electrical current weights a particle by its charge $q = -e$, while heat current weights it by $E - \mu$. Three moments consequently suffice for the two linear responses:

$$K_j = \int \Sigma(E)(E-\mu)^j[-\partial_E f(E)]\,dE,$$

$$j = 0, 1, 2.$$

Let $\mathcal{E}$ denote the electric/electrochemical field with any chemical-potential gradient properly included. The resulting currents along one direction are $j_e = q^2K_0\mathcal{E} - qK_1\nabla T/T$ and $j_Q = qK_1\mathcal{E} - K_2\nabla T/T$. The factor $E - \mu$ in the thermal drive follows from differentiating the Fermi function with respect to temperature. The repeated first moment is a manifestation of reciprocal coupling of electrical and thermal transport in this zero-field model.

Impose the actual electrical boundary condition

A voltage probe draws negligible charge current. Set $j_e = 0$ to obtain $\mathcal{E} = S\nabla T$ with $S = K_1/(qTK_0)$. Substitute this field into the heat current to get

$$\kappa_e = (K_2 - K_1^2/K_0)/T.$$

Under the positive measure $\Sigma(-f')dE/K_0$, the bracket divided by K_0 is the variance of $E - \mu$. Hence the open-circuit electronic thermal conductivity is nonnegative. Measuring with zero imposed electric field would retain a thermally driven charge current and a different heat-flow coefficient.

For a degenerate metal with smooth Σ, the weight $-f'$ is a narrow even peak. Replacing Σ by $\Sigma(\mu)$ gives $K_1 = 0$, $K_0 = \Sigma(\mu)$, and $K_2 = (\pi^2/3)(k_BT)^2\Sigma(\mu)$. The dimensionless second moment follows by integrating $x^2/[4\cosh^2(x/2)]$ over the real axis. With $\sigma = q^2K_0$,

$$\frac{\kappa_e}{\sigma T} = L_0 = \frac{\pi^2 k_B^2}{3e^2} = 2.44 \times 10^{-8}\,\mathrm{W\,\Omega/K^2}.$$

At $T = 100\,\mathrm{K}$, a metal with $\sigma = 2.0 \times 10^7\,\mathrm{S/m}$ has estimated electronic $\kappa_e = 48.8\,\mathrm{W/(m\,K)}$. A total measured conductivity of 60 would leave about $11\,\mathrm{W/(m\,K)}$ for other carriers if the assumptions hold. Phonons, inelastic scattering or strong energy dependence may change that inference.

Select a narrow range of transport energies

A model energy filter $\Sigma \propto \exp[-(x-c)^2/(2w^2)]$, with $x = (E-\mu)/(k_B T)$, c, w dimensionless, introduces a tunable departure from this metallic limit. The Seebeck coefficient is $(k_B/q)\langle x \rangle$ and the normalised Lorenz ratio is $\mathrm{var}(x)/(\pi^2/3)$. A very narrow filter centred above the chemical potential produces strong signed thermopower but little heat transport at zero charge current, because current-carrying particles have almost identical excess energies. Those statements refer to the ratio and thermal moments; the absolute conductivity also depends on the overall transmission.

The experiment integrates these moments numerically and plots their normalised weight against the reference Fermi derivative. Vary c from positive to negative at fixed width. Predict which observable reverses sign and which is invariant before reading the numerical results.

Open this experiment online

Try the calculation

Filter centre (E – μ)/(k$_\beta$T) 1

Filter width / (k$_\beta$T) 2

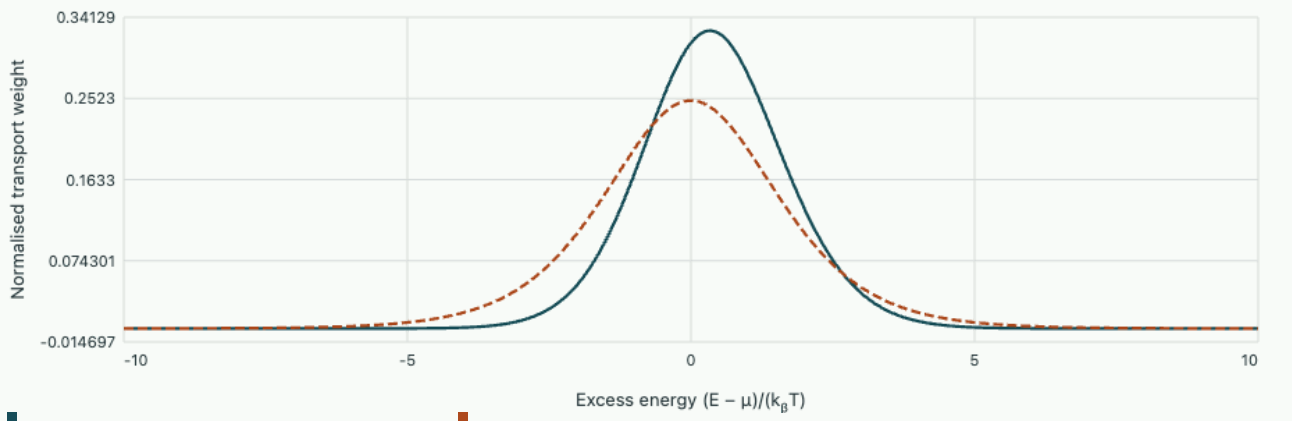

Mean excess energy 0.4018 k$_\beta$T; electron Seebeck coefficient -34.63 μV/K. Open-circuit Lorenz ratio L/L$_0$ = 0.4937.

Tutorial on open-circuit transport

For a fitted transport weight with $\langle x \rangle = 0.50$, $\langle x^2 \rangle = 2.25$, determine the Seebeck coefficient and L/L_0. Use $k_B/e = 86.17\,\mu\mathrm{V/K}$.

▼ **Work through the solution**

For electrons $S = -43.1\,\mu\mathrm{V/K}$. The variance is

$$2.25 - 0.50^2 = 2.00.$$

Hence

$$L/L_0 = 2.00/(\pi^2/3) = 0.608.$$

Using the second moment without subtracting the square of the first would represent the wrong electrical boundary condition.

Expand a smooth transport function to first order about μ. What controls the low-temperature sign of S?

▼ **Work through the solution**

The constant term contributes nothing to K_1 by symmetry. The slope gives

$$K_1 \simeq \Sigma'(\mu)(\pi^2/3)(k_B T)^2.$$

Dividing by qTK_0 yields

$$S \simeq [\pi^2 k_B^2 T/(3q)]\partial_E \ln \Sigma|_\mu.$$

For electrons a rising transport function gives negative S in the stated field convention.

Further reading · ETH Zürich, electronic transport coefficients · University of Pisa, transport integrals and Wiedemann–Franz law.

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

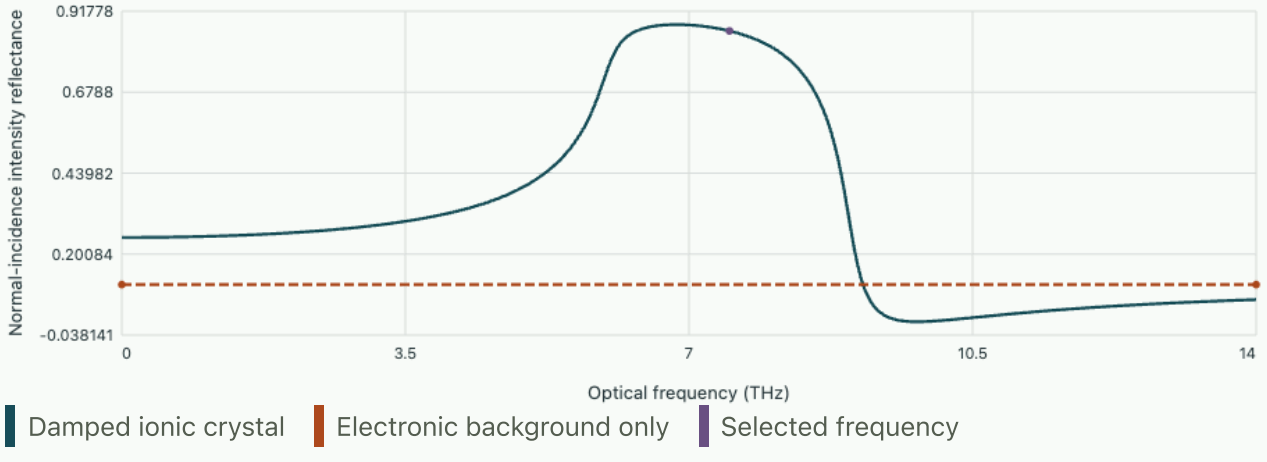

Lossless TO pole 6 THz and LO zero 9 THz. At 7.5 THz the chosen damping gives reflectance 0.8596.

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.

Exchange and two-particle energies

Quantum mechanics

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**

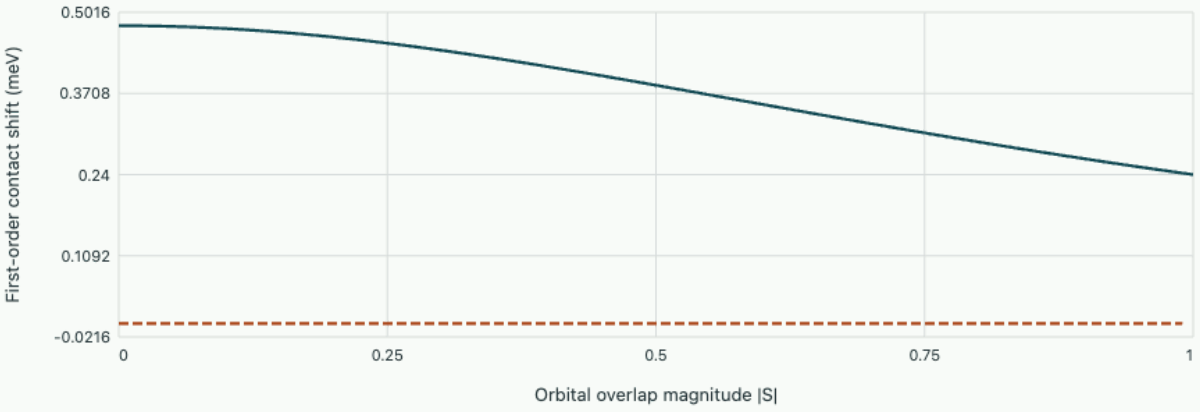

Symmetric spatial state Antisymmetric spatial state

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.

Density matrices and loss of phase

Quantum mechanics

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

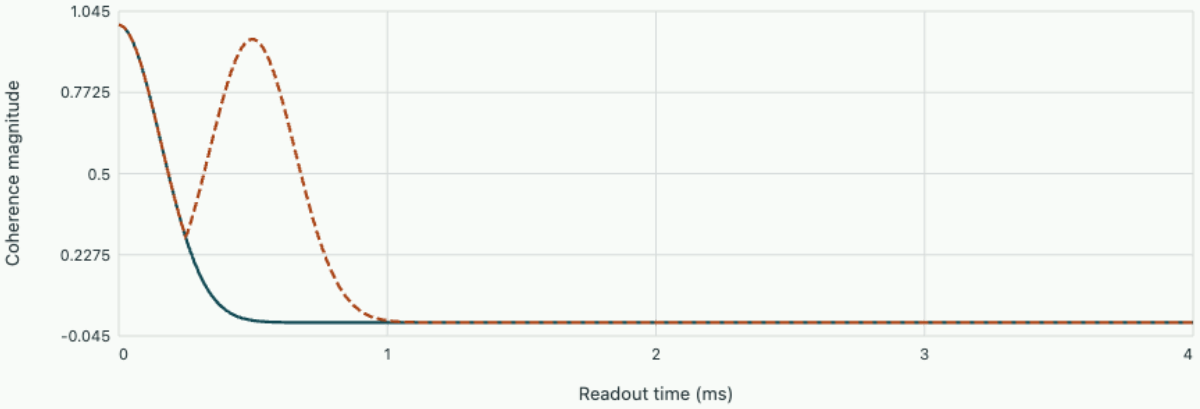

Echo centre at t=2τ=0.5 ms; its height exp(−2τ/T2)=0.95123. Static detuning reverses; homogeneous decay remains.

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.

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

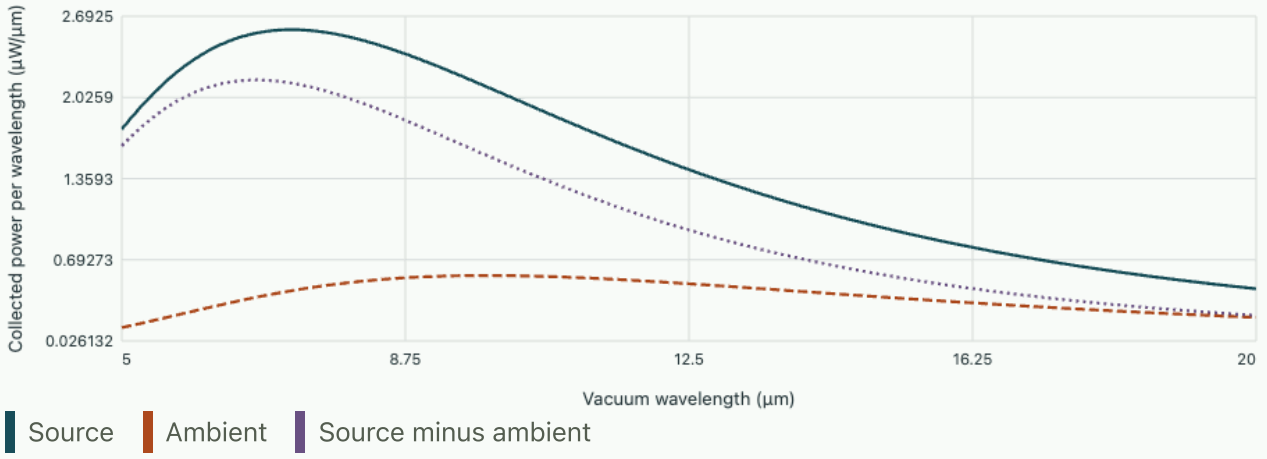

The net 8–12 µm signal is 6.053 µW for a 1 mm² area and 65% throughput. Ideal source peak wavelength is 7.24 µm.

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

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.

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 $m = 2.00\times10^{-25}$ kg and 100 kHz as above, let $F_0 = 2.00\times10^{-21}$ 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**

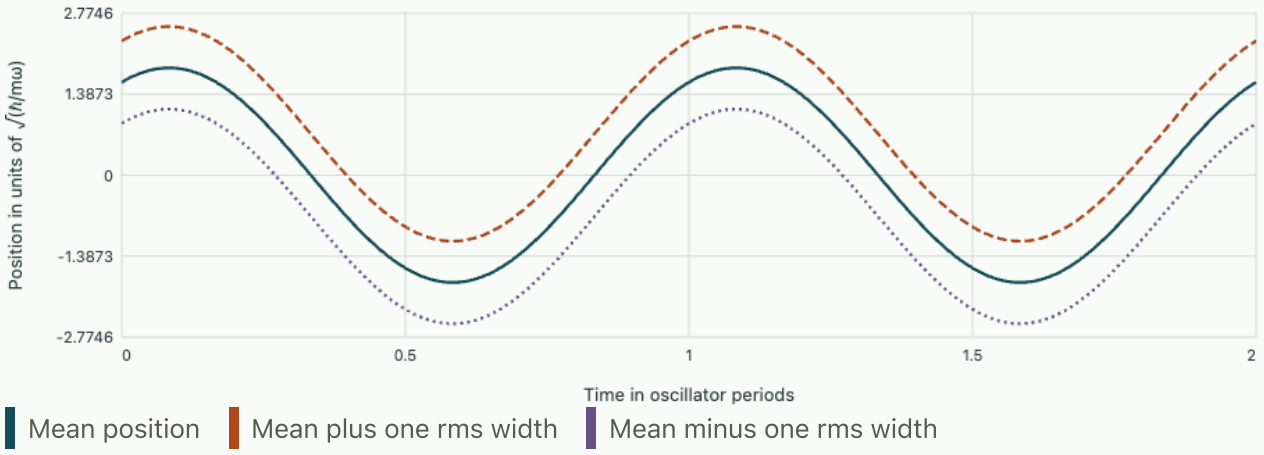

Mean occupation 1.69 quanta, standard deviation 1.3. Ground-state probability 0.18452. Position variance in the plotted units is 1/2.

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

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.

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 $\mu = m_e M/(m_e + M)$. 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 $|u(r)|^2 dr$ 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 $\langle r^{-2} \rangle = 1/[a^2 n^3 (l + 1/2)]$ 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

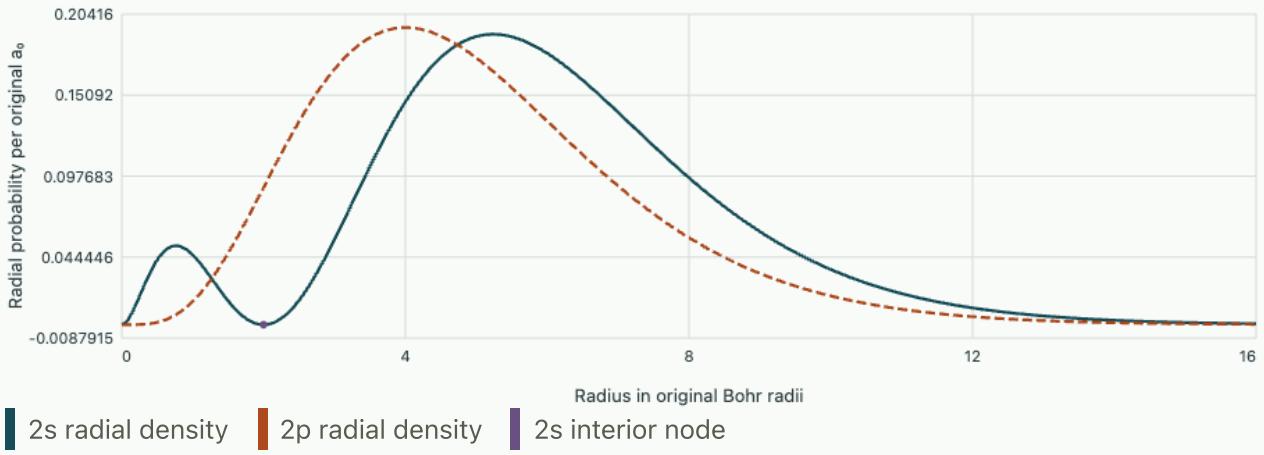

The 2s node is at 2 original Bohr radii. The most probable 2p radius is 4 original Bohr radii.

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

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.

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**

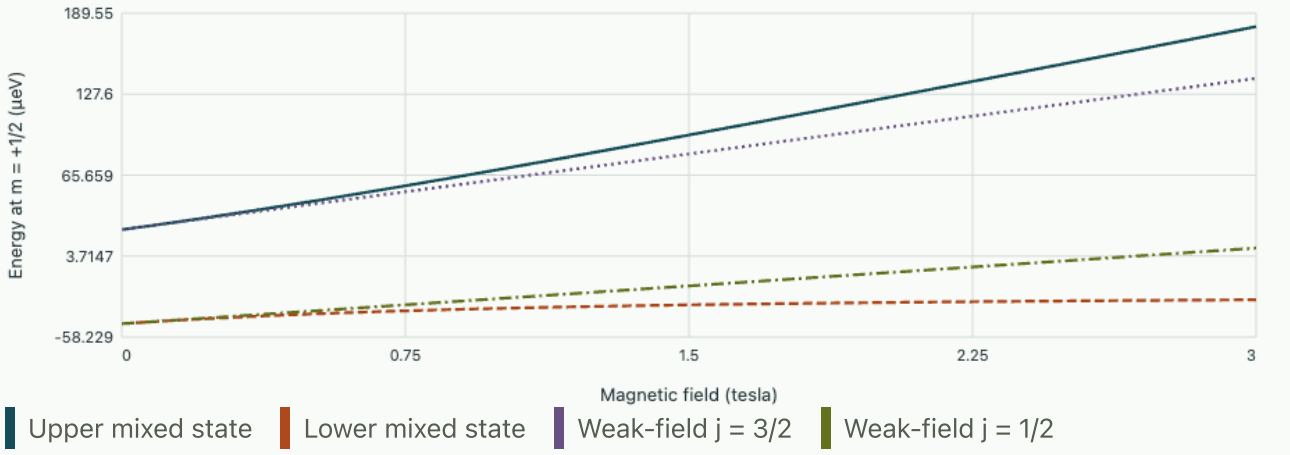

Selected energies -39.516 and 50.246 μeV. The ratio μB B/A is 0.72355.

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

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.

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

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.

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**

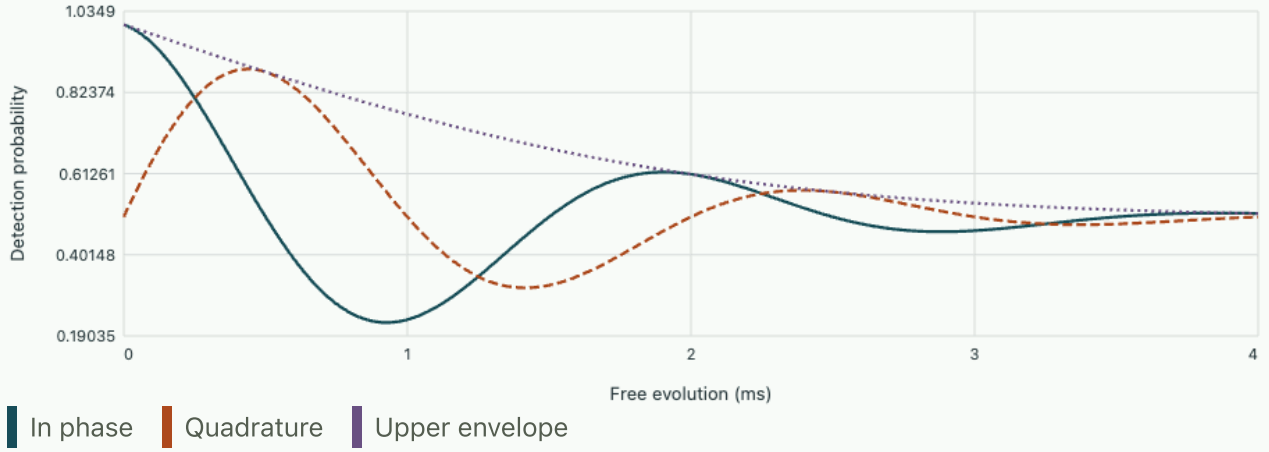

Contrast after 1 ms = 0.53455. The two probabilities at 0.50 ms are 0.5 and 0.87729.

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

Measure tunnelling from a four-site spectrum

Computational quantum experiments

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

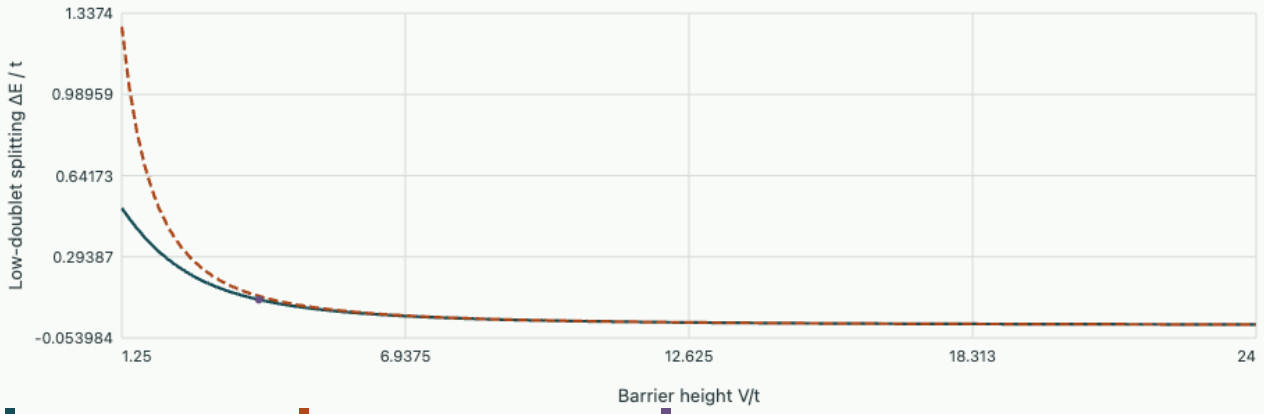

Exact parity sectors | Three-hop approximation | Selected barrier

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.

Irreversible heat transfer and fluctuating signals

Measure a transport coefficient using a steady gradient and a decay. Derive the fluctuation that survives in equilibrium and work out which part of it a finite-bandwidth instrument measures.

1. Steady conduction, thermal resistance and insulation
2. Thermal diffusion and transient surface cooling
3. Heat-exchanger effectiveness and finite capacity rates
4. When another centimetre of cooling fin stops helping
5. A moving boundary in a freezing layer
6. Diffusion with a stationary carrier gas
7. Measure how quickly a dielectric follows an electric field
8. Measure equilibrium field fluctuations with a capacitor
9. Measuring magnetic core loss by a hysteresis loop and by calorimetry
10. Optical trapping and Brownian calibration of a dielectric bead
11. How an unobserved environment changes a quantum state
12. Recover a transition frequency from two separated pulses

Steady conduction, thermal resistance and insulation

Engineering thermodynamics and heat transfer

Derive the heat leak through a cylindrical wall and identify when increasing insulation thickness can increase heat transfer.

FROM FIRST PRINCIPLES

Derive series thermal resistance from conservation of heat rate

Consider conduction along a prescribed path through a stationary material, with local area $A(x)$ normal to the heat flow and conductivity $k(x) > 0$. Assume a one-dimensional steady state, no side losses and no heat generation. A balance on an elementary slice gives a single heat rate $\dot{Q}$ at every position. Fourier's measured constitutive relation links that heat rate to the temperature gradient:

$$\frac{d\dot{Q}}{dx} = 0,$$

$$\dot{Q} = -k(x)A(x)\frac{dT}{dx}.$$

Separate variables and integrate from the hot to the cold surface. An exterior fluid film described by $\dot{Q} = hA_s(T_s - T_\infty)$ introduces an additional drop proportional to the same heat rate:

$$T_{\text{hot}} - T_\infty = \dot{Q}\left[\int_{\text{solid}} \frac{dx}{k(x)A(x)} + \frac{1}{hA_s}\right].$$

The bracket is a thermal resistance in K/W. For a plane homogeneous layer it gives $L/(kA)$; for radial conduction in a cylinder the local area $2\pi rL$ varies, giving a logarithm. Adding resistances is justified by a common heat rate. Heat-generating layers or side branches require a separate balance before using a series formula.

DERIVATION AND APPLICATIONS

How the changing area sets resistance in a spherical shell

Heat passing through the insulation of a spherical vessel spreads over a larger surface as it travels outward. Treat an inner surface at radius r_i as isothermal, a concentric outer radius r_o, constant thermal conductivity k, and uniform outside convection coefficient h. Assume steady radial conduction with no internal heat generation, perfect interfaces and a uniform ambient temperature T_∞. With no accumulation between concentric spheres, the heat rate is constant even though local heat flux decreases outward.

Fourier's empirical law specifies flux from temperature gradient. Multiply it by the local area and rearrange so that each radial slice contributes its own temperature drop:

$$\dot{Q} = -4\pi r^2 k\, dT/dr,$$

$$T_i - T_o = \frac{\dot{Q}}{4\pi k}\int_{r_i}^{r_o} \frac{dr}{r^2}.$$

The outer convective drop equals the same heat rate divided by the outer film conductance. Adding both drops gives a series relation:

$$\frac{T_i - T_\infty}{\dot{Q}} = R = \frac{1}{4\pi k}\left(\frac{1}{r_i} - \frac{1}{r_o}\right) + \frac{1}{4\pi h r_o^2}.$$

Increasing r_o adds conduction resistance but also raises film area. The sign of the derivative decides which effect dominates:

$$\frac{dR}{dr_o} = \frac{1}{4\pi k r_o^2} - \frac{2}{4\pi h r_o^3},$$

$$r_{\text{critical}} = 2k/h.$$

For an initially smaller radius, adding a thin layer can reduce total resistance; after this minimum the resistance grows. The result presumes fixed h and surface temperature T_i. If changes in geometry substantially change natural convection or radiation, the assumed film law must be revisited.

WORKED EXAMPLE

Evaluate an insulated sensor bead

Change the inputs and test this calculation.

Take $r_i = 3.00$ mm, $r_o = 6.00$ mm, $k = 0.0400$ W/(m K) and $h = 10.0$ W/(m² K). The shell contributes 331.6 K/W and the film 221.0 K/W. A 40.0 K inner-to-ambient difference gives 0.0724 W. The uncoated bead had outer resistance 884 K/W and lost only 0.0452 W. At 6 mm the coating is still below the 8.00 mm critical radius.

SECOND APPLICATION

Determine whether a change in airflow reverses the trend

Change the inputs and test this calculation.

With forced airflow raising h to 40.0 W/(m² K), the critical radius falls to 2.00 mm, inside the bare bead. The bare and coated resistances become 221 and $331.6 + 55.3 = 386.9$ K/W. The coating now reduces heat loss from 0.181 to 0.103 W. The same layer can change the sign of its effect because the surrounding film is part of the experiment.

Try it · Spherical insulation

Set the inputs to the values in the worked example, then vary one quantity while holding the others fixed. Compare the curve and numerical balance with the second application. The calculation uses the assumptions derived above; its readout identifies bounds such as choking, phase-region changes and inconsistent measurements.

Open this live calculation

Try the calculation

Insulation conductivity (W/m/K)	**0.04**
Outer convection coefficient (W/m²/K)	**10**
Selected outer radius (mm)	**6**

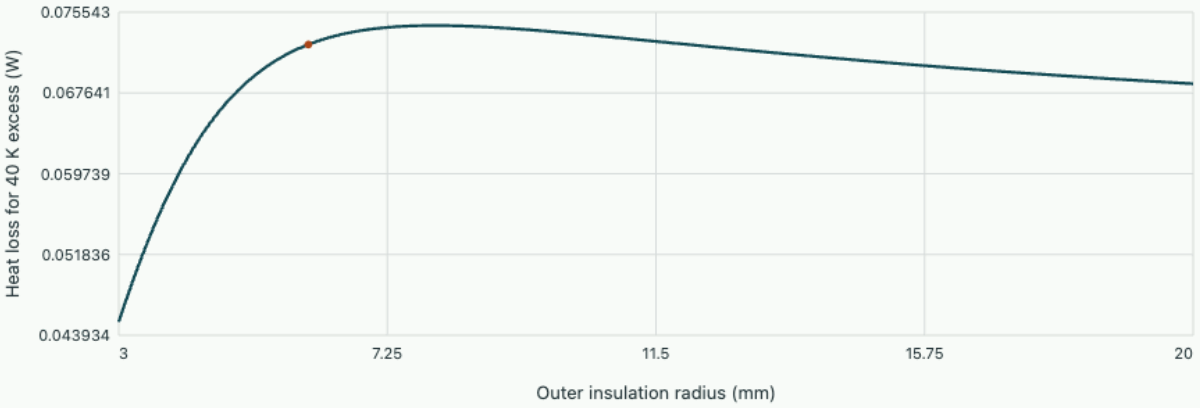

Spherical shell + surface film Selected coating

Inner radius 3 mm. Critical radius 2k/h = 8 mm. Bare loss 0.0452389 W; selected coated loss 0.0723823 W. The film coefficient and inner temperature are held fixed.

Turn heat flux into a temperature field

Energy moves through a stationary homogeneous solid down its temperature gradient. For isotropic conduction Fourier's law is

$$\mathbf{q} = -k\nabla T,$$

where $\mathbf{q}$ is heat flux in W/m² and k is thermal conductivity. A steady one-dimensional wall with no internal heat generation passes the same heat rate at every cut. Heat flux is uniform only when the area of the cuts is uniform. That distinction matters for pipes and spherical shells.

Let an inner cylindrical surface at radius r_i be warmer than the outer one at r_o. Assume a long concentric wall of length L, constant k, negligible end loss and axisymmetry. Fourier's law gives

$$\dot{Q} = -2\pi r L k \, dT/dr,$$

with outward rate taken positive. The area grows with radius, so the magnitude of dT/dr decreases outward even though the heat rate is unchanged.

Derive and combine thermal resistances

Separate and integrate over the cylindrical wall. This yields

$$T_i - T_o = \dot{Q} \ln(r_o/r_i)/(2\pi k L).$$

The cylindrical conduction resistance is the factor multiplying $\dot{Q}$. At an outer convection boundary $\dot{Q} = h\, 2\pi r_o L (T_o - T_\infty)$, giving another resistance between the surface and the bulk surroundings. For a prescribed inner surface temperature,

$$\dot{Q} = \frac{T_i - T_\infty}{\ln(r_o/r_i)/(2\pi k L) + 1/(2\pi r_o L h)}.$$

For multiple concentric layers, add the conduction resistance of each. The drops across series resistances sum because the same rate passes through them all. Poor contact between layers can add another temperature jump. A parallel heat leak bypassing the insulation carries part of the total heat under a common temperature difference, so its conductance adds instead.

Estimate the heat leak from a warm pipe

Take $L = 1.0$ m, $r_i = 0.020$ m and $r_o = 0.040$ m. The insulation has $k = 0.050$ W/(m K) and the surroundings have $h = 10$ W/(m² K), $T_\infty = 293$ K. The inner insulation surface is held at 353 K. The conduction resistance is $\ln 2/[2\pi(0.050)(1)] = 2.206$ K/W, and the outer film resistance is

$$1/[10(2\pi)(0.040)(1)] = 0.398\,\mathrm{K/W}.$$

Total heat loss is $60/(2.206 + 0.398) = 23.0$ W per metre of pipe.

The outer surface is

$$293 + (23.0)(0.398) = 302.2\,\mathrm{K}.$$

The rest of the 60 K difference is within the insulation. The heat flux at the outer surface is half that at the inner surface, because the heat-transfer area is twice as large. Assigning the outer circumference to the entire wall in a flat-wall formula would miss this variation. When thickness is very small relative to radius, expanding the logarithm recovers the plane-wall approximation.

Explain the critical radius

Increasing r_o makes the conduction resistance grow but makes the convection resistance shrink. At fixed T_i, T_∞, h, k, differentiating the total resistance gives

$$dR/dr_o = 1/(2\pi L k r_o) - 1/(2\pi L h r_o^2).$$

Its minimum occurs at $r_c = k/h$; that radius maximises heat loss. Here $r_c = 0.0050$ m, already below the bare pipe radius, so adding this insulation reduces loss from the outset. A much thinner wire could initially lose more heat because its exposed area increases faster than the wall resistance. This conclusion assumes the same external coefficient as the diameter changes. If natural convection or radiation changes appreciably, the boundary model and its temperature dependence must be reconsidered.

Recover an unknown contact resistance

A 0.010 m² test sandwich passes 8 W at steady state between measured outer face temperatures of 50 and 40 °C. Its measured total resistance is

$$10/8 = 1.25\,\mathrm{K/W}.$$

The two known solid layers together account for 0.75 K/W, leaving 0.50 K/W attributable to the interface if other parasitic paths are small. The area-normalised contact resistance is $R'' = AR = 0.0050\,\mathrm{m^2K/W}$, corresponding to conductance $1/R'' = 200\,\mathrm{W/(m^2K)}$. A pair of thermocouples on opposite sides of the interface should then show a 4 K temperature difference. That independent measurement helps distinguish contact resistance from an unaccounted external heat loss.

Derive the temperature field along a cooling fin

A fin adds area for convection but must first conduct heat from its base to that area. Treat a straight fin with constant cross-sectional area A_c, exposed perimeter P, conductivity k and uniform film coefficient h. Its base is held at T_b and surrounding air at T_∞. Assume the temperature is nearly uniform across each fin section; a small transverse Biot number supports this approximation. Take the tip as adiabatic, appropriate when tip area is small or insulated. Define excess temperature $\theta = T - T_\infty$.

A slice dx loses heat $hP\theta\,dx$ to convection. Its decrease in axial heat rate must supply that loss. Fourier’s law $Q_x = -kA_c\,d\theta/dx$ then yields $\theta'' - m^2\theta = 0$, where $m = \sqrt{hP/(kA_c)}$. The boundary conditions are

$$\theta(0) = \theta_b,\ \theta'(L) = 0.$$

Using a hyperbolic cosine referenced to the insulated end gives

$$\theta(x) = \theta_b \frac{\cosh[m(L-x)]}{\cosh(mL)},$$

$$Q_b = \sqrt{hPkA_c}\,\theta_b \tanh(mL).$$

If the whole fin stayed at base temperature, heat removal would be $hPL\theta_b$. The ratio $\eta_f = \tanh(mL)/(mL)$ is the fin efficiency under the stated tip model. A very long fin can have low efficiency although it removes more heat than a short one; efficiency measures use of installed area, while total heat removal measures cooling duty.

Decide whether doubling fin length is worthwhile

An aluminium fin has rectangular section 20 mm by 2 mm, length 50 mm and conductivity 180 W/(m K). Air at 25 °C flows over it with $h = 35\,\mathrm{W/(m^2K)}$, and the base is at 75 °C. Thus

$$A_c = 4.0 \times 10^{-5}\,\mathrm{m^2},\ P = 0.044\,\mathrm{m},\ \theta_b = 50\,\mathrm{K}.$$

The transverse length $A_c/P = 0.909\,\mathrm{mm}$ gives $hA_c/(kP) = 0.000177$, consistent with small cross-sectional gradients. The fin parameter is $m = \sqrt{35(0.044)/[180(4\times10^{-5})]} = 14.6\,\mathrm{m^{-1}}$; hence $mL = 0.731$ and $\tanh(mL) = 0.624$.

The base heat rate is

$$\sqrt{35(0.044)(180)(4\times10^{-5})}(50)(0.624) = 3.29\,\mathrm{W}.$$

Fin efficiency is

$$0.624/0.731 = 0.854.$$

Doubling length makes $mL = 1.46$ and raises heat removal to about $5.27\tanh(1.46) = 4.73\,\mathrm{W}$. Twice the material therefore gives only 44% more cooling. The original covered base area by itself would convect $35(4\times10^{-5})(50) = 0.070\,\mathrm{W}$, so the fin still offers a large improvement. Adding several fins can change the film coefficient through crowding, and finite contact resistance can lower the effective base temperature. Those effects must be checked before multiplying a single-fin heat rate by the number installed.

A surface thermometer offers a second way to test the heat-transfer model. The insulated-tip solution gives

$$T(L) = 25 + 50/\cosh(0.731) = 64.1^\circ\mathrm{C}.$$

If a thin thermocouple instead measures nearly 75 °C all along a fin while the base still sheds only the predicted few watts, the effective film coefficient may be smaller than assumed. A much colder measured tip could indicate strong local tip convection or poor base contact. Interpret a profile together with the base heat rate: several combinations of conductance, base temperature and film conditions can reproduce one measured temperature, but they need not reproduce both the profile and the total cooling duty.

The root of the fin offers a particularly direct simultaneous measurement of axial temperature gradient and supplied heat rate.

Infer wall cooling from a measured skin-friction force

The film coefficient in a thermal-resistance calculation carries information about transport within the flow. For a smooth isothermal plate, take a steady parallel free stream of speed U and temperature T_e. Define $\theta = (T - T_w)/(T_e - T_w)$. Both u/U and θ rise from zero at the wall to one in the outer stream. If momentum and heat diffuse at equal rates, their transport equations and boundary conditions match for a constant-property, zero-pressure-gradient boundary layer starting at the same leading edge. Under this approximation their profiles and their normalized wall gradients coincide. In turbulent flow the analogy additionally relies on similar turbulent diffusivities.

For a wall hotter than the stream the heat flux into the fluid and the magnitude of skin friction are $q'' = k(T_w - T_e)\partial_y\theta|_w$ and $\tau_w = \mu U \partial_y(u/U)|_w$. Divide the equations. At $Pr = \mu c_p/k = 1$, the unknown gradient cancels. With $q'' = h(T_w - T_e)$ the Reynolds analogy becomes

$$h \simeq \frac{c_p \tau_w}{U},$$

$$St = \frac{h}{\rho U c_p} \simeq \frac{C_f}{2},$$

$$C_f = \frac{\tau_w}{\frac{1}{2}\rho U^2}.$$

A small heat-transfer tunnel has effective Pr close to one. At $U = 10.0\,\mathrm{m\,s^{-1}}$, the stream has $\rho = 1.20\,\mathrm{kg\,m^{-3}}$ and $c_p = 1000\,\mathrm{J\,kg^{-1}K^{-1}}$. A force balance isolates the skin friction on one flush plate face of area $0.50\,\mathrm{m^2}$, and measures $0.180\,\mathrm{N}$. The mean shear is $0.360\,\mathrm{Pa}$, hence $\bar{C}_f = 0.00600$ and $\bar{h} \simeq 36.0\,\mathrm{W\,m^{-2}K^{-1}}$. A uniform wall 20 K above the stream loses approximately $36(0.50)(20) = 360\,\mathrm{W}$ by convection. Average drag produces an average coefficient because both quantities are integrated over the same face. Pressure drag from a protruding instrument, an unheated entry length or a fluid with a very different Prandtl number would invalidate that identification.

Can adding insulation initially increase heat loss?

A long wire of radius 2 mm is held 50 K above ambient. Insulation of conductivity 0.08 W/(m K) adds radial resistance but increases exterior convection area. Add cylindrical conduction and convection resistances per unit length before converting back to heat rate.

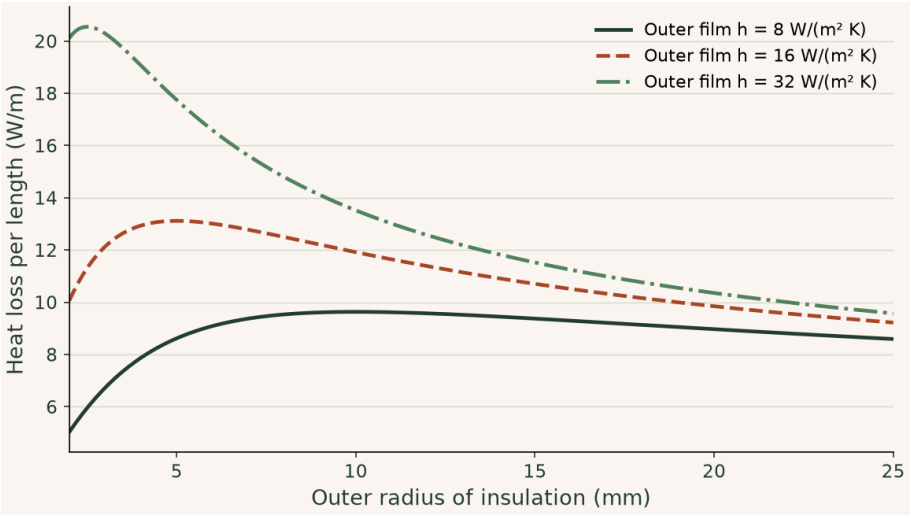

At h = 8 W/(m² K) heat loss peaks at the critical radius k/h = 10 mm; at h = 32 it peaks at 2.5 mm. This effect assumes constant surface film coefficient and prescribed inner temperature. A fixed-power heater poses a different question.

▶ **Read values from the plot**

Interactive calculation · change the inputs of this example

Raise the film coefficient until the critical radius falls below the wire radius. Check whether adding insulation can still increase heat loss.

The curves below are recalculated from this section's equations in your browser. The fixed reference plot above remains available for comparison.

Open this live calculation

Try the calculation

Conductivity (W/m K)	**0.08**
Film coefficient (W/m² K)	**8**
Wire overtemperature (K)	**50**

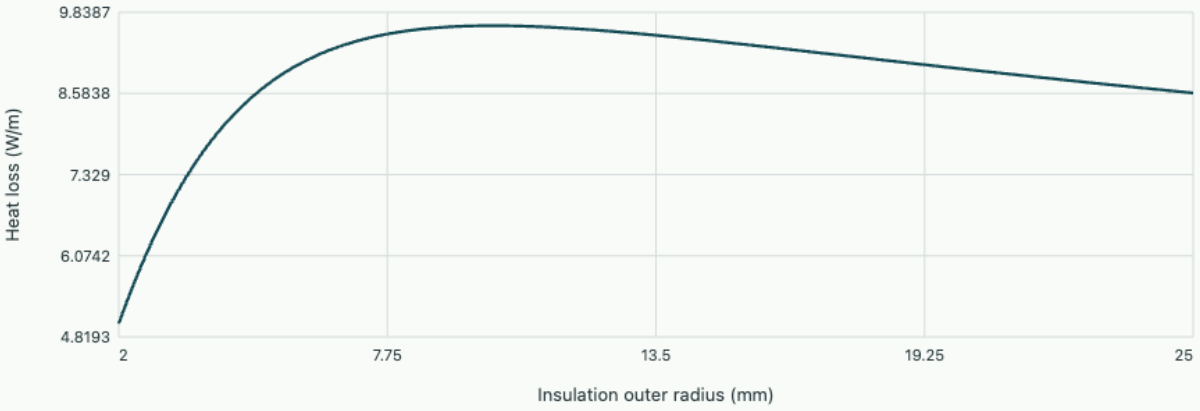

Conduction and convection

Critical radius k/h = 10.000 mm. The core radius is fixed at 2 mm and its temperature is prescribed.

Thermal diffusion and transient surface cooling

Engineering thermodynamics and heat transfer

Use diffusion time, Biot number and penetration depth to choose between a uniform-body model and a spatial temperature field.

FROM FIRST PRINCIPLES

Derive thermal diffusion from stored energy and conductive flux

Let a stationary homogeneous solid have density ρ, specific heat c and thermal conductivity k. Assume no heat generation and consider one-dimensional temperature variation. A slice of area A and thickness dx stores energy at rate $\rho c A\,dx\,\partial_t T$. The difference between the conductive flux entering and leaving gives

$$\rho c A\,dx\,\partial_t T = A[q''(x) - q''(x+dx)] \quad \Longrightarrow \quad \rho c\,\partial_t T = -\partial_x q''.$$

Fourier's law $q'' = -k\partial_x T$ closes this balance. For constant properties the diffusion equation and its characteristic time over distance L follow:

$$\partial_t T = \alpha\,\partial_{xx} T,$$
$$\alpha = \frac{k}{\rho c}, \qquad t_{\mathrm{diff}} \sim L^2/\alpha.$$

Initial temperatures and boundary heat-flux or temperature conditions are needed to select a specific solution. If conduction inside the solid is sufficiently rapid relative to transfer at its surface, one can integrate the balance using an almost uniform interior temperature. Comparing the internal and external resistance scales yields the Biot number hL_c/k used for that simplification.

Background · Biot number

$$\mathrm{Bi} = hL_c/k$$

Here h is a surface heat-transfer coefficient and k is conductivity. A small Biot number supports a nearly uniform specimen temperature for the specified boundary condition.

DERIVATION AND APPLICATIONS

Predict how a thermal probe follows a changing fluid temperature

A sensor calibrated by an abrupt temperature step can lag behind a fluid that is itself heating steadily. Take a body of mass m, heat capacity c, exposed area A and film coefficient h with negligible internal temperature gradients. This lumped description requires internal conduction to be fast compared with surface heat transfer. Treat properties and h as constant and neglect self-heating, radiation and lead conduction. Denote the body temperature by T, fluid temperature by T_f, and set $\tau = mc/(hA)$.

An instantaneous energy balance gives a first-order differential equation. For a linearly increasing fluid temperature $T_f = T_0 + \beta t$, define the lag $D = T_f - T$:

$$mc\,\dot{T} = hA(T_f - T),$$
$$\dot{D} + D/\tau = \beta.$$

Suppose the probe and fluid start in equilibrium at T_0. Multiply the lag equation by $e^{t/\tau}$ and integrate from zero. The integrating factor converts the left side to a single derivative:

$$e^{t/\tau} D(t) = \beta \int_0^t e^{s/\tau}\, ds = \beta\tau(e^{t/\tau} - 1).$$

Subtract the resulting lag from the prescribed fluid history to obtain the reading:

$$T(t) = T_0 + \beta[t - \tau(1 - e^{-t/\tau})].$$

For $t \ll \tau$ the rise begins quadratically, since the initial film temperature difference was zero. For $t \gg \tau$ the probe warms at the same rate as the fluid but reads low by $\beta\tau$. Cooling ramps reverse the sign. A measured constant lag is therefore evidence for this model only after the startup transient has decayed.

WORKED EXAMPLE

Evaluate the reading during a ramp calibration

Change the inputs and test this calculation.

A probe has $mc = 0.800$ J/K and $hA = 0.100$ W/K, hence $\tau = 8.00$ s. Start at 20.0°C and heat the fluid by 0.500 K/s. At $t = 16.0$ s the fluid is at 28.0°C and the lag is $4.00(1 - e^{-2}) = 3.46$ K, so the probe reads 24.5°C. Simply subtracting the eventual 4.00 K would overcorrect during startup.

SECOND APPLICATION

Estimate a time constant without an abrupt temperature step

Change the inputs and test this calculation.

During a long calibration ramp of 0.200 K/s, a second probe remains 1.20 K below an independent fluid reference. Its inferred time constant is

$$\tau = 1.20/0.200 = 6.00\,\mathrm{s}.$$

If the fluid is then held fixed, the remaining lag decays exponentially. Waiting $6.00\ln(1.20/0.100) = 14.9$ s brings the error below 0.100 K. The hold and the ramp test two different responses using the same fitted parameter.

Try it · Probe ramp

Set the inputs to the values in the worked example, then vary one quantity while holding the others fixed. Compare the curve and numerical balance with the second application. The calculation uses the assumptions derived above; its readout identifies bounds such as choking, phase-region changes and inconsistent measurements.

Open this live calculation

Try the calculation

Probe time constant (s)	8
Fluid ramp rate (K/s)	0.5
Ramp duration before hold (s)	16

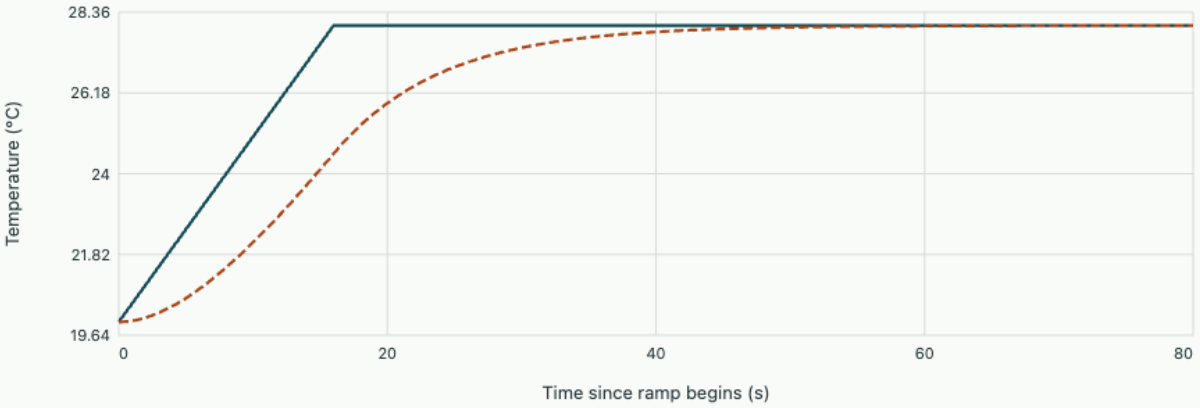

At end of ramp fluid 28 °C, probe 24.5413 °C, signed lag 3.45866 K. Long-ramp lag βτ=4 K; the hold tests the same time constant independently. Internal temperature gradients are neglected.

Allow energy to accumulate in the material

A surface thermocouple responds rapidly to a cold wash, while the interior of a thick part can remain hot. To predict the time needed to cool the centre, steady thermal resistance alone is insufficient. Enclose a differential slice of material. Incoming conductive heat minus outgoing heat plus local generation equals the rate of change of its stored energy. Substituting Fourier's law in a homogeneous stationary solid gives

$$\rho c_p \frac{\partial T}{\partial t} = k \nabla^2 T + \dot{q}_v,$$

$$\alpha = \frac{k}{\rho c_p}.$$

Thermal diffusivity α in m²/s compares a material's ability to spread heat with its ability to store it. A large conductivity does not alone guarantee fast adjustment if volumetric heat capacity is also large. Without heat generation, a length L gives a characteristic diffusion time $t_d = L^2/\alpha$. The dimensionless Fourier number $Fo = \alpha t/L^2$ measures how long diffusion has acted relative to that scale.

Decide whether the interior stays nearly uniform

At a convecting boundary $-k\partial T/\partial n = h(T_s - T_\infty)$. Comparing the resistance of the interior with a fluid film leads to $Bi = hL_c/k$, often using $L_c = V/A_s$ for a lumped model. If this ratio is small, commonly below about 0.1 as a first engineering check, the internal temperature difference is small compared with that across the film. Integrating energy over a uniform body then gives

$$\rho c_p V\, dT/dt = -hA_s(T - T_\infty),$$

$$(T - T_\infty)/(T_0 - T_\infty) = e^{-t/\tau},$$

with $\tau = \rho c_p V/(hA_s)$.

The Biot and Fourier numbers answer different questions. A small Fourier number says a disturbance has penetrated only a short distance compared with L. A small Biot number says the boundary transfers heat slowly enough for internal gradients to remain small. Choose the length for each comparison explicitly, especially for a slab cooled on one side rather than both.

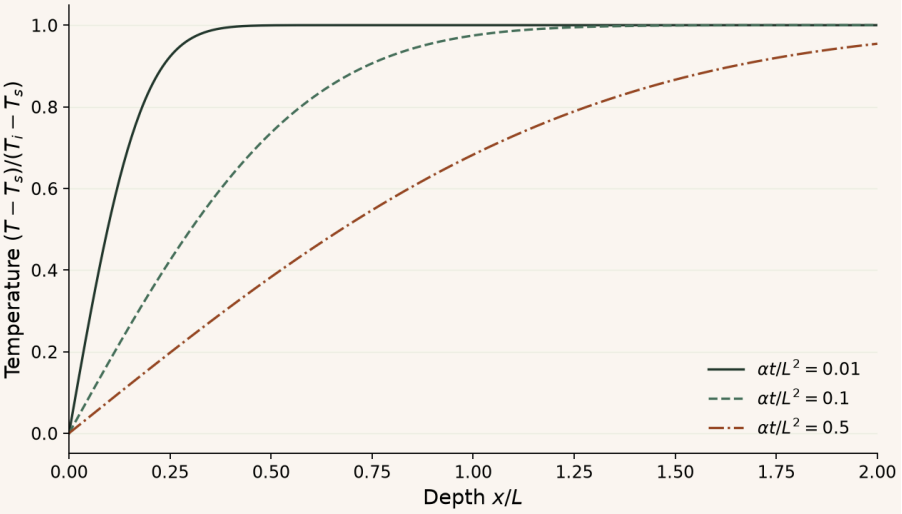

Semi-infinite cooling profiles plotted against a fixed reference length L. Increasing time by a factor of ten increases the depth of a given temperature ratio by the square root of ten.

Solve a sudden prescribed-temperature change

Consider initially uniform material occupying $x > 0$ whose surface is suddenly held at T_s, with its far interior still at T_i. Neglect internal generation and use constant properties. The initial and boundary conditions have no geometric length scale. Try the similarity coordinate $\eta = x/(2\sqrt{\alpha t})$ and set $\Theta = (T - T_s)/(T_i - T_s)$. Differentiating and substituting into the diffusion equation yields $\Theta'' + 2\eta\Theta' = 0$, with derivatives in η. Integrating once gives

$$\Theta' = Ce^{-\eta^2}.$$

The limits $\Theta(0) = 0,\ \Theta(\infty) = 1$ give

$$\Theta = \operatorname{erf}(\eta).$$

At $x = 0$ the imposed surface temperature holds exactly. At a fixed depth and very early time η is large, so the temperature tends to its initial value. The distance of any fixed fractional temperature change grows as $\sqrt{t}$. Quadrupling exposure time doubles that distance, a scaling quite different from a body travelling at a constant speed.

Predict depth and validate the time window

A thick solid initially at 373 K is held at 293 K at its surface. With $\alpha = 1.0 \times 10^{-5}$ m²/s and $t = 10$ s, $2\sqrt{\alpha t} = 0.020$ m. At depth 10 mm, $\eta = 0.50$ and $\operatorname{erf}(0.50) = 0.5205$, giving $T = 293 + 80(0.5205) = 334.6$ K. At 20 mm, $\operatorname{erf}(1) = 0.8427$ gives 360.4 K. The surface has cooled by 80 K while the deeper sensor has cooled by only 12.6 K.

At 40 s the 20 mm sensor reaches the same fraction previously measured at 10 mm. This prediction assumes the remote boundary has not begun to influence the measurement and that the surface temperature is actually prescribed. A finite slab or a convective wash may require a finite-domain diffusion solution instead. Record surface and interior histories to distinguish uncertain boundary conditions from uncertain diffusivity.

Interpret a penetration depth

If two embedded sensors are at 2 and 6 mm in the same semi-infinite specimen, the deeper sensor reaches the same normalised temperature at nine times the shallow sensor’s time because $x/\sqrt{\alpha t}$ must match. An observed factor much smaller than nine can indicate a heat leak along a probe or a far boundary that has begun to influence the temperature field.

Resolve a slab with a finite surface film

A prescribed surface temperature and a uniformly cooling lump are two useful limits. A finite slab exchanging heat through a finite film generally satisfies neither. Let L be half-thickness, $x = 0$ the midplane, and T_i the initial uniform temperature. Both sides suddenly see fluid at T_∞ with constant h. Using $X = x/L,\ Fo = \alpha t/L^2$ and $\Theta = (T - T_\infty)/(T_i - T_\infty)$, the diffusion equation becomes

$$\Theta_{Fo} = \Theta_{XX},$$

with symmetry $\Theta_X(0) = 0$ and film condition $-\Theta_X(1) = Bi\,\Theta(1)$, where $Bi = hL/k$.

Seek a separated mode $\Theta = \cos(\lambda X)e^{-\lambda^2 Fo}$. Its derivative vanishes at the midplane. Imposing the convective surface condition gives

$$\lambda \tan \lambda = Bi.$$

The uniform initial condition fixes modal amplitudes by cosine orthogonality on $0 < X < 1$:

$$\Theta = \sum_n A_n \cos(\lambda_n X) e^{-\lambda_n^2 Fo},$$

$$A_n = \frac{4 \sin \lambda_n}{2\lambda_n + \sin(2\lambda_n)}.$$

For $Bi \ll 1$, the smallest root has $\lambda_1^2 \simeq Bi$, its shape approaches a constant, and $A_1 \simeq 1$. The decay becomes $e^{-Bi\,Fo}$, exactly the lumped time scale because $Bi\,Fo = ht/(\rho c L)$. For a large film coefficient the first root tends towards $\pi/2$, corresponding to a prescribed surface temperature. The same boundary-value model thus recovers both limiting behaviours.

Predict when the core and surface meet different targets

A 40 mm thick flat slab with $k = 2.0\,\mathrm{W/(m\,K)}$ and $\alpha = 1.0 \times 10^{-6}\,\mathrm{m^2/s}$ is cooled on both sides by $h = 50\,\mathrm{W/(m^2 K)}$. It starts at 100 °C in fluid at 20 °C. Here $L = 0.020$ m and $Bi = 0.50$. The first root of $\lambda \tan \lambda = 0.5$ is 0.6533 and $A_1 = 1.070$. After 600 s, $Fo = 1.50$. The first term gives centre excess ratio $1.070 \exp[-(0.6533)^2(1.50)] = 0.564$, hence $T_c = 65.1^\circ\mathrm{C}$. At the surface multiply by $\cos(0.6533) = 0.794$, giving about 55.9 °C. Higher roots are strongly attenuated at this Fourier number.

A process requiring the centre to fall below 50 °C needs centre ratio $(50 - 20)/(100 - 20) = 0.375$. Solve $Fo = \ln(1.070/0.375)/0.6533^2 = 2.46$, or approximately 984 s. By contrast a lumped estimate gives $Fo = -\ln(0.375)/0.5 = 1.96$, about 785 s. It underestimates the core cooling time by over three minutes because the film and internal gradients both matter. An early-time surface prediction requires more modes than this late-time centre criterion. In particular the first term alone gives a centre ratio above unity at time zero, signalling that the remaining terms are necessary to reconstruct the original uniform state.

If product quality depends on average energy content rather than the hottest spot, average the first-mode shape over half the thickness. Its factor $\int_0^1 \cos(\lambda_1 X) dX = \sin \lambda_1 / \lambda_1 = 0.930$ gives mean excess ratio $0.564(0.930) = 0.525$ after 600 s. The predicted bulk temperature is therefore about 62.0 °C, between centre and surface as expected. A calorimetric measurement made after removing and mixing the material is closer to the mean than to either point thermometer. A decision to stop cooling based on such a reading can therefore satisfy an average enthalpy requirement while leaving the centre several kelvins above a local temperature limit.

The surface flux can be checked against the rate of thermal energy loss. At 600 seconds each face transfers $50(55.9 - 20) = 1795\,\mathrm{W/m^2}$ to the air. Per square metre of plan area the slab therefore loses approximately 3590 W. The volumetric heat capacity obtained independently from k/α is $2.0 \times 10^6\,\mathrm{J/(m^3 K)}$. Multiplying by total 0.040 m thickness gives an areal capacity of 80000 J/(m² K), and a predicted instantaneous mean cooling rate $3590/80000 = 0.0449\,\mathrm{K/s}$. Differentiating the one-term mean temperature gives the same decay rate to within rounding.

This independent balance is especially useful if measurements suggest two different film coefficients on opposite sides. An unequal pair of measured fluxes cannot be fitted by a strictly symmetric half-slab model even if it reproduces one centre temperature. The full-thickness conduction problem then requires both surface boundary conditions; simply replacing them by their arithmetic mean may match total heat over one interval while shifting the hottest location away from the midplane.

Interactive calculation · Change the inputs of this example

Change the parameters of this reference plot and compare its limiting cases. The fixed illustration remains beside the derivation; the adjustable calculation exposes the same dependence on geometry, material properties or loading.

Open this live calculation

Try the calculation

Thermal diffusivity (mm²/s)	5
Elapsed time (s)	10
Initial temperature (°C)	100

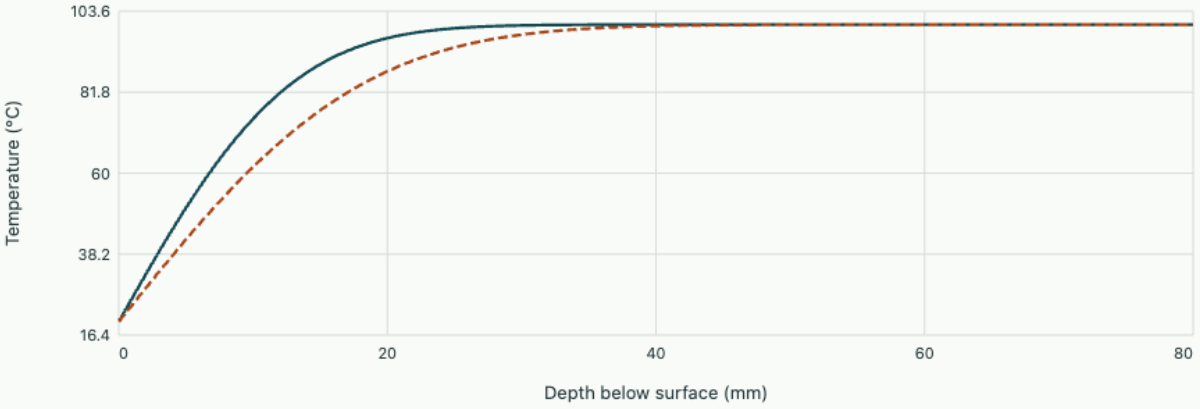

Selected time Twice elapsed time

The diffusion scale2√(αt) is 14.142 mm. Semi-infinite solid, initially uniform, surface suddenly held at20°C.

Heat-exchanger effectiveness and finite capacity rates

Engineering thermodynamics and heat transfer

Derive the counterflow effectiveness relation and determine outlet temperatures when neither stream temperature is fixed.

FROM FIRST PRINCIPLES

Derive exchanger effectiveness from the limiting heat balance

Consider a steady two-stream heat exchanger with no external heat gain, constant specific heats and no phase change. Each stream has a capacity rate $C = \dot{m}c_p$, the power required for one kelvin of temperature change. Applying the first law independently to each stream gives a common transferred heat rate

$$\dot{Q} = C_h(T_{h,i} - T_{h,o}) = C_c(T_{c,o} - T_{c,i}).$$

In passive exchange neither exit can cross the opposite stream's inlet temperature. Therefore the maximum available temperature change on either stream is bounded by the same inlet difference $T_{h,i} - T_{c,i}$. The smaller capacity rate limits the largest conceivable common duty. Define effectiveness from this physically accessible upper scale:

$$\dot{Q}_{\max} = C_{\min}(T_{h,i} - T_{c,i}),$$

$$\varepsilon = \dot{Q}/\dot{Q}_{\max},$$

$$\mathrm{NTU} = UA/C_{\min}.$$

The bound need not be reached by a finite exchanger or by every flow arrangement. Local conductance $U\,dA$ and the changing local temperature difference determine the actual effectiveness. Knowing the two capacity rates and ε suffices to reconstruct both outlet temperatures using the separate energy balances.

Background · Heat-exchanger effectiveness

$$\epsilon = \dot{Q}/[C_{\min}(T_{h,i} - T_{c,i})]$$

Heat-capacity rate is mass flow times specific heat for a single-phase stream. The smaller rate limits the maximum possible sensible heat transfer.

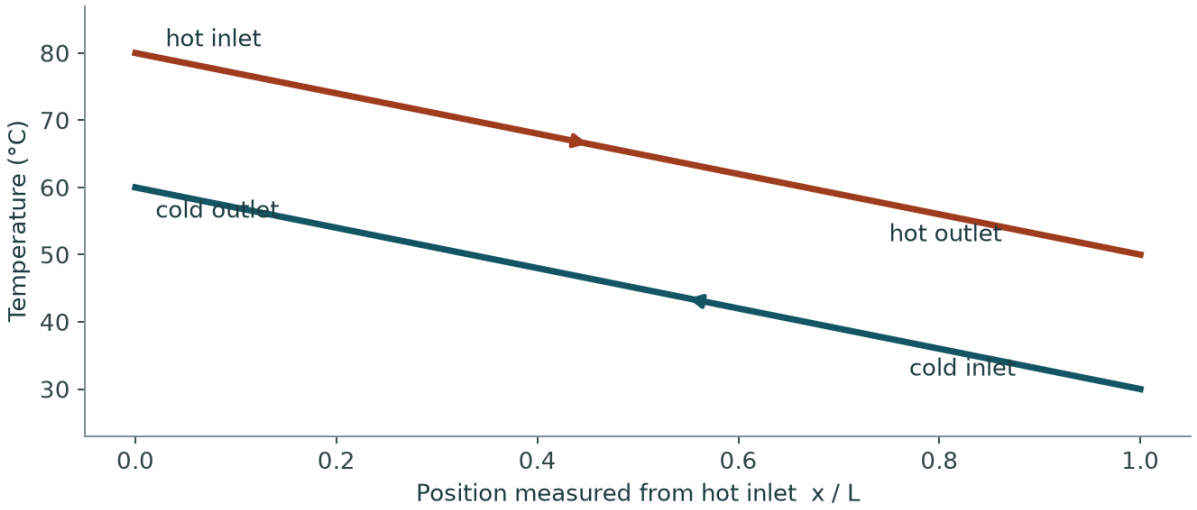

Read exchanger terminal temperatures at the same physical end

The hot stream enters at 80 °C and leaves at 50 °C. The cold stream enters at 30 °C on the right and leaves at 60 °C on the left. Equal heat-capacity rates and uniform conductance per length give parallel lines. The local driving difference is 20 K throughout, even though the cold outlet is hotter than the hot outlet.

DERIVATION AND APPLICATIONS

Where the logarithmic mean temperature difference comes from

A local heat-transfer coefficient multiplies the temperature difference at that location. Integrating one exchanger therefore requires the changing driving difference across its whole area. Take a steady counterflow exchanger with no phase change, constant positive capacity rates $C_h = \dot{m}_h c_h$, $C_c = \dot{m}_c c_c$, constant overall coefficient U and no external heat loss. Measure area coordinate A' along the hot flow. Let $d\dot{Q}$ be the positive heat transfer from hot to cold over the increment. Both fluid temperatures decrease along this chosen coordinate because the cold stream flows the other way.

$$d\dot{Q} = U(T_h - T_c)\,dA',$$

$$dT_h = -d\dot{Q}/C_h,$$

$$dT_c = -d\dot{Q}/C_c.$$

Thus $\Delta = T_h - T_c$ changes linearly with cumulative heat transfer. Write its two physical terminal values as $\Delta_1 = T_{h,i} - T_{c,o}$ and $\Delta_2 = T_{h,o} - T_{c,i}$:

$$d\Delta = (1/C_c - 1/C_h)d\dot{Q},$$

$$\Delta_2 - \Delta_1 = (1/C_c - 1/C_h)\dot{Q}.$$

To recover total area, integrate $dA' = d\dot{Q}/(U\Delta)$ using the preceding linear relation. Eliminating the reciprocal capacity difference leaves

$$\dot{Q} = UA\,\frac{\Delta_1 - \Delta_2}{\ln(\Delta_1/\Delta_2)}.$$

For equal terminal differences the apparent zero divided by zero has limit Δ_1. When one positive approach tends toward zero at a fixed heat duty, the necessary area diverges. Terminal differences must pair the temperatures at the same physical end, not two inlet temperatures chosen from opposite ends.

WORKED EXAMPLE

Size an exchanger from measured terminal temperatures

Change the inputs and test this calculation.

A hot stream cools from 90 to 55°C while the counterflow cold stream warms from 20 to 60°C. The terminal differences are 30 and 35 K, so their logarithmic mean is

$$5/\ln(35/30) = 32.4\,\mathrm{K}.$$

For a duty of 70.0 kW and $U = 500$ W/(m² K), the required area is 4.32 m². The capacities 2.00 and 1.75 kW/K independently deliver the same 70.0 kW.

SECOND APPLICATION

Compare equal duty at a tighter terminal approach

Change the inputs and test this calculation.

A redesign has measured terminal differences of 5.00 and 35.0 K at the same duty and coefficient. The logarithmic mean drops to $30/\ln 7 = 15.4$ K and the required area rises to 9.08 m². Using the arithmetic mean 20.0 K would give only 7.00 m². The region with low driving difference consumes disproportionately more area per kilowatt transferred.

Try it · Lmtd area

Set the inputs to the values in the worked example, then vary one quantity while holding the others fixed. Compare the curve and numerical balance with the second application. The calculation uses the assumptions derived above; its readout identifies bounds such as choking, phase-region changes and inconsistent measurements.

Open this live calculation

Try the calculation

Selected first terminal difference (K)	30
Second terminal difference (K)	35
Overall coefficient (W/m²/K)	500

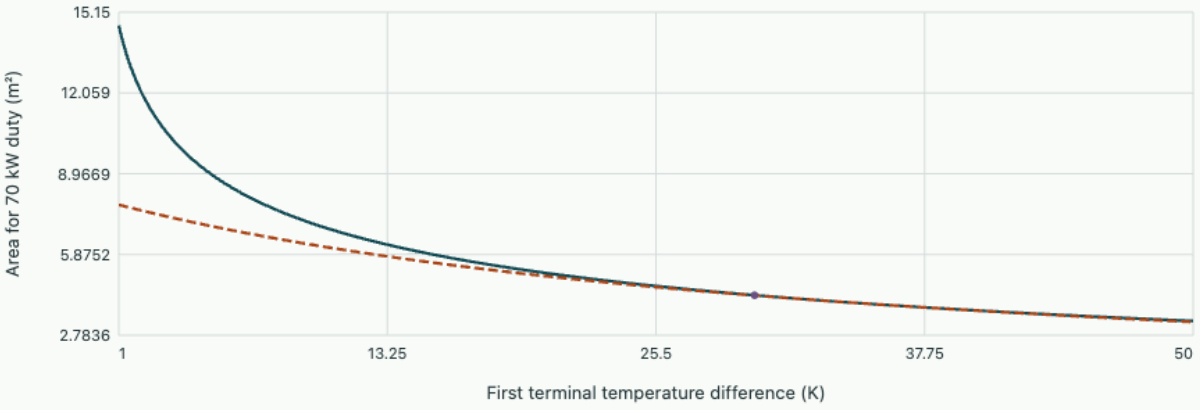

LMTD 32.4358 K; required area 4.31622 m². Equal terminal differences use the continuous limit. Positive paired approaches at each physical end are required.

Set the maximum possible heat rate

In a two-stream heat exchanger, the hot temperature falls as the cold temperature rises. Treat each stream as single phase with constant heat capacity rate $C = \dot{m}c_p$ in watts per kelvin. Neglect heat leakage, axial conduction and changes in kinetic energy. A steady overall balance gives

$$Q = C_h(T_{hi} - T_{ho}) = C_c(T_{co} - T_{ci}).$$

Performing this balance alone leaves both outlet temperatures and the transferred heat rate unknown.

The stream with smaller capacity rate would undergo the larger temperature change for a given heat rate. Even in an ideal counterflow exchanger it cannot pass the inlet temperature of the other stream at its nearest end. Thus $Q_{\max} = C_{\min}(T_{hi} - T_{ci})$. Effectiveness is $\varepsilon = Q/Q_{\max}$, which immediately gives both outlet temperatures once the heat transfer law fixes ε.

Apply a balance to an area element

Let A increase along the hot stream and write the local temperature difference $\Delta = T_h - T_c$. An element transfers $dQ = U\Delta\, dA$. In counterflow both temperature profiles decrease when followed in the hot-stream direction, giving $dT_h = -dQ/C_h$, $dT_c = -dQ/C_c$. Hence $d\ln\Delta/dA = U(1/C_c - 1/C_h)$. Integrate over area and express each end difference in terms of inlet temperatures and total heat rate using the stream balances.

$$NTU = UA/C_{\min},$$

$$C_r = C_{\min}/C_{\max},$$

$$\varepsilon = \frac{1 - \exp[-NTU(1 - C_r)]}{1 - C_r \exp[-NTU(1 - C_r)]}.$$

The sign in the differential balance follows from the chosen spatial direction. In parallel flow the cold stream warms along that coordinate and the local temperature difference decays more rapidly. For equal capacity rates the counterflow formula has a removable zero over zero. Its limit is $NTU/(1 + NTU)$, which can also be obtained from the constant temperature difference in that special case.

OneSimulation · Can the cold outlet leave hotter than the hot outlet?

The hot stream enters at 80 °C with a heat-capacity rate of 2 kW/K. The cold stream enters at 20 °C with a capacity rate of 1 kW/K. Compare the same heat exchanger conductance in parallel flow and counterflow. Obtain the two discharges from the same transferred heat rate $\dot{Q}$, using $\dot{C} = \dot{m}c_p$.

$$T_{h,o} = T_{h,i} - \dot{Q}/\dot{C}_h,$$

$$T_{c,o} = T_{c,i} + \dot{Q}/\dot{C}_c.$$

Equal outlets occur at 60 °C and require 40 kW, or effectiveness $2/3$. Parallel flow approaches this common temperature asymptotically. Counterflow reaches it at $\mathrm{NTU} = 2\ln 2$. At NTU 2 the two OneSimulation models predict cold and hot outlets of 66.5 °C and 56.8 °C in counterflow, against 58.0 °C and 61.0 °C in parallel flow. The counterflow outlets sit at opposite ends; heat still travels locally from the hotter fluid to the colder one.

```
"""Check for an outlet temperature cross in a two-stream heat exchanger."""
import matplotlib.pyplot as plt
import numpy as np
from onesim.physics.systems.process import counterflow_effectiveness, parallel_flow_effectiveness

hot_in, cold_in = 80., 20.       # °C
hot_rate, cold_rate = 2., 1.     # kW/K
capacity_min = min(hot_rate, cold_rate)
ratio = capacity_min/max(hot_rate, cold_rate)
ntu = np.linspace(0, 4, 201)

fig, ax = plt.subplots(figsize=(7, 4), constrained_layout=True)
for name, model, style in [("Counterflow", counterflow_effectiveness, "-"),
                           ("Parallel flow", parallel_flow_effectiveness, "--")]:
    heat = model(ntu, ratio)*capacity_min*(hot_in-cold_in)  # kW
    hot_out = hot_in - heat/hot_rate
    cold_out = cold_in + heat/cold_rate
    assert np.allclose(hot_rate*(hot_in-hot_out), cold_rate*(cold_out-cold_in))
    ax.plot(ntu, hot_out, style, color="#a5442e", label=f"{name}, hot outlet")
```

```
    ax.plot(ntu, cold_out, style, color="#19596a", label=f"{name}, cold outlet")

ax.axvline(2*np.log(2), color="0.5", linestyle=":", label="Counterflow outlets equal")
ax.set(xlabel="Transfer units UA / minimum capacity rate", ylabel="Outlet temperature (°C)")
ax.legend(fontsize=8)
fig.savefig("engineering-thermodynamics.png", dpi=170)
```

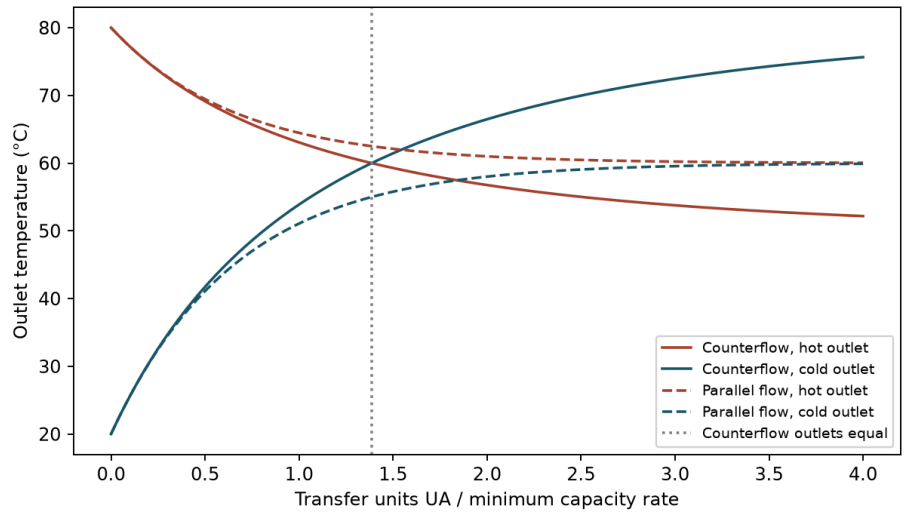

Compare outlets at the same horizontal coordinate. A temperature cross in counterflow describes two different positions along the exchanger.

Download the Python example

Change the inputs of this exact example and compare the recalculated curves with the recorded OneSimulation output above. The browser evaluates the stated analytical or modal reference.

Open this live calculation

Try the calculation

Hot inlet (°C)	80
Cold inlet (°C)	20
Cold / hot capacity rate	0.5
Selected NTU	2

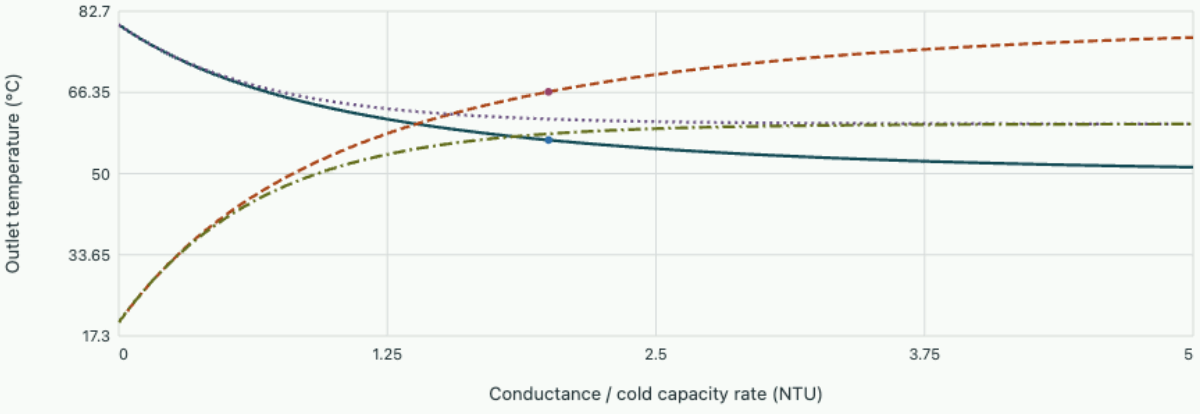

Hot, counterflow · Cold, counterflow · Hot, parallel flow · Cold, parallel flow · Chosen hot counterflow · Chosen cold counterflow

At selected NTU, counterflow cold/hot=66.5/56.8°C; parallel=58.0/61.0°C. Counterflow outlets occur at opposite physical ends. Constant properties, no ambient loss.

Predict both outlet temperatures

A counterflow cooler receives a hot stream at 90 °C with $C_h = 2.0\,\mathrm{kW/K}$ and a cold stream at 20 °C with $C_c = 1.0\,\mathrm{kW/K}$. Its conductance UA is 1.50 kW/K. Thus $C_r = 0.5$, $NTU = 1.5$, and the exponential is

$$e^{-0.75} = 0.4724.$$

Effectiveness is 0.6908. Consequently $Q = 0.6908(1.0)(90 - 20) = 48.4\,\mathrm{kW}$,

$$T_{ho} = 90 - 48.4/2 = 65.8^{\circ}\mathrm{C},$$

$$T_{co} = 20 + 48.4 = 68.4^{\circ}\mathrm{C}.$$

The cold outlet is warmer than the hot outlet. They are at opposite physical ends, so this does not imply a local reversal of heat flow. The end differences are $90 - 68.4 = 21.6$ K and $65.8 - 20 = 45.8$ K, both positive. Their logarithmic mean is approximately $(45.8 - 21.6)/\ln(45.8/21.6) = 32.2\,\mathrm{K}$, independently recovering $Q = UA\Delta T_{\mathrm{lm}} \simeq 48.3\,\mathrm{kW}$. This local check catches errors caused by swapping the terminal temperatures of a counterflow exchanger.

Predict the counterflow experiment

A smaller stream capacity rate raises NTU when the exchanger conductance is held fixed, yet it can reduce heat transfer in watts because less thermal capacity passes each second. For inlet temperatures 350 and 290 K and $C_r = 0.60$, use the counterflow expression to predict the effectiveness at NTU 1.5 and 3.0. Then halve both capacity rates, as the model does to obtain the second NTU, and predict the ratio of heat rates. Check both outlet temperatures and the hot-stream heat rate, paying attention to which stream has the smaller capacity. Open the heat-exchanger experiment.

Identify which side controls an exchanger’s conductance

The effectiveness relation uses a single conductance UA, but a designer may need to determine whether the hot film, cold film or fouling dominates it. Write the total thermal resistance in K/W as $R_{\mathrm{tot}} = 1/(h_h A_h) + R_w + R_f + 1/(h_c A_c)$. Areas must be kept explicit if the two sides use different reference areas. Suppose a test can vary hot-side flow while holding cold-side flow, wall resistance and deposit condition fixed. Let a supported correlation in the tested region imply $h_h \propto \dot{m}_h^{\,n}$. With $m_r = \dot{m}_h/\dot{m}_{\mathrm{ref}}$, the relation takes the linear form

$$(UA)^{-1} = R_{\mathrm{fixed}} + R_{h,\mathrm{ref}}\, m_r^{-n}.$$

This is the basis of a Wilson-plot style separation. Each measured UA can be obtained from balanced heat duty and the correct logarithmic mean temperature difference. The intercept estimates the combined fixed resistance; the slope estimates the hot-film resistance at reference flow. One set of measurements alone does not separate cold film from wall and deposits individually. Temperature-dependent material properties and changes in the opposite stream can spoil the assumed linear relation.

Treating UA as proportional to flow whenever convection rises would ignore the series resistances left in its path. As m_r becomes large the hot-film term tends to zero, leaving an asymptote $1/R_{\mathrm{fixed}}$ rather than unbounded conductance. Pumping power typically rises over this same change of flow.

Use a flow sweep to distinguish fouling from film resistance

A clean counterflow rig at reference hot-side mass flow has $UA = 500\,\mathrm{W/K}$, giving total resistance $0.00200\,\mathrm{K/W}$. At twice the flow it has $UA = 600\,\mathrm{W/K}$, or $0.001667\,\mathrm{K/W}$. Over this interval use measured exponent $n = 0.80$, so

$$2^{-0.8} = 0.5743.$$

Subtracting the two equations gives

$$R_{h,\mathrm{ref}} = (0.00200 - 0.001667)/(1 - 0.5743) = 0.000783\,\mathrm{K/W}.$$

The combined fixed resistance is $0.001217\,\mathrm{K/W}$. Even an ideal hot film could raise conductance only to about 822 W/K under the unchanged fixed-side assumptions.

A later test at reference flow gives $UA = 400\,\mathrm{W/K}$, so total resistance has grown to 0.00250 K/W. If a second flow sweep confirms the same hot-film slope, the added fixed resistance is 0.00050 K/W. For a heat-

transfer area of 2.0 m² on the chosen basis this corresponds to fouling resistance $R''_f = 0.0010\,\mathrm{m^2K/W}$. If instead the slope has changed, assigning the entire difference to deposits is unwarranted; changed flow distribution or fluid properties may contribute. Temperature measurement also matters strongly when the two terminal differences are close to sensor uncertainty. Compare the hot and cold enthalpy duties before fitting the conductance and retain the residual across all measured flow settings.

A third clean test can challenge the fitted exponent. At four times reference hot flow, the predicted resistance is $0.001217 + 0.000783(4^{-0.8}) = 0.001475\,\mathrm{K/W}$, or $UA \simeq 678\,\mathrm{W/K}$. A systematic departure that grows at high flow could reflect changing cold-side properties, thermocouple bias or failure of the hot-film correlation. Two calibration points always fit a two-parameter straight line; the independent third point actually tests the assumed dependence.

Determine whether a specified exchanger effectiveness is attainable

Fix both inlet temperatures and $C_{\min}$. Changing area changes NTU. For constant properties and no losses the effectiveness divides actual heat rate by $C_{\min}(T_{h,i} - T_{c,i})$. Compare two flow arrangements and the isothermal-stream limit.

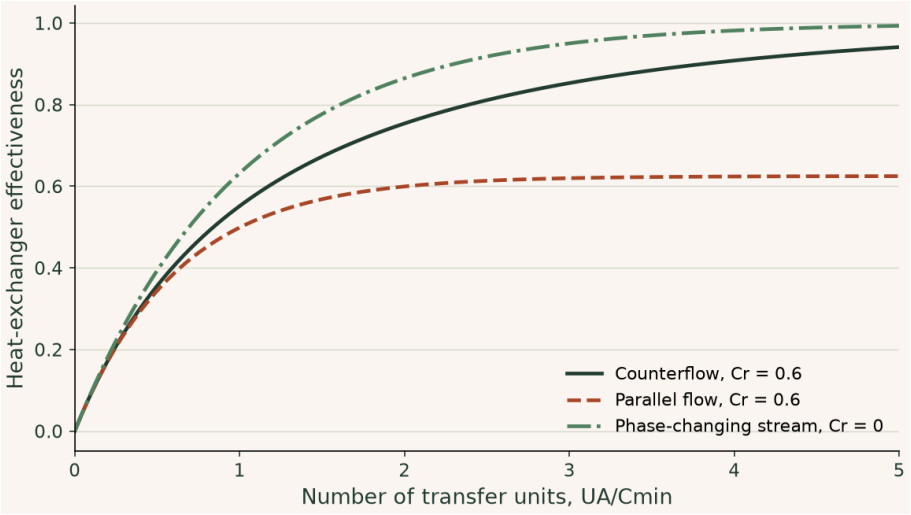

With capacity ratio 0.6, parallel flow approaches only $1/1.6 = 0.625$ however large the area. Counterflow can exceed 0.8. An area calculation that ignores the arrangement can return a target that the selected exchanger cannot reach.

▶ **Read values from the plot**

Interactive calculation · change the inputs of this example

Move the capacity-rate ratio toward one and compare the finite-area behaviour with the phase-change limit.

The curves below are recalculated from this section's equations in your browser. The fixed reference plot above remains available for comparison.

Open this live calculation

Try the calculation

Capacity-rate ratio	0.6
Hot inlet (°C)	100
Cold inlet (°C)	20

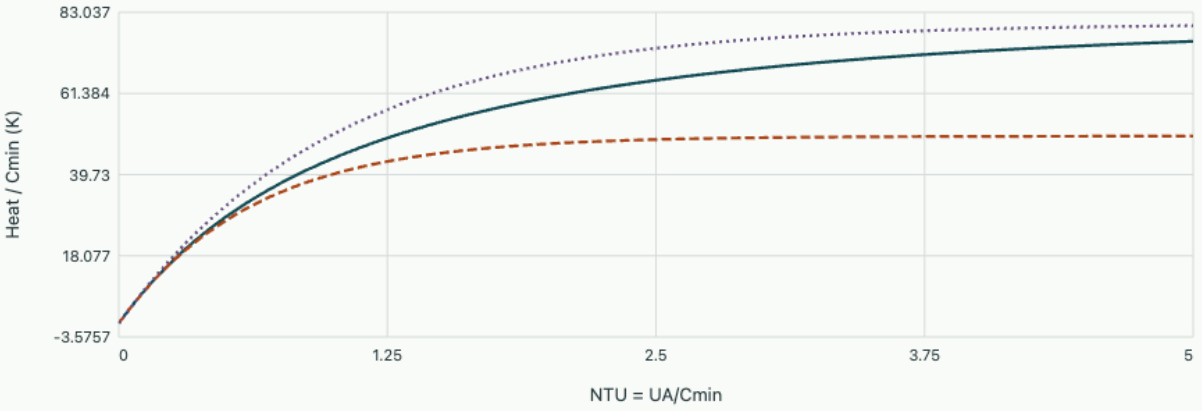

Counterflow Parallel flow Isothermal-stream limit

At NTU 2, counterflow effectiveness 0.754. Parallel flow cannot exceed 0.625.

When another centimetre of cooling fin stops helping

Engineering thermodynamics and heat transfer

A fin adds surface area, but heat must travel down its length before it can reach that area. Derive the temperature profile from an energy balance and quantify how much of the nominal area contributes to cooling.

Balance a short slice with distributed convection

Consider a straight prismatic fin of conductivity k, cross-section A, exposed perimeter P and length L. Hold the root at T_b, let a fluid at T_∞ surround the fin and use a uniform convection coefficient h. Temperature is assumed nearly uniform across each cross-section. A small transverse Biot number $h(A/P)/k$ supports that assumption. Ignore radiation and heat loss from the tip. Define $\theta = T - T_\infty$. The axial heat entering a slice, less the heat leaving its far face, must equal $hP\theta\,dx$.

$$q_x = -kA\,\theta',$$

$$-\frac{dq_x}{dx} = hP\theta,$$

$$\theta'' - m^2\theta = 0,$$

$$m = \sqrt{hP/(kA)}.$$

A positive temperature excess makes the heat flow diminish downstream. For an insulated tip $q_x(L) = 0$, so a hyperbolic cosine centered at the tip satisfies the zero-slope condition. Its coefficient follows from the root temperature.

$$\frac{\theta(x)}{\theta_b} = \frac{\cosh[m(L-x)]}{\cosh(mL)},$$

$$Q_b = kAm\theta_b \tanh(mL).$$

Efficiency follows by comparing with an isothermal fin

If every part of the lateral area stayed at T_b, the heat transfer would be $hPL\theta_b$. The actual-to-ideal ratio is therefore

$$\eta_f = \tanh(mL)/(mL).$$

For small mL the temperature falls little along the wall and doubling the area nearly doubles the heat flow. For large mL the root heat flow approaches $kAm\theta_b$. Near the long-fin limit increasing conductivity or the root cross-section can buy more cooling than adding length. Efficiency alone does not measure total heat removal because it contains the nominal area in its denominator.

Two fins made from the same stock

Use a rectangular aluminium strip, 20 mm wide, 2 mm thick, $k = 200$ W/(m K), $h = 25$ W/(m² K) and $\theta_b = 50$ K. All four long faces are exposed. With $A = 4.0 \times 10^{-5}$ m² and $P = 0.044$ m, $m = 11.73$ m⁻¹. At 50 mm

length, $mL = 0.586$, $Q_b = 2.47$ W and $\eta_f = 0.899$. The tip remains about 42.5 K above ambient.

At 150 mm, the heat rate is 4.43 W and the efficiency drops to 0.536. The material and exposed area tripled, while the heat transfer increased by about 79%. A fin root bonded through a finite contact resistance R_c would run cooler than the heated base. Replacing the temperature excess by $T_{\text{base}} - T_\infty - Q_b R_c$ in the root heat equation gives a direct way to include the bond in a thermal circuit.

Choose a short fin and inspect how far the infinite-fin approximation is from the insulated-tip temperature. Increase length until the root heat rate is almost saturated. Estimate the added cooling from another centimetre at that point.

Open this experiment online

Try the calculation

Length (mm) **50**

Convection coefficient (W/m² K) **25**

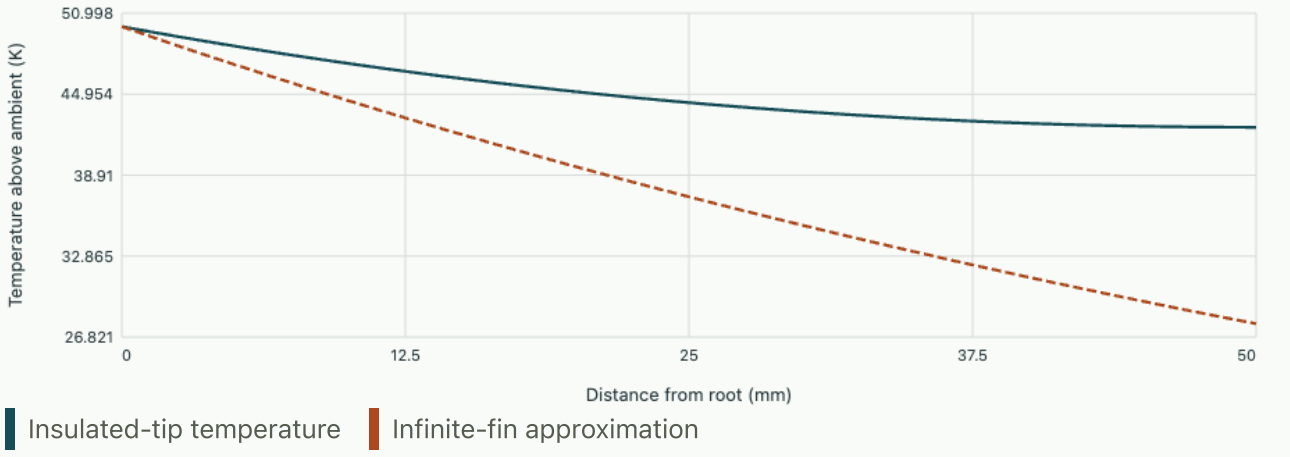

The finite tip is 42.5 K above ambient. Root heat flow is 2.47 W and the fin efficiency is 89.9%.

Tutorial · select a length from a marginal-benefit criterion

For the same strip, require 95% of the infinitely long fin's heat rate. Find the required length. Then estimate the heat rate if the base is held 50 K above ambient through a contact resistance of 1 K/W.

▼ **Show the inverse calculation**

Set $\tanh(mL) = 0.95$, giving $mL = \tfrac{1}{2}\ln(1.95/0.05) = 1.832$ and $L = 0.156$ m. The root conductance there is $G = kAm(0.95) = 0.0891$ W/K. Adding the contact resistance gives

$$Q = 50/(1/G + 1) = 4.09\,\text{W}.$$

The temperature drop across the contact is about 4.09 K. This check distinguishes the fin-root temperature from the base temperature reported by a sensor under the bond.

Further reading · Lienhard and Lienhard · A Heat Transfer Textbook, chapter 4 on conduction and fins

A moving boundary in a freezing layer

Engineering thermodynamics and heat transfer

A flat cold surface grows a layer of ice from water initially at its melting temperature. Estimate when it will reach a specified thickness, then discover how a finite heat-transfer resistance changes its early growth.

Write a balance at a boundary whose position is unknown

Let the solid occupy $0 < x < s(t)$. The boundary at x=0 is held at $T_w < T_m$, the liquid remains at T_m, and the solid-liquid interface is at its equilibrium melting temperature. Treat density, latent heat of fusion L_f and solid conductivity k as constants. Ignore volume change and fluid convection. When the front moves by ds, each unit area releases latent energy $\rho L_f \, ds$. In this one-phase problem conduction into the cold side removes that energy. Taking the magnitude of heat extracted, the interface condition reads

$$\rho L_f \, \frac{ds}{dt} = k \, \frac{\partial T}{\partial x}\bigg|_{s^-}.$$

The derivative is positive because temperature increases towards the liquid. Solving the full problem requires time-dependent conduction in a domain with a moving outer end. A useful approximation emerges when the sensible heat needed to cool the new solid is small beside its latent heat. Measure that ratio with the Stefan number $\mathrm{Ste} = c_p(T_m - T_w)/L_f$.

Derive the square-root law and check the approximation

For small Ste the temperature profile in the solid adjusts quickly compared with front growth. Approximate the gradient by $(T_m - T_w)/s$. Multiplying the interface balance by s separates the variables. Integrating from an initial layer s_0 gives

$$s \, ds = \frac{k(T_m - T_w)}{\rho L_f} \, dt,$$

$$s^2 = s_0^2 + \frac{2k(T_m - T_w)}{\rho L_f} t.$$

With zero initial thickness the predicted speed starts unbounded because the model inserts a finite temperature difference across zero resistance. A finite surface heat-transfer coefficient removes this singularity. If coolant at T_c reaches the wall through a film $1/h$, the heat flux is $(T_m - T_c)/(1/h + s/k)$. For a layer growing from zero, integration gives

$$t = \frac{\rho L_f}{T_m - T_c} \left(\frac{s}{h} + \frac{s^2}{2k} \right).$$

The linear term dominates while the film presents the larger resistance. Once the solid is thick, the quadratic term recovers diffusion-controlled growth. These two regimes can be distinguished by recording thickness at equal time intervals.

An ideal wall and a finite film

For water freezing at 0 °C take $k = 2.2$ W/(m K), $\rho = 917$ kg/m³, $L_f = 334$ kJ/kg and $c_p = 2.1$ kJ/(kg K). With a wall at −5 °C the Stefan number is 0.0314. The ideal-wall time to form 10 mm is $917(334000)(0.01)^2/[2(2.2)(5)] = 1392$ s, about 23 minutes. Reaching 20 mm requires four times that duration. The increasing conduction distance accounts for the slowing front.

With coolant at −5 °C and $h = 200$ W/(m² K), the added time for 10 mm is 3063 s. The resulting total is about 74.2 minutes. At that point the film resistance is 0.005 m² K/W and the ice contributes 0.00455 m² K/W. Both terms are of the same order. Vary h and find when a second thickness measurement would provide more information about ice conductivity than a first one.

At a fixed coolant temperature compare the early and late parts of the thickness histories. Increase the film coefficient strongly. Predict where the finite-film curve should approach the fixed-wall result first in relative terms, then read the completion time.

Open this experiment online

Try the calculation

Coolant depression below freezing (K) **5**

Film coefficient (W/m² K) **200**

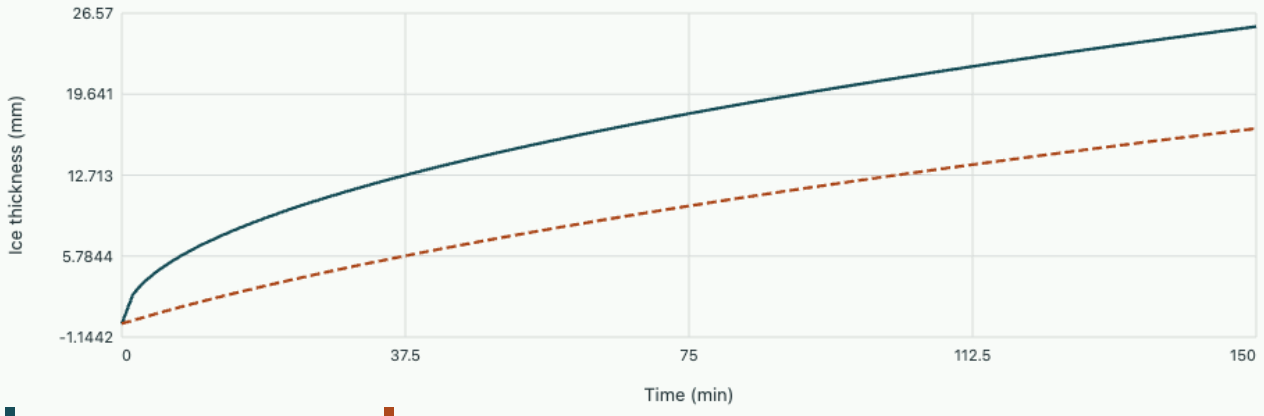

At the selected film coefficient 10 mm requires 74.2 minutes. The fixed-wall prediction is 23.2 minutes.

Tutorial · recover the coolant coefficient from one observation

A 15 mm layer is observed after 8000 s with the same properties and coolant temperature. Estimate h from the integrated equation. Would a time below the ideal-wall value be compatible with a positive h in this model?

▼ **Show the resistance fit**

The ideal-wall contribution is about 3132 s. The remaining 4868 s equals $\rho L_f s/[h(T_m - T_c)]$, giving $h \simeq 189$ W/(m² K). Both resistances are positive. A shorter time than the ideal-wall prediction would require checking the temperature, initial ice thickness, geometry or material data before assigning a heat-transfer coefficient.

Further reading · RWTH Aachen · Computational Modeling of Multiphysics, chapter 12 on the Stefan problem

Diffusion with a stationary carrier gas

Engineering thermodynamics and heat transfer

A solvent evaporates into a gas layer. Its molecules diffuse relative to the mixture, while evaporation also induces a bulk molar flow. Work through the species balances to see why the flux ceases to scale linearly with surface concentration.

Choose the reference frame for a flux

Let A be the evaporating species and B an insoluble carrier. A planar film of thickness L lies between the evaporating surface, at mole fraction y_s, and a well-mixed reservoir at $y_\infty < y_s$. Keep total pressure, temperature, total molar concentration $c = p/(RT)$ and binary diffusivity D uniform. In a stationary laboratory frame the total species flux equals diffusion relative to the molar-average velocity plus transport by that velocity. If positive z points towards the reservoir, write

$$N_A = -cD\frac{dy_A}{dz} + y_A(N_A + N_B).$$

Steady conditions in one dimension make N_A, N_B independent of z. Carrier B cannot enter the liquid and has zero net flux there. This fixes $N_B = 0$ throughout the film. B still diffuses towards the interface relative to the moving mixture. Its diffusive flux balances its outward convective flux. The constraint $N_A + N_B = 0$ would instead describe equimolar counterdiffusion between two exchanging species.

Separate variables and enforce both boundary concentrations

Collect the terms in N_A, divide by $1 - y_A$ and integrate from the surface to the reservoir. The remaining integral is a logarithm of the carrier mole fraction.

$$N_A(1 - y_A) = -cD\frac{dy_A}{dz},$$

$$N_A = \frac{cD}{L}\ln\frac{1 - y_\infty}{1 - y_s}.$$

$$1 - y_A(z) = (1 - y_s)\left(\frac{1 - y_\infty}{1 - y_s}\right)^{z/L}.$$

The concentration profile curves because a given net flux requires a different concentration gradient at different carrier fractions. A dilute expansion of the logarithm gives

$$N_A \simeq (cD/L)(y_s - y_\infty).$$

The exact-to-dilute flux ratio exceeds unity when the surface emits A. The discrepancy becomes measurable long before the carrier disappears. At fractions very close to unity, the assumptions of a stagnant, isothermal film also deserve new scrutiny.

Small and large amounts of vapour at the surface

At 300 K and 1 bar, $c = 40.09$ mol/m³. Take $D = 2.0 \times 10^{-5}$ m²/s and $L = 5$ mm. The coefficient $cD/L = 0.1604$ mol/(m² s). With a vapour-free reservoir and $y_s = 0.05$, the logarithm is 0.05129 and the predicted flux is 0.00823 mol/(m² s). A linear gradient would give 0.00802, about 2.5% lower.

For $y_s = 0.50$, the flux rises to 0.111 mol/(m² s), versus 0.0802 for the dilute estimate. At mid-depth the exact fraction is $1 - \sqrt{0.5} = 0.293$, while linear interpolation gives 0.25. An experimental check can therefore use either a mass-loss measurement or a local concentration measurement. Evaporative cooling changes the saturation mole fraction at the surface, so a measured surface temperature matters even if room temperature seems stable.

Use a vapour-free reservoir and increase the surface fraction from 0.10 to 0.80. Estimate which midpoint concentration lies farther from linear interpolation. Compare the resulting flux errors and identify the physical transport term responsible.

Open this experiment online

Try the calculation

Surface vapour mole fraction **0.5**

Reservoir vapour mole fraction **0**

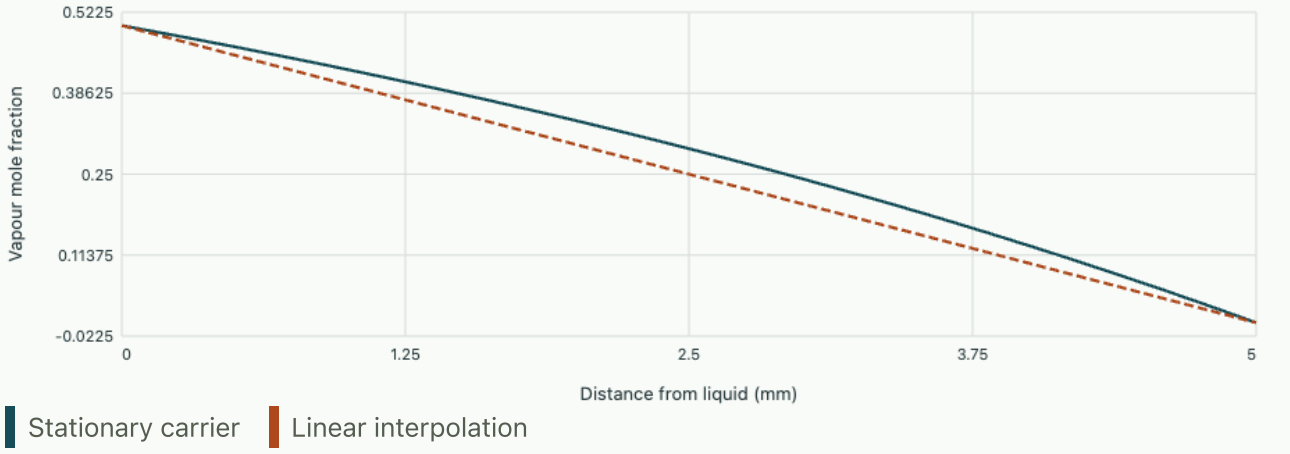

The evaporation flux is 0.1112 mol/(m² s). The dilute estimate is 0.0802 mol/(m² s).

Tutorial · recover the effective film thickness

At the same temperature and diffusivity the surface mole fraction is 0.30, the reservoir fraction is 0.05 and a balance reports $N_A = 0.050$ mol/(m² s). Estimate L. What film thickness would be inferred if the outward molar flow were ignored?

▼ **Show the two fits**

Solving the exact expression gives

$$L = cD\ln(0.95/0.70)/N_A = 4.90\,\mathrm{mm}.$$

The linear fit yields

$$L = cD(0.25)/N_A = 4.01\,\mathrm{mm}.$$

The latter fit makes the film thinner to compensate for omitted advection. In an experiment that effective thickness would predict the wrong flux when the surface composition changes appreciably.

Further reading · Lienhard and Lienhard · A Heat Transfer Textbook, chapter 11 on mass transfer by diffusion

Measure how quickly a dielectric follows an electric field

Fields and the response of matter

Derive the time and frequency responses of a relaxing polarisation, locate the loss maximum and decide whether a measured sample contains more than one relaxation process.

From a delayed polarisation to a constitutive equation

A slab of insulating material is placed between parallel electrodes. At first the measured charge increases very quickly when a small voltage is switched on. It then continues to rise slowly even though the applied voltage is fixed. A constant real relative permittivity can describe the final charge but cannot predict this record. The experiment reveals at least two relevant time scales. For a measurement window that resolves only the slower process, treat the fast electronic response as instantaneous and describe the slower polarisation P by its approach to an equilibrium value.

At fixed weak field the equilibrium slow polarisation is $\epsilon_0 \Delta\epsilon E$, with $\Delta\epsilon > 0$. Assume the rate of change is proportional to the distance from that equilibrium, with one positive relaxation time τ. Then

$$\tau \dot{P} + P = \epsilon_0 \Delta\epsilon E(t).$$

This is a phenomenological material law. Its assumptions include linear response about the prepared state, a time-independent τ, and a sample whose constituents can be described by one dominant relaxation. Testing those assumptions is part of the experiment.

For $E(t) = E_s$ switched on at zero from an unpolarised state, solve using the integrating factor $e^{t/\tau}$. The result is

$$P = \epsilon_0 \Delta\epsilon E_s (1 - e^{-t/\tau}).$$

After one τ, 63.21% of the slow response has developed; after 3τ, 95.02%. Measuring an absolute fraction requires subtracting the fast initial jump and knowing the eventual plateau. A long baseline drift can look like a second relaxation if those references are not controlled.

Derive the complex dielectric response

Now use peak phasors with time dependence $e^{-i\omega t}$. Differentiation contributes $-i\omega$, so

$$P/E = \epsilon_0 \Delta\epsilon / (1 - i\omega\tau).$$

Include the high-frequency relative permittivity ϵ_∞, where “high” means above the slow process but below additional electronic resonances. The real and imaginary parts of relative permittivity are

$$\epsilon' = \epsilon_\infty + \frac{\Delta\epsilon}{1 + (\omega\tau)^2},$$

$$\epsilon'' = \frac{\Delta\epsilon\, \omega\tau}{1 + (\omega\tau)^2}.$$

At fixed field amplitude the polarisation current is

$$\dot{P} = -i\omega P.$$

The mean electrical work per unit time and volume is

$$\tfrac{1}{2}\mathrm{Re}(E^{*}\dot{P}) = \tfrac{1}{2}\omega\epsilon_0\epsilon''|E|^2.$$

Loss therefore has positive ϵ'' in the chosen time convention. Using $e^{+i\omega t}$ reverses the sign of the imaginary part of permittivity but leaves mean absorbed energy positive when the current convention is changed consistently.

Set $u = \omega\tau$. Differentiating $u/(1+u^2)$ gives $(1-u^2)/(1+u^2)^2$. Hence ϵ'' peaks at $u = 1$, with height $\Delta\epsilon/2$, and the real part lies halfway between its low- and high-frequency plateaux. The peak identified here is in the imaginary *permittivity*. A plot of dissipated power at fixed field includes an additional factor of frequency and need not have a peak at the same location.

Tie a frequency sweep to a step test

Suppose a material has $\epsilon_\infty = 2$, $\Delta\epsilon = 6$, and a measured loss peak at 1 kHz. Then $\tau = 1/(2\pi\,1000) \simeq 159\ \mu\mathrm{s}$. At the peak the relative permittivity is $5 + 3i$ in our convention. A time-domain voltage step should show a slow charge transient reaching 63.2% of its added charge in roughly 159 µs. This comparison connects two experiments that weight electrode offsets and parasitic impedances differently.

An independent shape check follows by eliminating u. Let $X = \epsilon' - \epsilon_\infty$ and $Y = \epsilon''$. From the formulas $X^2 + Y^2 = \Delta\epsilon X$. Thus $(X - \Delta\epsilon/2)^2 + Y^2 = (\Delta\epsilon/2)^2$: a single Debye relaxation traces a semicircle in the complex plane. Departures from a semicircle can reflect several relaxation times, spatially inhomogeneous samples, leakage or imperfect fixture correction. They do not uniquely identify the microscopic mechanism.

Resolve a second process only when the data require it

Open the live experiment

Try the calculation

Slow fraction of total dielectric strength **0.5**

Slow / fast relaxation time **12**

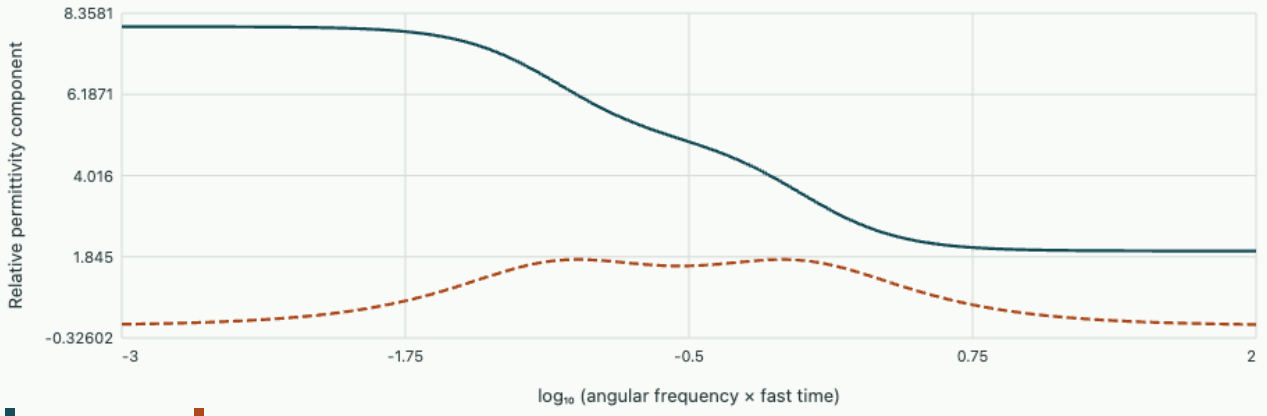

Largest sampled loss 1.7724 near ωτ_fast = 0.1. Relative permittivity tends to 8 at low frequency and 2 at high frequency.

The synthetic sample has two groups contributing a fixed total $\Delta\epsilon = 6$. Change the fraction in the slower group and its relaxation time relative to the faster group. At equal relaxation times the groups give exactly the single-process response whatever their fractions. When the times differ, the imaginary curves broaden and may separate. Observe the real dispersion steps as well. A broad lossy peak alone can be fitted by more than one pair of times and strengths.

In a practical measurement the low-frequency tail can contain DC conductivity. In the $e^{-i\omega t}$ convention,

$$\sigma E - i\omega\epsilon_0\epsilon E = -i\omega\epsilon_0[\epsilon + i\sigma/(\epsilon_0\omega)]E.$$

This added imaginary part is proportional to $1/\omega$, whereas the low-frequency side of an isolated Debye relaxation rises as ω. Change electrode separation and use a DC conduction measurement to challenge the interpretation. An electrode interface can add a further time scale, so an observed fit parameter deserves an uncertainty and a physical validation step.

Tutorial

The loss maximum is at 2 kHz and its height is 4 in relative-permittivity units. The high-frequency plateau is 3. Predict the low-frequency plateau and the slow response at $t = 100\ \mu$s after a step.

▼ **Work through the solution**

The single-time model gives $\Delta\epsilon = 8$ and the low-frequency plateau 11. The relaxation time is $1/[2\pi(2000)] \simeq 79.6\ \mu$s. The slow fraction after 100 µs is $1 - e^{-100/79.6} \simeq 0.715$. Repeat at another step amplitude to check linearity.

At the loss peak, is the phase difference between the *total* displacement and the electric field always 45°?

▼ **Work through the solution**

The slow polarisation alone has a 45° phasor angle at $\omega\tau = 1$. The total displacement includes $\epsilon_\infty E$. Its angle is $\arctan[(\Delta\epsilon/2)/(\epsilon_\infty + \Delta\epsilon/2)]$, generally less than 45°.

Why does observing only a high-frequency tail constrain a ratio of parameters more tightly than the individual parameters?

▼ **Work through the solution**

For large $\omega\tau$,

$$\epsilon'' \simeq \Delta\epsilon/(\omega\tau).$$

Many values with the same ratio $\Delta\epsilon/\tau$ have very similar tails. Measure the peak or both plateaux to distinguish the two parameters.

Compare with material measurements

Inspect the orientational response and dielectric-loss treatment in MIT's material-properties course. Use the phasor convention of each source to predict the sign assigned to loss.

Measure equilibrium field fluctuations with a capacitor

Fields and the response of matter

Derive an electrical noise measurement in two ways, determine what fraction reaches a finite-bandwidth detector and use agreement to calibrate an experiment.

A zero-average voltage can store measurable energy

Connect a capacitor C across a resistor R at a uniform absolute temperature T, with no external drive. Ideal circuit equations predict that any deliberately deposited charge decays on time scale RC. At thermal equilibrium, microscopic fluctuations continually exchange charge and energy with the capacitor. Symmetry gives zero mean voltage, but the stored energy involves V^2. In the classical limit a quadratic degree of freedom has mean energy $k_B T/2$, giving $\langle V^2 \rangle = k_B T/C$. The resistor sets the time scale of equilibration; equilibrium mean-square voltage does not depend on resistance in this ideal model.

Check the dimensions: $k_B T$ is an energy in joules and capacitance has units coulombs per volt. Their ratio is volts squared. At $T = 300$ K and $C = 100$ pF the RMS voltage is $\sqrt{k_B T/C} \simeq 6.44\ \mu\mathrm{V}$. A voltmeter that averages the signed voltage for a long time reports a value near zero. Measuring the noise requires an RMS, variance or appropriately calibrated spectral measurement.

Derive the same variance from a spectral measurement

Use a one-sided voltage power spectral density per hertz at positive frequencies. The classical Johnson–Nyquist open-circuit result for a resistor is $S_R(f) = 4k_B T R$ V²/Hz. Its square root is an amplitude spectral density in V/$\sqrt{\mathrm{Hz}}$, not a voltage variance. Use an equivalent series voltage noise source and the RC transfer function $H(f) = 1/[1 + 2\pi i f RC]$. A stationary linear system multiplies the voltage power spectrum by $|H|^2$, so

$$S_C(f) = \frac{4k_B T R}{1 + (2\pi f RC)^2},$$

\int_0^\infty S_C(f)\,df=\frac{k_BT}{C}.

The integral follows by $u = 2\pi RCf$ and $\int_0^\infty du/(1+u^2) = \pi/2$. It furnishes a detailed check on the factor four, the use of one-sided density, the conversion from hertz to angular frequency and the amplitude-versus-power distinction. If your result disagrees by a constant factor of two or 2π, audit these conventions before assigning a new noise temperature.

For $R = 10$ kΩ, $T = 300$ K, the low-frequency amplitude spectral density is approximately 12.9 nV/$\sqrt{\mathrm{Hz}}$. The RC corner at $C = 100$ pF is $f_c = 1/(2\pi RC) \simeq 159$ kHz. The open-circuit low-frequency density rises with R while the bandwidth shrinks in inverse proportion. Their integrated product gives the earlier resistance-independent variance.

Include a real observation bandwidth

Suppose a first measurement adds an ideal low-pass cutoff B. The recorded fraction of equilibrium variance is $F = (2/\pi)\arctan(B/f_c)$. Half of the noise variance lies below the RC corner. That may be surprising if

the corner was introduced as the frequency at which the *spectral density* reaches half its zero-frequency value. These two properties coincide for this single-pole spectrum for different mathematical reasons.

At $B = 0.1 f_c$, $F \simeq 0.06345$, so the recorded RMS is about 25.2% of the infinite-band RMS. At $B = 10 f_c$, $F \simeq 0.93655$, giving about 96.8% of the RMS. In an experiment, taking a square root one step too early changes this comparison. A real detector transfer function replaces the sharp cutoff by an integral of $S_C(f)|H_{\mathrm{det}}(f)|^2$.

Open the live experiment

Try the calculation

Resistance (kΩ) **10**

Temperature (K) **300**

Detector bandwidth (kHz) **100**

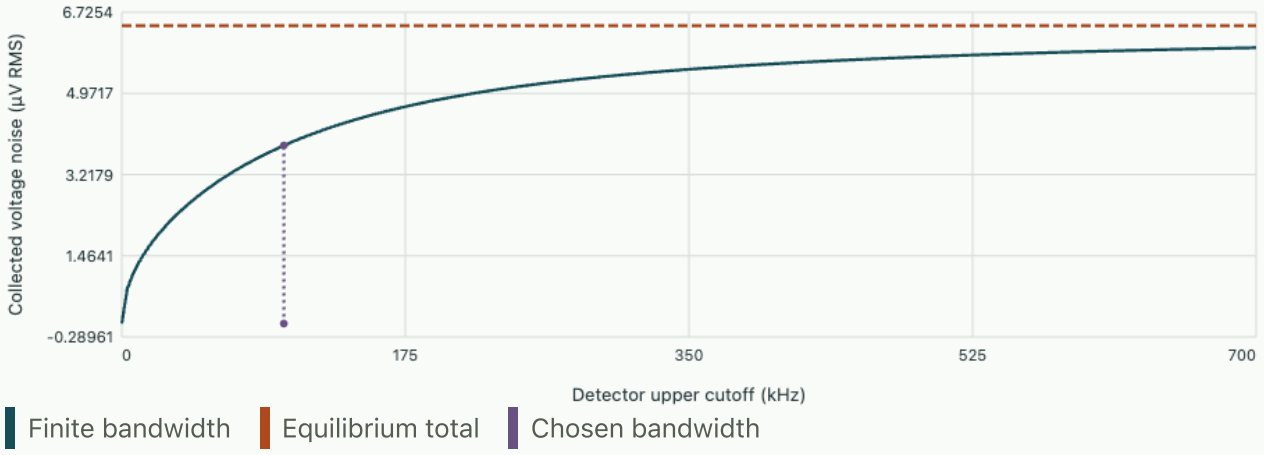

RC corner 159.15 kHz. Detector reads 3.8461 µV RMS, collecting 35.713% of equilibrium variance.

Sweep the resistance by a factor ten. Observe how the corner moves along the frequency axis and how the fraction measured in a fixed bandwidth changes. Then vary temperature. In the classical regime spectral density and integrated variance are proportional to T, while RMS voltage grows only as $\sqrt{T}$. A plot labelled “noise” without units cannot distinguish these two dependencies.

Background · How does a longer voltage exposure alter the measured noise?

An equilibrium resistor–capacitor circuit has voltage covariance $C_v(s) = (k_B T/C)\exp(-|s|/\tau)$, where $\tau = RC$. Suppose a digitiser returns the mean voltage during an exposure T_m. Squaring a sample of this mean measures a different variance from averaging the instantaneous voltage squared. In the classical stationary model the two-time integral becomes

$$\mathrm{var}(\bar{v}) = \frac{2}{T_m^2}\int_0^{T_m}(T_m - s)C_v(s)\,ds = \frac{2k_BT}{C}\,\frac{u - 1 + e^{-u}}{u^2},$$

$$u = \frac{T_m}{\tau}.$$

Short exposures recover k_BT/C. At ten correlation times only about 18% of that variance remains. For very long exposures the variance decreases as $2\tau/T_m$. Taking more closely spaced samples of one exposure cannot undo that averaging.

A physical check and a calibration check

An RC time-domain experiment provides a second observable beyond integrated power. Under the same single-pole equilibrium model the voltage autocovariance is

$$\langle V(t)V(t+\tau)\rangle = (k_BT/C)e^{-|\tau|/(RC)}.$$

Its variance at zero lag agrees with equipartition, and its decay time is obtained from the homogeneous charge-relaxation equation. Repeating the measurement with a known extra parallel capacitance lowers the RMS and lengthens the correlation time. Seeing only one of those changes suggests that the observed noise may include substantial amplifier input voltage or an independent background.

The classical approximation assumes $hf \ll k_BT$ over the relevant band and a resistor that can be treated as frequency independent. A MHz room-temperature experiment is far inside this regime. At sufficiently high frequencies or cryogenic temperatures, a quantum noise model, an explicit definition of the measured ordering of the spectrum and the detector's energy exchange are needed. Extending an ideal classical white source to infinite frequency is only permissible as an approximation inside the convergent low-frequency RC calculation; it is not a microscopic model at arbitrarily high frequency.

Tutorial

At fixed C, T, doubling R in a detector with $B \ll f_c$ changes the measured RMS by what factor? Will the same factor apply when $B \gg f_c$?

▼ **Work through the solution**

At very small bandwidth, $S_C \simeq 4k_BTR$ is nearly constant, so variance doubles and RMS increases by $\sqrt{2}$. At large bandwidth the measurement approaches k_BT/C and becomes independent of R.

A two-sided spectrum in hertz is used instead. What density and integration limits preserve the same variance?

▼ **Work through the solution**

For this real voltage, the density for each sign of frequency is half the one-sided positive-frequency value. Integrate it over $(-\infty, \infty)$. Integrating the one-sided density over both signs would count the power twice.

A measured RMS is 10 µV, including independent amplifier noise of 6 µV RMS. Can 6 µV simply be subtracted to infer the sample contribution?

▼ **Work through the solution**

Independent zero-mean voltages add in variance. The inferred sample RMS is

$$\sqrt{10^2 - 6^2} = 8\ \mu\mathrm{V}.$$

A deterministic offset or correlated noise would require a different treatment.

Compare with an actual laboratory

The MIT Junior Lab noise experiment treats the detector as part of the measurement. Identify its calibration observable before comparing the theoretical density with an instrument reading.

Measuring magnetic core loss by a hysteresis loop and by calorimetry

Electromechanics, optics and measurement

A closed loop on a graph of magnetic induction against field encloses an area with units of energy per unit volume. A thermometer measures kelvins. Bring these two observations onto the same scale, account for the winding heat, and determine when the two experiments actually test the same material loss.

Infer the field and the induction from independent windings

Wind a toroidal specimen with a primary winding of N_p turns and a lightly loaded pickup of N_s turns. Let A be the representative iron cross-sectional area, ℓ the mean magnetic path length and $V_c = A\ell$ the magnetic volume. Neglect a gap and leakage initially. Ampère's law around a path through the core gives

$$H \simeq N_p i_p / \ell.$$

This relation is accurate when the toroid is slender enough that the spread of path radii is small and the pickup current is negligible. In a thick toroid $H(r) = N_p i_p/(2\pi r)$, so one point on a nominal loop can actually represent a range of local field values.

Faraday's law supplies the second observable. Choose winding polarity so the reported voltage is

$$v_s = N_s A \, dB/dt.$$

Integrating the voltage gives the change in cross-section-averaged magnetic induction. The pickup cannot determine a constant induction offset, because static flux induces no voltage. A closed loop area is insensitive to adding a constant to B or H, but a dc offset in the voltage integrates into a time-dependent ramp and prevents the reconstructed trajectory from closing. In steady periodic operation one can subtract the voltage mean over a complete period before integration. Applying that correction to a genuinely evolving magnetisation process would remove part of a real flux change.

Multiply the two measured field quantities in the appropriate order. The field work supplied per unit material volume on a small induction change is $H\,dB$. The source term associated with the magnetising flux is $i_p v_{p,\mathrm{induced}}\,dt$. Using $v_{p,\mathrm{induced}} \simeq (N_p/N_s)v_s$, its value is

$$i_p N_p A \, dB = V_c H \, dB.$$

Integrating around a settled cycle gives the magnetic loss

$$W_c = V_c \oint H \, dB,$$

$$P_c = f V_c \oint H \, dB.$$

The field can absorb positive work and return some of it later in the same cycle. The area measures the algebraic remainder. Reversing the polarity of a single signal flips the loop orientation. A passively heated specimen must receive positive net energy when signs are defined for power entering the core. Simply taking the absolute value of all instantaneous products would count reversible storage twice and can grossly overestimate the heat.

Work an elliptical loop from time-domain signals

Suppose a nearly sinusoidal field is $H = H_p \cos \omega t$ and the induction lags it by a positive angle δ, $B = B_p \cos(\omega t - \delta)$. Differentiate the latter and multiply by H. Over one cycle $\cos \omega t \sin \omega t$ averages to zero, and $\cos^2 \omega t$ averages to one half. The resulting mean power density is $\omega H_p B_p \sin \delta / 2$, giving $\oint H \, dB = \pi H_p B_p \sin \delta$. With zero lag the two signals trace a line and the area vanishes, even if the peak values are large. A saturation-distorted loop need not be elliptical; the original time integral remains valid with synchronised data.

Take $H_p = 120\ \mathrm{A/m}$, $B_p = 0.200\ \mathrm{T}$ and $\delta = 12^\circ$. The energy per cycle and volume is $15.68\ \mathrm{J/m^3}$. A toroid with $\ell = 0.20\ \mathrm{m}$, $A = 2.0 \times 10^{-4}\ \mathrm{m^2}$ and $V_c = 40\ \mathrm{cm^3}$ dissipates $0.627\ \mathrm{mJ}$ per cycle. At $f = 200\ \mathrm{Hz}$ its power is $0.125\ \mathrm{W}$. With $N_p = 100$, the necessary current peak is $0.240\ \mathrm{A}$. With $N_s = 50$, the pickup voltage has peak $N_s A (2\pi f) B_p = 2.51\ \mathrm{V}$. These last two predictions allow the raw oscilloscope traces to be checked before a calculated loss in watts is trusted.

For sampled periodic data use $w = \sum_n (H_n + H_{n+1})(B_{n+1} - B_n)/2$, including the segment from the last point back to the first. This trapezoidal line integral computes an oriented polygon area. Plot the starting and ending points before adding the closing segment. Connecting widely separated endpoints by hand hides drift and can generate a fictitious wedge whose area depends on acquisition timing. Repeat at higher sampling rates to see whether corners or switching features carry appreciable extra area. When a periodic excitation has harmonics the time increment must resolve the highest significant harmonic in both measured signals.

Compare a thermal measurement with the magnetic area

Consider first a calorimeter that receives heat only from the core. Let the heat capacity of specimen, holder and sensor be C_{eff}. Let the leakage conductance to an ambient bath be G, and let $\Delta T = T - T_a$. The thermal balance $C_{\mathrm{eff}} \dot{\Delta T} = P_c - G \Delta T$ predicts a steady offset P_c / G and an initial heating slope P_c / C_{eff}. A small thermometer lag or poor thermal contact can make the earliest measured slope misleading even though environmental losses are then small. A separate electrical heater placed where the core deposits its heat can calibrate the effective thermal response.

For $C_{\mathrm{eff}} = 20\ \mathrm{J/K}$, $G = 0.015\ \mathrm{W/K}$, and a $2.0\ \mathrm{K}$ rise above ambient, the magnetic loss $0.125\ \mathrm{W}$ would give a slope of about $0.00477\ \mathrm{K/s}$. The heat already leaking to the bath is $0.030\ \mathrm{W}$, too large to discard in a ten-percent comparison. If the winding is inside the same calorimeter, its Ohmic heat must enter the balance as well. With primary resistance $1.50\ \Omega$ and the computed sinusoidal current, winding heating is

$$(0.240/\sqrt{2})^2 (1.50) = 0.0432\ \mathrm{W}.$$

The expected slope of the combined package at the same temperature is now $0.00693\ \mathrm{K/s}$. Subtraction of winding heat recovers the core result only when the estimated resistance represents the winding at its operating temperature and frequency.

A thermal estimate drawn from a whole interval is often more robust than a derivative from two adjacent temperature samples. Integrating gives

$$Q_{\mathrm{deposited}} = C_{\mathrm{eff}} [\Delta T(t_2) - \Delta T(t_1)] + G \int_{t_1}^{t_2} \Delta T \, dt.$$

Divide by the duration for mean power, then remove other measured heat sources. A drifting ambient must be recorded explicitly because the calorimeter stores energy relative to its own absolute temperature. If the core and coil equilibrate slowly with each other, a single fitted heat capacity can depend on the fitting time window. A two-temperature model or a longer dwell then supplies an independent diagnostic.

Separate reversible exchange from hysteretic and induced currents

A single measured loop at finite frequency gives the total energy dissipated by the core at that prescribed waveform and temperature. The loop can include irreversible domain processes, eddy currents and other rate-dependent effects. It does not by itself split these contributions. To probe their different scaling hold the peak induction and waveform shape fixed, then change frequency. Under a restricted low-frequency model the quasi-static hysteresis energy per cycle is approximately constant, while classical thin-lamination eddy-current power scales as $f^2 B_p^2$. Its energy per cycle therefore scales roughly as f. Plotting the measured loss per cycle and volume against f tests whether such an intercept-plus-slope description fits the chosen range.

Suppose the area is $12\,\mathrm{J/m^3}$ at $100\,\mathrm{Hz}$ and $20\,\mathrm{J/m^3}$ at $300\,\mathrm{Hz}$, at the same peak induction. A straight line gives intercept $8\,\mathrm{J/m^3}$ and slope $0.040\,\mathrm{J\,s/m^3}$. It predicts $16\,\mathrm{J/m^3}$ at $200\,\mathrm{Hz}$, near the elliptical-loop calculation. Two points always determine a line. A third and fourth frequency, together with uncertainties, are required to test the law. Different losses can share similar frequency scaling over a short interval, and temperature-dependent conductivity can bend a nominally simple curve.

Use the pickup to reject the reversible part of terminal power

Measuring primary terminal voltage and current with a power analyser provides a separate electrical route to loss, but the wanted number can be a small difference between terms. The applied voltage contains winding resistance, induced magnetising voltage, leakage-inductance voltage and possible parasitic contributions. In steady periodic operation the energy stored in a fixed linear leakage inductance returns after each cycle, while winding resistance dissipates energy permanently. The secondary pickup preferentially measures the core-linked flux. Forming $p_c(t) = (N_p/N_s) i_p(t) v_s(t)$ and averaging over a full cycle estimates core power without directly including the primary's copper voltage drop. It relies on both windings linking the same relevant flux and on negligible power drawn by the pickup.

For the elliptical example the primary magnetising voltage has peak $5.03\,\mathrm{V}$ and leads the induction by ninety degrees. Its phase relative to the current is

$$90^\circ - \delta = 78^\circ.$$

The core branch's apparent power is $V_{\mathrm{rms}} I_{\mathrm{rms}} \approx 0.603\,\mathrm{VA}$; its mean dissipative part is about $0.125\,\mathrm{W}$, and the reactive part is about $0.590\,\mathrm{var}$. Only the mean part contributes net heating of the magnetic volume. Including the winding voltage adds about $0.043\,\mathrm{W}$ of real power to the terminal reading. A simple product of peak current and peak secondary voltage, without the turns ratio, factor one half and phase factor, has no reason to agree with either measured loss.

Consider a nearly ideal core with the same peak signals but a lag of only one degree. Its loss is about $10.5\,\mathrm{mW}$, while the much larger reversible power exchange is almost unchanged. A time offset of $10\,\mu\mathrm{s}$ at $200\,\mathrm{Hz}$ adds 0.72°, causing a relative loss error of approximately 72% for that small lag. Calibrate the relative channel delays using a reference signal through the actual current transducer and voltage chain, including

their filters. Adjusting traces on screen until a core loop appears narrow assumes the very loss one is trying to measure. A synchronous measurement over many periods suppresses random noise but does not remove a persistent delay bias.

Test reconstruction on a waveform with a third harmonic

Let $H(t) = H_1 \cos \omega t + H_3 \cos 3\omega t$ while $B(t) = B_1 \cos(\omega t - \delta_1) + B_3 \cos(3\omega t - \delta_3)$. With these definitions δ_n is the phase lag of the nth harmonic in its own sinusoid. Multiplying H by $\dot{B}$ and integrating over an integer number of fundamental periods eliminates products of different harmonics. Each surviving harmonic contributes the number of its own cycles in one fundamental period. The area per fundamental period is therefore $\pi[H_1 B_1 \sin \delta_1 + 3 H_3 B_3 \sin \delta_3]$. This identity tests whether a numerical area calculation tracks the full waveform. It does not assume that a real material is linear; here the two signals are supplied as measured periodic functions.

Use the earlier fundamental values and add $H_3 = 20\ \mathrm{A/m}$, $B_3 = 0.030\ \mathrm{T}$ and $\delta_3 = 30^\circ$. The third harmonic contributes $2.83\ \mathrm{J/m^3}$ in each fundamental period. This represents an additional $0.0226\ \mathrm{W}$ for the chosen specimen and frequency. A fundamental-only sine fit would miss nearly fifteen percent of the resulting total area. An apparent jagged loop can also arise from asynchronous sampling or an offset. Distinguish genuine harmonics by checking that they are phase locked to the drive and persist under a change in sampling rate and analogue bandwidth.

Sweep phase and read the oriented loop

The live experiment reconstructs the H, B trajectory and compares it with a zero-lag reference. It uses the same peak field and induction for the two curves. At fixed peaks changing frequency leaves an assigned phase lag and therefore the prescribed energy per cycle unchanged, while watts rise with the number of cycles per second. A real core can change its lag and waveform with frequency; the graph separates this independent variable so that the meaning of the two scaled quantities is clear. Add a timing offset and predict the direction of the area error before checking its reported size. Near zero lag even the direction of an apparent loop can be controlled by a small electronic offset.

Open the live experiment

Try the calculation

Physical induction lag (°) **12**

Frequency (Hz) **200**

Induction channel delay (µs) **0**

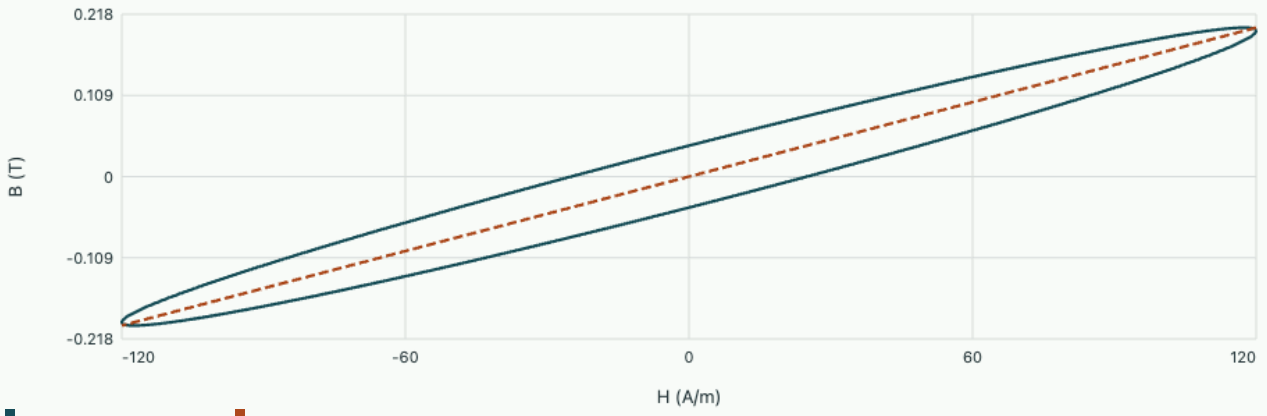

True loop area 15.68 J/m³ per cycle. The recorder indicates 15.68 J/m³. Physical heating of the 40 cm³ specimen is 0.125 W.

Two toroids are made of the same nominal material with equal cross sections but different mean path lengths. How should a result reported only as watts be compared?

▼ **Work through the solution**

Determine actual magnetic volume for each toroid and calculate the energy per unit volume per cycle, $P/(fV_c)$. Also check that both experienced the same induction waveform, frequency and temperature. Equal primary currents need not give equal field when the path lengths differ; equal primary applied voltages need not give equal core induction if turns, winding voltage drops or area differ. Reproducing the actual $B(t)$ with feedback on the pickup is a more direct way to match excitation. Report dimensions and peak induction with the result so a reader can reconstruct what material condition was tested.

A saturated sample shows a small loop during a biased excitation that never crosses zero induction. Does the area formula still apply?

▼ **Work through the solution**

A settled periodic minor loop encloses an area representing the net cyclic energy under exactly that bias and excursion. Adding a static induction offset does not change $\oint H\,dB$, but the bias can alter the trajectory, differential permeability, and width of the minor loop itself. Do not infer the full major-loop energy from its area. The pickup returns only changes of induction unless an independent dc measurement or a known magnetic preparation supplies the offset. A bias winding may additionally heat the calorimeter, so include its measured power when comparing with thermal loss.

Tutorial on comparing two loss measurements

A data logger reports H_p and B_p correctly but adds an uncorrected $10\,\mu\mathrm{s}$ delay to the induction channel. Estimate the fractional area error of the elliptical example.

▼ **Work through the solution**

At $200\,\mathrm{Hz}$ the extra apparent lag is

$$2\pi f \Delta t = 0.01257\,\mathrm{rad} = 0.72^\circ.$$

The reported area is multiplied by $\sin(12.72^\circ)/\sin(12^\circ)$, about 1.059. An uncorrected delay of only one hundred-thousandth of a second overstates the core loss by about six percent. For a very low-loss specimen with much smaller physical lag, the same delay produces a far larger relative error.

A thermal run on the combined winding and core reports a slope $0.0070\,\mathrm{K/s}$ at $\Delta T = 2.0\,\mathrm{K}$. Recover the core power and the volume-specific loss area at $200\,\mathrm{Hz}$.

▼ **Work through the solution**

Heat deposition in the package is

$$20(0.0070) + 0.015(2.0) = 0.170\,\mathrm{W}.$$

Subtract the estimated $0.0432\,\mathrm{W}$ of winding heating to obtain $0.1268\,\mathrm{W}$ in the core. Dividing by $200 \times 4.0 \times 10^{-5}$ gives $15.85\,\mathrm{J/m^3}$ per cycle. Its closeness to $15.68\,\mathrm{J/m^3}$ should be evaluated against the uncertainties of heat capacity, thermal conductance, resistance, dimensions and electrical phase alignment.

Vary the magnetic circuit and identify the induction scale used by the specimen.

Optical trapping and Brownian calibration of a dielectric bead

Electromechanics, optics and measurement

An image of a bead remaining near a focus does not establish the strength of its trap. Derive the force pulling a small dielectric sphere into a bright region, compare it with the forward force from scattering and learn what a motion record can tell you about stiffness, drag and temperature. The model here uses an inert dielectric sphere in a homogeneous Newtonian liquid.

From two dielectric boundaries to the energy of an induced dipole

A sphere of radius a and relative permittivity ϵ_b sits in a liquid of relative permittivity ϵ_m. Assuming the sphere is small compared with a wavelength in the liquid, solve its local electric problem in an almost uniform applied field E. Rotational symmetry about the field admits an interior potential proportional to $r\cos\theta$. The exterior consists of the imposed potential $-Er\cos\theta$ and an induced dipolar term proportional to $\cos\theta/r^2$. Equality of potential and of normal displacement at the surface gives the local-field and dipole factors

$$E_{\mathrm{in}} = \frac{3\epsilon_m}{\epsilon_b + 2\epsilon_m} E,$$

$$p = \alpha E,$$

$$\alpha = 4\pi\epsilon_0\epsilon_m a^3 \frac{\epsilon_b - \epsilon_m}{\epsilon_b + 2\epsilon_m}.$$

The field doing the polarising work is the external field in the host at the particle position before inserting it. To turn that work into a potential, build up the dipole from zero rather than assuming a permanent dipole moment. Quasistatic assembly gives

$$-\int_0^E \alpha E'\, dE' = -\alpha E^2/2.$$

For a harmonic peak field average the square over an optical period. In a nearly lossless liquid of refractive index n_m, the peak field is related to travelling-wave intensity by $I = n_m\epsilon_0 c|E|^2/2$. The resulting conservative force is

$$U = -\frac{\mathrm{Re}\,\alpha}{4}|E|^2 = -\frac{\mathrm{Re}\,\alpha}{2n_m\epsilon_0 c} I,$$

$$\mathbf{F}_{\mathrm{grad}} = -\nabla U.$$

The two factors of two have different origins: induced-dipole assembly and time averaging of a sinusoid. Using a permanent-dipole potential with a peak field would miss both. A bead of higher refractive index than the liquid has positive real polarizability and is attracted to increasing intensity. A lower-index inclusion reverses this local tendency. For large particles, strong absorption or rapidly structured fields, a dipole with real static polarizability no longer specifies the force; a field or multipole calculation of the complete momentum balance is required.

Investigate the small-sphere dipole response and its dependence on particle size

Compute a transverse stiffness and test how local it is

Assume the beam near its focal plane has intensity $I(x) = I_0 e^{-2x^2/w^2}$, with power P, waist w defined at $1/e^2$ intensity and $I_0 = 2P/(\pi w^2)$. Put $U_0 = \mathrm{Re}\,\alpha I_0/(2n_m\epsilon_0 c) > 0$. The potential measured above its minimum is $U_0[1 - e^{-2x^2/w^2}]$. Expanding to second order and equating its leading term to $\kappa_x x^2/2$ gives

$$\kappa_x = 4U_0/w^2.$$

Stiffness therefore scales linearly with power, as the third power of sphere radius and as the inverse fourth power of waist at fixed indices. Each scaling applies within the stated approximations.

For a 40 nm radius bead with $n_b = 1.59$, take $n_m = 1.33$, $w = 0.65$ micrometres, vacuum wavelength 1064 nm, $P = 90$ mW and $T = 298$ K. The index contrast factor is 0.1252 and the waist intensity is 1.36×10^{11} watts per square metre. The predicted depth 3.03×10^{-20} J is 7.36 times the thermal energy. The lateral stiffness is 2.87×10^{-7} N/m. A displacement of 50 nm then requires approximately 14.3 fN of transverse balancing force. At a displacement of half a waist, the exact restoring force is the harmonic estimate multiplied by $e^{-1/2} = 0.607$; a Hookean calibration cannot be extrapolated that far.

Open the live experiment

Try the calculation

Bead radius (nm) **40**

Optical power (mW) **90**

Beam waist (µm) **0.65**

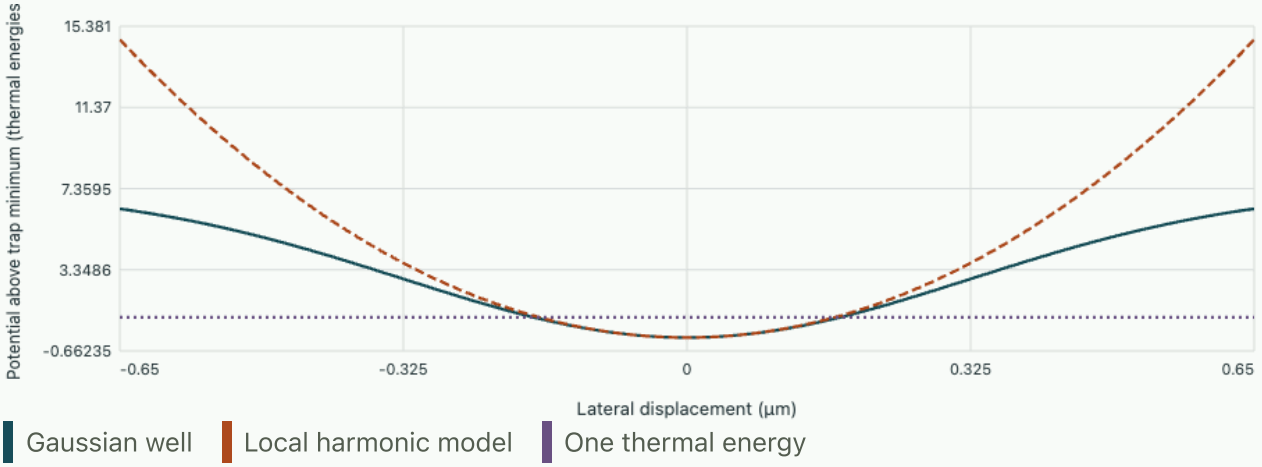

The well depth is 7.36 kBT. The local model gives an rms displacement of 119.8 nm and a corner frequency of 60.5 Hz at 298 K.

Use the two curves to identify the interval in which the local quadratic energy is within ten percent of the Gaussian prediction. First lower power at fixed waist. Both curves shrink together in energy units while the harmonic thermal displacement grows. Next increase bead radius at fixed power and watch the well deepen, but also evaluate the size parameter $2\pi n_m a/\lambda_0$. It is 0.314 at the default and grows linearly with radius. Stronger trapping predicted by this simplified model does not establish the accuracy of the dipole approximation at arbitrary sizes.

Compare axial restoring force with radiation pressure

An induced dipole reradiates a portion of the incident field. In a uniform host use the Rayleigh scattering cross section $\sigma_s = (8\pi/3)k_m^4 a^6 |(\epsilon_b - \epsilon_m)/(\epsilon_b + 2\epsilon_m)|^2$. In a common travelling-wave momentum-flux model for a bead in a stationary nondispersive liquid the axial scattering force is approximately $F_s = n_m I \sigma_s / c$; absorption would add a further contribution. This estimate uses the total scattering cross section because the leading dipole pattern has vanishing mean forward cosine. For a larger forward-scattering sphere a transport cross section weights the deflected momentum instead.

In the default case the cross section is roughly 2.05×10^{-18} square metres. The forward force at the waist is about 1.23 fN. A Gaussian axial profile has $I(0,z) = I_0/[1 + (z/z_R)^2]$ with $z_R = \pi n_m w^2/\lambda_0 = 1.66$ micrometres. Its local axial stiffness $2U_0/z_R^2 = 2.20 \times 10^{-8}$ N/m is much smaller than its transverse stiffness. Balancing the constant forward force with that local stiffness predicts an equilibrium about 56 nm downstream of the waist. The small displacement relative to the Rayleigh length supports the linearisation.

Both the axial stiffness and forward scattering force scale with power in this simple model. A power sweep can thus increase confinement while barely shifting the equilibrium offset, a useful discriminating prediction. At fixed focus the offset grows rapidly with bead size because the scattering force scales with the sixth power and the gradient force with the third power. A stable lateral position alone does not establish stable three-dimensional capture. Check the axial restoring slope at a force zero and verify that thermal escape and fluid flow are negligible over the acquisition interval.

Extract stiffness from fluctuations or from a relaxation time

Near a stable point in one Cartesian direction, let x denote displacement from the mean and ζ the viscous drag. At sufficiently low frequencies the inertial term is negligible and the overdamped Langevin equation is

$$\zeta \dot{x} = -\kappa x + \xi.$$

At equilibrium the fluctuating force has zero mean and white covariance $2k_B T \zeta \delta(t - t')$ over the frequencies considered. The deterministic relaxation time is

$$\tau = \zeta/\kappa.$$

Multiplying the equation at positive delay by $x(0)$ and averaging removes the uncorrelated later noise and gives the decaying covariance. Equipartition sets its initial value.

$$C(t) = \frac{k_B T}{\kappa} e^{-|t|/\tau},$$

$$f_c = \frac{\kappa}{2\pi\zeta},$$

$$S_x^{(1)}(f) = \frac{4 k_B T \zeta}{\kappa^2 + (2\pi f \zeta)^2}, \quad f \geq 0.$$

The stated spectrum is one-sided in ordinary frequency. Its integral from zero to infinity equals $C(0)$; doubling another time would double the variance. In an unbounded fluid with viscosity 1.00 mPa s, Stokes drag is $6\pi\eta a = 7.54 \times 10^{-10}$ N s/m. The default predicts $f_c = 60.5$ Hz and rms displacement 120 nm. The latter is roughly 0.18 waist, so the outer part of a long thermal trajectory may sample appreciable

anharmonicity. A finite-depth well in an unbounded bath does not have a globally normalisable trapped equilibrium. These local equilibrium predictions apply to a sufficiently long-lived, locally sampled trapped state.

If a voltage detector gives $V = gx$, the spectral corner and normalised autocorrelation time do not depend on the displacement gain g. They determine only κ/ζ. Equipartition using a calibrated spatial variance gives κ independently of drag, provided mechanical fluctuations are in equilibrium and sufficiently resolved. If the viscosity or nearby-wall drag is uncertain, combining a known stage oscillation with the Brownian spectrum can determine further calibration quantities. Keep a stage-driven coherent peak out of the noise integral used for equipartition.

Account for exposure, sample size and unobserved motion

A camera typically reports the position averaged during an exposure of duration t_e. Its variance is the double integral of the covariance over the exposure interval divided by t_e^2. Evaluating the exponential local model and introducing $u = t_e/\tau$ gives a multiplicative factor $2(u - 1 + e^{-u})/u^2$. At $u = 1$ the reported physical-motion variance is only 0.736 of the instantaneous variance. Interpreting it as the full variance would overestimate stiffness by approximately 36 percent. Independent pixel-localisation noise adds additional variance with the opposite sign of bias.

Estimate that localisation noise with an effectively immobilised reference bead at a comparable brightness. Test the averaging correction by changing exposure duration at constant sampling frequency where possible. A camera that reports overlapping exposures has correlated measurement errors and needs a different covariance model. A long series of closely spaced frames also contains fewer independent thermal configurations than its frame count suggests. Compare estimates across several blocks substantially longer than the relaxation time; residual low-frequency drift can dominate even when high-frequency position precision looks excellent.

A separate check uses a known uniform liquid flow. At a steady small displacement x_s, force balance gives

$$\kappa x_s = \zeta u_{\mathrm{flow}}.$$

For the default drag and stiffness a flow of 10 micrometres per second gives an offset of 26.3 nm. Repeat with reversed flow and use half the displacement difference to cancel an unknown zero-force position. Near a wall, the local flow at the bead and wall-corrected drag may differ from a stage velocity and the unbounded Stokes value. Those effects alter the mechanical calibration even if the optical field is unchanged.

A known stage oscillation calibrates position in a fluid

Brownian fluctuations alone determine the product of several unknown calibration parameters. A weak sinusoidal translation of the sample cell adds a known mechanical signal at a chosen frequency. Work in laboratory coordinates with a fixed optical trap and with fluid moving uniformly at velocity $u(t)$ near the bead, well away from walls in the idealised model. The drag is now $-\zeta(\dot{x} - u)$. Superposing a stage displacement $A\cos\omega t$ on thermal fluctuations gives a deterministic response whose complex amplitude satisfies $(\kappa - i\omega\zeta)\widetilde{x} = -i\omega\zeta A$ in the negative-exponential convention. Therefore its amplitude ratio is

$$\frac{|\widetilde{x}|}{A} = \frac{\omega\tau}{\sqrt{1+\omega^2\tau^2}}.$$

At the trap corner the bead displacement amplitude is the known stage amplitude divided by the square root of two. A calibrated 80.0 nm stage oscillation there produces a predicted bead amplitude of 56.6 nm. If a position-sensitive receiver records a coherent peak amplitude of 0.170 V, its displacement gain is 3.01 mV/nm. At very low drive frequencies the fluid velocity is small for fixed stage amplitude and the trap holds the bead nearly stationary. At high enough frequency the bead follows the fluid displacement within the overdamped model, but finite hydrodynamic, optical and electronic bandwidths eventually matter.

After determining the gain, convert the surrounding voltage spectrum into spatial units and use the area of the stochastic fluctuation spectrum to find stiffness. Exclude the calibrated line and account for the finite spectral averaging window; a narrow drive that leaks into neighbouring bins can inflate the inferred thermal variance. Dividing measured stiffness by the independently determined corner rate then gives drag, which can be compared with the known viscosity and sphere radius. This avoids assuming free-space Stokes drag in the first step of a near-wall measurement.

There are two limiting systematic effects that a good fit to a single drive frequency cannot resolve. A frequency-dependent stage calibration changes the actual fluid displacement, and motion of the optics relative to the cell generates a direct detector background even with the bead immobilised. Repeat the drive above and below the corner, compare coherent phase as well as amplitude and acquire a trapped and immobilised reference. Correct conversion should explain the entire complex response with one spatial gain and one corner time. A single calibration number obtained at one frequency can hide a phase lag in the stage controller or in the optical position channel.

One last cross-check reverses the stage drive phase by half a period. The coherent bead response must reverse, while thermal background power remains unchanged. A component that fails this reversal could come from ambient vibration unrelated to the stage. Averaging the two traces with and without a sign change separates these contributions before either is used in an optical-force calibration.

Tutorial · Separate a unit-conversion error from excess noise

In a single-axis trap at 298 K a calibrated detector reports a variance of 2500 square nanometres after removal of localisation noise and exposure correction. An independent spectral fit gives a corner of 175 Hz. Infer stiffness, drag and the equivalent unbounded hydrodynamic radius for a liquid of viscosity 1.00 mPa s. Compare against a nominal 80 nm radius bead.

▼ Infer the three quantities and test their consistency

A nanometre squared is 10^{-18} square metres, so the variance is 2.50×10^{-15} square metres. Equipartition gives 1.65×10^{-6} N/m. Dividing by $2\pi f_c$ gives drag 1.50×10^{-9} N s/m and radius $\zeta/(6\pi\eta) = 79.5$ nm. The agreement supports the local overdamped description under the assumed temperature and viscosity. Forgetting the factor of 2π would predict a radius roughly six times larger and could be mistaken for a huge near-wall correction.

Now imagine that additional uncorrelated read noise had been left in the variance. This raises apparent variance, lowers the equipartition stiffness and would lower the inferred drag at the same fitted corner.

Nearby-wall drag for a given radius instead normally raises hydrodynamic resistance for parallel translation. Their effects on the inferred radius thus differ in sign under the stated assumptions. Varying height and independently calibrating the measurement noise are appropriate diagnostics before attributing a discrepancy to bead dimensions.

Tutorial · Test a force-law inference with power and concentration controls

At fixed bead radius a trap is operated at two powers differing by a factor of two. Measured corner frequencies increase from 42 to 84 Hz, but rms displacement falls only from 110 to 87 nm. What does the harmonic model with fixed drag and temperature predict? Estimate a power-independent additive position-noise variance that could reconcile the two readings.

▼ **Use two variances to separate scale and offset**

The doubled corner gives doubled stiffness if drag is unchanged. Physical thermal variance should halve. The predicted rms at higher power would be 77.8 nm if the lower-power value were entirely physical. Let the low-power measured variance be $v + b$ and the high-power one $v/2 + b$, where b is a constant additive noise variance. Solving from 12,100 and 7,569 square nanometres gives $v = 9062$ and $b = 3038$ square nanometres, or approximately 55.1 nm rms noise.

This is a large inferred detector noise and demands an independent immobilised reference check. Power-dependent heating of the liquid changes temperature and viscosity, so an unchanged product inferred from only two quantities cannot prove that heating is absent. Test several powers, reverse their order to reveal slow drift and vary bead concentration to exclude overlapping particles in the tracking window. The inferred additive term should remain constant across those controls if it comes from the specified detector model.

Tolić-Nørrelykke and colleagues · position and force calibration using thermal and driven measurements

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.

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**

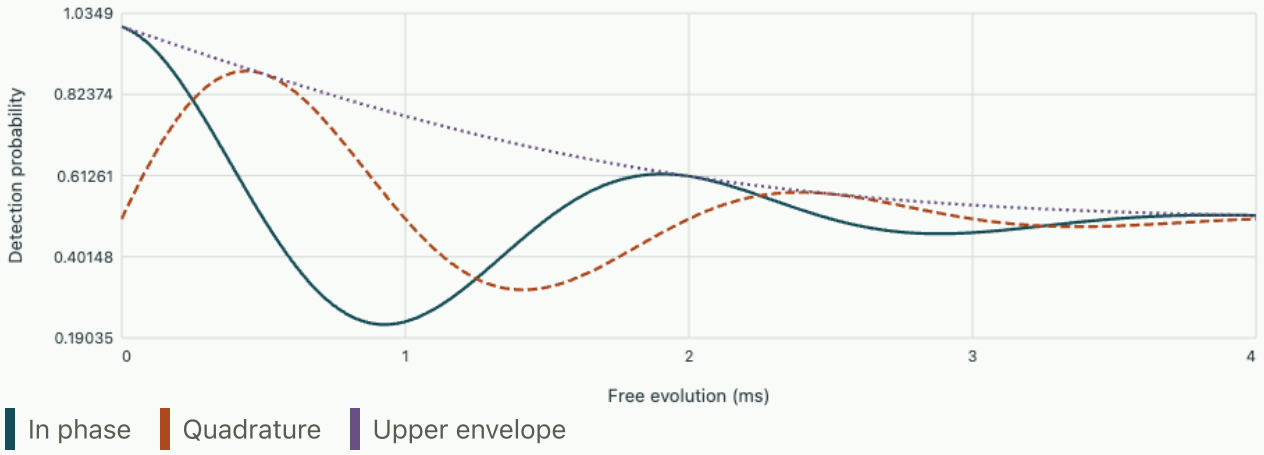

Contrast after 1 ms = 0.53455. The two probabilities at 0.50 ms are 0.5 and 0.87729.

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 $-1 \le c \le 1$ 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 $t \ge 0$ 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 $t \ge 0$. 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

Recover a transition frequency from two separated pulses

Computational quantum experiments

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**

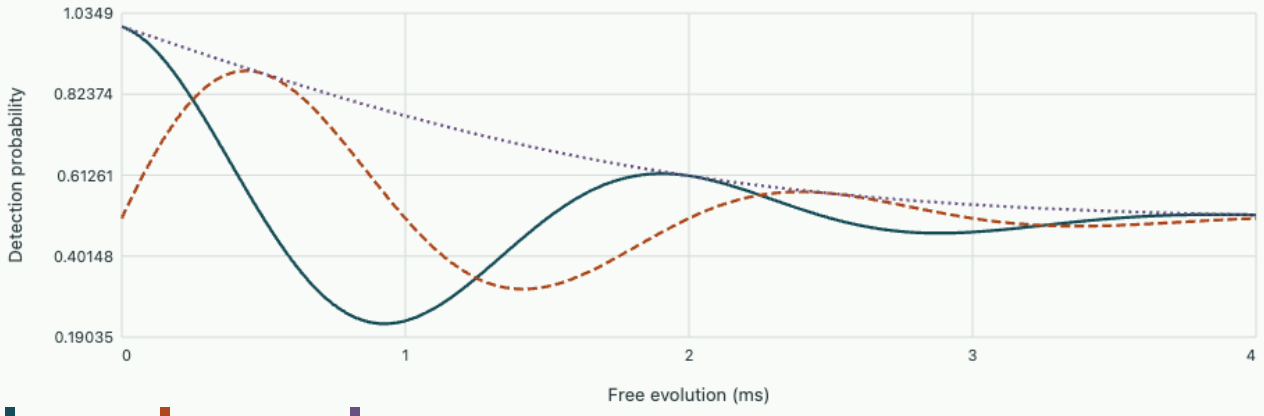

Contrast after 1 ms = 0.53455. The two probabilities at 0.50 ms are 0.5 and 0.87729.

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.

Ensembles, diffusion and finite data

Contrast conservation along a microscopic orbit with apparent relaxation after coarse sampling. Test a diffusion law and infer an unknown thermal state from a sequence of noisy observations.

1. Liouville flow and what an ensemble can remember
2. Time-dependent diffusion
3. What a periodic grid does to a decaying Fourier mode
4. Splitting a stiff relaxation from a slowly varying drive
5. Learning a drifting temperature from a sequence of readings
6. Predicting a new measurement when the fitted parameters are correlated
7. Symplectic transformations, Poisson brackets and the geometry of measurements
8. Action, angle and the slowing clock at a separatrix
9. Action density, stress waves and the boundary that reflects them

Liouville flow and what an ensemble can remember

Classical mechanics

A cloud of identical noninteracting beads is released with slightly different initial positions and momenta. Its position distribution widens. Does the system occupy more phase space, or has its uncertainty merely changed shape?

Follow a small phase-space element

For one canonical degree of freedom the velocity of a point (q, p) is $(\partial_p H, -\partial_q H)$. A short-time map takes $\boldsymbol{z}$ to $\boldsymbol{z} + \boldsymbol{v}(\boldsymbol{z}, t)\, dt$. Its Jacobian matrix is $I + (\partial \boldsymbol{v}/\partial \boldsymbol{z})dt$; expanding the determinant gives $1 + (\nabla_z \cdot \boldsymbol{v})dt + O(dt^2)$. Hamilton's equations make the divergence zero because the two mixed second derivatives of a smooth Hamiltonian cancel. Each additional canonical pair contributes a similar cancellation. A time-dependent Hamiltonian may change energy but still preserve the instantaneous phase-space volume along the flow.

Conservation of the number of ensemble members in a moving phase-space box also gives the continuity equation $\partial_t f + \nabla_z \cdot (f\boldsymbol{v}) = 0$. The vanishing divergence leaves $df/dt = 0$ on each trajectory. A stationary mechanical system can admit $f = F(H)$ as a stationary distribution because trajectories stay on energy surfaces. Such stationarity alone does not imply that one trajectory explores an entire energy surface. For example, further integrals of motion restrict regular multidimensional trajectories.

Propagate a cloud released into free flight

For mass m, the exact map is $p = p_0,\ q = q_0 + t p_0/m$. Its determinant is one regardless of time. An initially uncorrelated cloud with variances s_q^2, s_p^2 has evolved covariance $C_{qp} = t s_p^2/m$. The new position variance is $s_q^2 + (t s_p/m)^2$, while the momentum variance stays fixed. Compute the determinant of this covariance matrix:

$$\sigma_q^2 \sigma_p^2 - C_{qp}^2 = s_q^2 s_p^2.$$

For a Gaussian the ellipse at a fixed multiple of the covariance radius thus retains its area while it shears. With $m = 1\,\mathrm{kg}$, $s_q = 1\,\mathrm{m}$, $s_p = 0.5\,\mathrm{kg\,m/s}$, and $t = 4\,\mathrm{s}$, the observed position standard deviation is $\sqrt{5} = 2.236\,\mathrm{m}$. The momentum standard deviation remains 0.5 and the covariance becomes $1.0\,\mathrm{kg\,m^2/s}$. Their uncorrected width product is 1.118 in action units. The square root of the covariance determinant is still 0.5. Multiplication of measured marginal widths alone falsely suggests an area increase.

Compare reversible shear and real drag

Now apply a linear drag $\dot{p} = -\gamma p$ to each bead with constant $\gamma > 0$. Its phase-space divergence is $-\gamma$, hence the map determinant after time t is $e^{-\gamma t}$. Direct integration gives

$$p = e^{-\gamma t} p_0,\ q = q_0 + (1 - e^{-\gamma t}) p_0/(m\gamma).$$

The cloud now genuinely loses area in the reduced bead description. A larger Hamiltonian model of bead plus environment can remain volume preserving. Choose a drag rate in the graph and compare the original, freely

sheared and damped ellipses for the same observation time. Check there is a smooth zero-drag limit rather than substituting zero into a quotient.

Open this experiment online

Try the calculation

Time after release (s) **4**

Drag rate γ (1/s) **0.2**

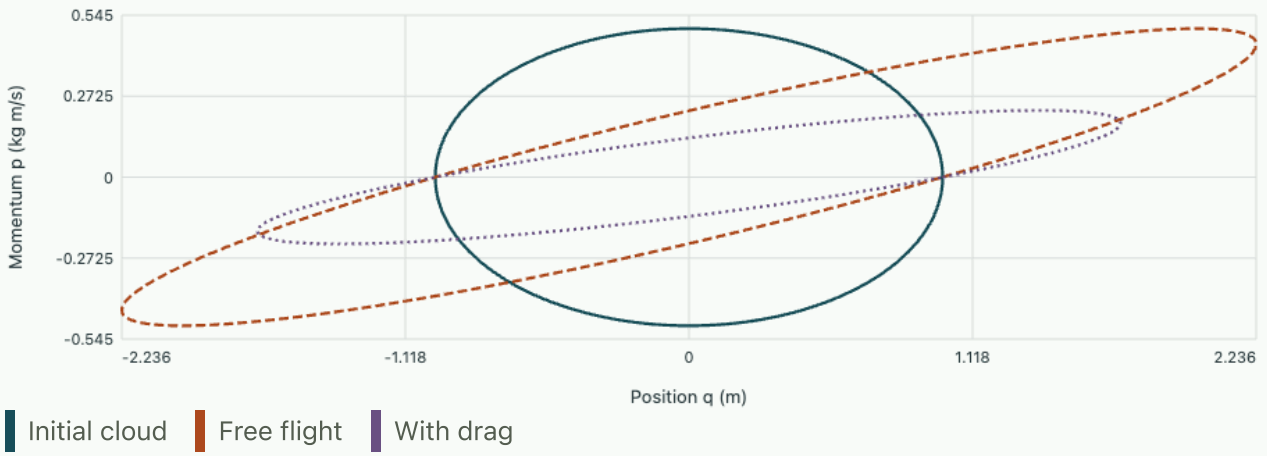

Free position spread 2.236 m. The damped ellipse retains 44.93% of its initial area and has position spread 1.702 m.

Tutorial on distinguishing a shear from dissipation

A harmonic oscillator maps its scaled quadratures by a rotation. Show why an initially anisotropic cloud can periodically sharpen in position without phase-space cooling.

▼ **Work through the solution**

Write its covariance as $R\Sigma_0 R^T$. Since $\det R = 1$, its determinant is unchanged. A rotation of one quarter turn exchanges the original widths, so position can sharpen when the original momentum width in scaled units was smaller. Energy extraction is unnecessary.

At $t = 3\,\mathrm{s}$ the measured ellipse area is 55 percent of the initial area in the linear-drag model. Infer γ. Which independent observable checks the inference?

▼ **Work through the solution**

$$\gamma = -\ln(0.55)/3 = 0.199\,\mathrm{s}^{-1}.$$

The ratio of final to initial momentum standard deviations should also be 0.55 for drag without random forcing. Added environmental fluctuations would modify that prediction.

Further reading · Cambridge classical dynamics, Hamiltonian flow and Liouville's theorem.

Time-dependent diffusion

Numerical partial differential equations

Derive the diffusion time-step restriction from two independent arguments and compare explicit and implicit decay of spatial modes.

FROM FIRST PRINCIPLES

Derive the explicit heat step and its grid-dependent time scale

Take the heat equation $T_t = \alpha T_{xx}$ with constant positive diffusivity on a uniform one-dimensional grid. Fourier conduction through each side of a cell produces a temperature update proportional to the difference of the two neighbour gradients. With forward Euler timing, collect those fluxes to obtain

$$T_i^{n+1} = T_i^n + r(T_{i-1}^n - 2T_i^n + T_{i+1}^n),$$
$$r = \alpha \Delta t / h^2.$$

When $0 \le r \le 1/2$ the new temperature is a convex average of three old ones, preserving bounds on a source-free grid. To check whether a higher step could still be stable for all periodic grid disturbances, substitute $T_j^n = g^n e^{ij\theta}$, where j is the node index. The amplification simplifies to

$$g = 1 + r(e^{-i\theta} - 2 + e^{i\theta}) = 1 - 4r\sin^2(\theta/2).$$

The alternating mode at $\theta = \pi$ gives $g = 1 - 4r$, so $|g| \le 1$ again requires $r \le 1/2$. Halving spatial spacing therefore quarters the largest explicit time step for this discretization. Implicit stepping changes this restriction but still has to resolve any transient relevant to the measurement.

DERIVATION AND APPLICATIONS

Can an energy-stable heat step create an alternating transient?

A diffusion mode should decay rather than oscillate. To distinguish decay of a norm from preservation of sign, examine the finite-element or finite-volume semidiscrete system $M\dot{u} + Ku = 0$, with symmetric positive definite mass matrix M and thermal stiffness K positive semidefinite. An eigenvector satisfying $Kv = \lambda Mv$ has scalar amplitude $y(t)$ obeying $y' = -\lambda y$. A generalized θ step applied to this mode gives

$$[1 + \theta z]y^{n+1} = [1 - (1-\theta)z]y^n,$$
$$z = \Delta t\, \lambda.$$

For Crank–Nicolson take $\theta = 1/2$. The factor $G = (1 - z/2)/(1 + z/2)$ has $|G| \le 1$ for all $z \ge 0$ but changes sign when $z > 2$. An initial grid-scale component can therefore alternate in time. To test a remedy, replace the first full step by two half backward-Euler steps. Their combined factor is

$$G_{\text{start}} = (1 + z/2)^{-2}.$$

Thereafter return to the second-order Crank–Nicolson update. After n further full steps the mode amplitudes, with and without the initial damping, are

$$y_{\mathrm{damped}} = G_{\mathrm{start}} G^n y^0,$$
$$y_{\mathrm{plain}} = G^{n+1} y^0.$$

For small z, both startup strategies agree through first order, while for large z the two half steps strongly suppress the unresolved mode. This argument explains a startup treatment for nonsmooth initial data; sign preservation of every physical grid value is a stronger matrix property than sign preservation of one mode.

WORKED EXAMPLE

Check damping of a near-grid-scale eigenmode

Take $\lambda = 2000$ s⁻¹ and $\Delta t = 0.0100$ s. Then $z = 20$. A single Crank–Nicolson step leaves $-9/11 = -0.818$ of the amplitude; two backward-Euler half steps leave $1/121 = 0.00826$. The exact amplitude factor is

$$e^{-20} = 2.06 \times 10^{-9}.$$

Both discrete results are bounded, but the startup treatment removes almost all of a mode that the time grid cannot resolve.

SECOND APPLICATION

Check the cost for a physically slow component

For the same mesh and step, a slow mode with $\lambda = 2.00$ s⁻¹ has $z = 0.0200$. Crank–Nicolson gives $0.99/1.01 = 0.980198$, the startup gives $1/1.01^2 = 0.980296$ and the exact factor is 0.980199. The slow component changes by less than 10^{-4} in this comparison. A timestep must still resolve a transient that is itself part of the desired output.

Discretise space and time separately

A heat pulse in a thin insulated rod smooths over time according to $u_t = D u_{xx}$ for constant positive diffusivity D. Use a uniform grid of spacing h, and initially examine interior points or periodic boundaries so the stencil is unchanged everywhere. Replacing the second derivative by a centred difference gives a system of ordinary differential equations for the nodal temperatures. The time method still has to resolve the fastest rate that this spatial operator can generate.

An explicit Euler step yields $u_i^{n+1} = r u_{i-1}^n + (1 - 2r) u_i^n + r u_{i+1}^n$ with $r = D\Delta t / h^2$. If $0 \le r \le 1/2$, the update is a convex combination of three temperatures. The maximum cannot rise and the minimum cannot fall in a source-free interior step. Starting with one hot node at 1 and neighbours at zero, $r = 0.4$ gives a centre value 0.2 and neighbours 0.4 after a single step. At $r = 0.6$ the centre would become negative although the continuum heat equation preserves nonnegativity.

Identify the spatial modes that set the limit

Insert $u_j^n = g^n e^{\mathrm{i}j\theta}$, where i is the imaginary unit and j indexes the grid. The two neighbouring values contribute factors $e^{\pm\mathrm{i}\theta}$, yielding $g = 1 - 4r\sin^2(\theta/2)$. Long waves have $\sin^2(\theta/2) \simeq \theta^2/4$, so their decay approaches the continuum rate as the mesh is refined. The alternating grid mode $\theta = \pi$ has the most negative amplification and enforces $r \le 1/2$ for bounded magnitude.

At the limiting value $r = 1/2$ the alternating mode flips sign without losing amplitude. A strict margin damps it. For $D = 10^{-4}$ m²/s and $h = 2$ mm, the maximum explicit step is

$$h^2/(2D) = 0.020\,\mathrm{s}.$$

Refining to 1 mm reduces that limit to 0.005 s. The mesh gives a sharper temperature profile yet can cost four times more steps over the same physical interval, besides the increased number of spatial degrees of freedom.

Compare backward Euler and Crank–Nicolson

Writing the semidiscrete system as $u' = -A_h u$, backward Euler solves $(I + \Delta t A_h)u^{n+1} = u^n$. Crank–Nicolson averages the old and new diffusive fluxes, giving $(I + \Delta t A_h/2)u^{n+1} = (I - \Delta t A_h/2)u^n$. For a Fourier mode put $s = 4r\sin^2(\theta/2)$. Their factors are $1/(1+s)$ and $(1 - s/2)/(1 + s/2)$, respectively. Both remain bounded at arbitrarily large s, but Crank–Nicolson may alternate for under-resolved sharp features.

With $s = 20$, backward Euler retains only $1/21$ per step. Crank–Nicolson yields $-9/11$, retaining much more of the fast component and reversing its sign. For a smooth low-frequency mode Crank–Nicolson generally gives more accurate temporal decay because it is second order. Decide whether early sharp transients or later smooth behaviour is the relevant physical quantity. To identify spatial and temporal error separately, refine Δt first on a fixed mesh, then refine space with time error kept well below spatial error. Merely halving both at once cannot reveal which term dominated the original discrepancy.

Apply a prescribed heat input without confusing it with temperature

An insulated periodic rod with uniform heating rate s measured in K/s has an average temperature that must rise by $s\Delta t$ each step. Sum the centred second differences over the periodic indices: each nodal value appears with coefficients $1, -2, 1$, which cancel. The source alone changes the mean. For $s = 3$ K/s and a step of 0.02 s the numerical average must increase by 0.06 K, independently of the diffusivity and initial profile.

A code can pass this global test while still mishandling short-wave damping, since an alternating error has zero spatial mean. Pair a mean-temperature test with a nonconstant mode-decay test. Use a homogeneous periodic run for isolates of diffusion, and a uniform temperature with prescribed source for the storage term. A failure in one test then identifies a different part of the mathematical balance than a failure in the other.

Prove decay of a discrete thermal energy

Let a closed semidiscrete diffusion model be $M\dot{u} + Ku = 0$, with M symmetric positive definite and K symmetric positive semidefinite. Assume $K\mathbf{1} = 0$, expressing zero flux for a uniform temperature. Backward Euler satisfies $M(u^{n+1} - u^n) + \Delta t K u^{n+1} = 0$. Multiplying by the constant vector and using symmetry gives

$$\mathbf{1}^T M u^{n+1} = \mathbf{1}^T M u^n.$$

Thus the capacity-weighted sum, proportional to stored sensible heat for constant capacities and a fixed reference, is conserved for any step size.

Now multiply by u^{n+1}. The identity $a^T M(a-b) = [\|a\|_M^2 - \|b\|_M^2 + \|a-b\|_M^2]/2$ yields the exact balance

$$\tfrac{1}{2}\|u^n\|_M^2 - \tfrac{1}{2}\|u^{n+1}\|_M^2 = \Delta t\,(u^{n+1})^T K u^{n+1} + \tfrac{1}{2}\|u^{n+1} - u^n\|_M^2.$$

Both terms are nonnegative. The quadratic quantity therefore decays although total sensible heat is conserved. Subtracting the conserved mean first yields the same decay statement for temperature variation around equilibrium. The first term represents spatial smoothing measured at the new time, and the second reflects the numerical time update. This derivation gives a stability certificate on nonuniform grids without assuming Fourier modes.

Audit two cells with different heat capacities

Choose capacities 1 and 3 J/K, conductance 2 W/K, and initial temperature deviations $(4, 0)^T$ K. The conserved sensible heat is 4 J, so the final equilibrium in a closed system must be one kelvin above reference. At $\Delta t = 0.5$ s the matrix $M + \Delta t K$ has rows $(2, -1)$ and $(-1, 4)$, and the right-hand side is $(4, 0)^T$. The new temperatures are $(16/7, 4/7)^T$. Checking their weighted sum gives $16/7 + 3(4/7) = 4$, as required. The unweighted average changes, which is physically appropriate for unequal capacities.

The quadratic quantity at the start is $4^2/2 = 8$. After the step it is

$$[(16/7)^2 + 3(4/7)^2]/2 = 152/49.$$

Its drop is $240/49$. The conductive contribution is $(0.5)(2)(12/7)^2 = 144/49$; the time-update contribution is

$$[(-12/7)^2 + 3(4/7)^2]/2 = 96/49.$$

Their sum recovers the measured drop exactly. Neither contribution equals a loss of the 4 J of sensible heat. The quadratic measure depends on squared temperatures and has different physical units.

A positive heat supply or a prescribed exterior temperature introduces a boundary or source work term into these identities. Applying the closed-system decay test to an externally heated experiment would therefore diagnose the wrong physical model. Verify which conserved or dissipated quantity is actually implied by the capacity matrix, boundary conditions and source terms before comparing numerical integrators.

The conserved mean suggests another check after many large steps. The steady state must have zero conductance dissipation, which for two connected cells requires equal temperatures. Together with the weighted heat total this fixes both temperatures at one kelvin above reference. A scheme that instead approaches the unweighted initial average of two would be converging to the wrong conserved quantity despite producing a visually smooth, uniform endpoint. The measured quadratic variation about the correct mean tends to zero; the uncentred squared quantity tends to a positive constant because the remaining uniform state still has temperature above the chosen reference.

What happens to a grid-scale temperature ripple?

The alternating spatial mode on a uniform heat grid has discrete Laplacian eigenvalue $-4/\Delta x^2$. Substitute this mode into each method and compare the amplitude after one step. The plot isolates the shortest resolvable pattern.

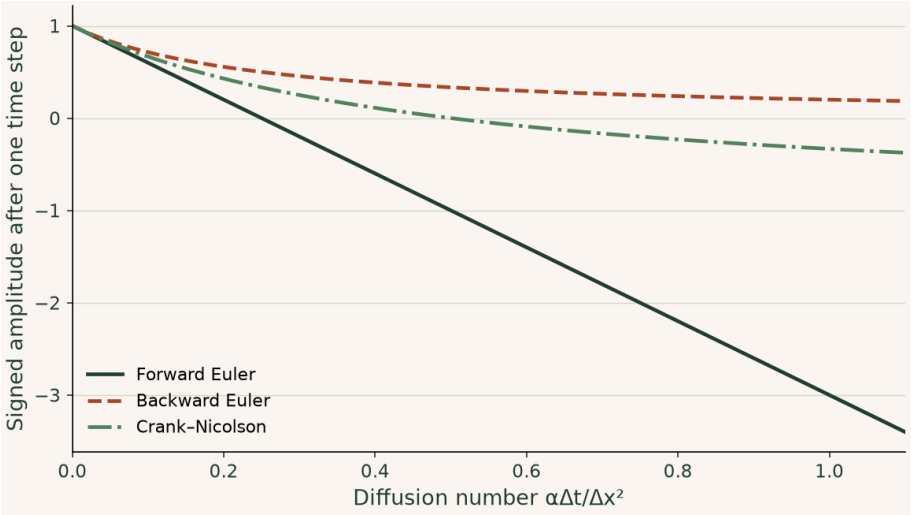

Forward Euler reverses the mode above 1/4 and becomes unstable above 1/2. Crank–Nicolson stays bounded for every plotted step but also reverses the mode above 1/2. Smooth initial data can hide this high-frequency diagnostic.

▶ **Read values from the plot**

Change parameters · What happens to a grid-scale temperature ripple?

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

Highest diffusion number **1.5**

Mode fraction of Nyquist **1**

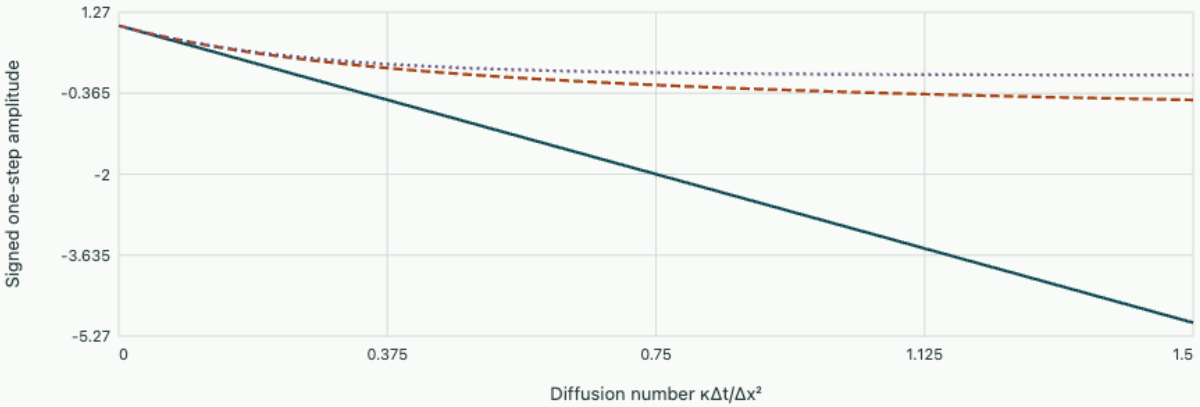

Forward Euler Crank–Nicolson Exact semidiscrete mode

Directional Laplacian magnitude times Δx² is 4. Forward Euler is stable up to r = 0.5 and nonnegative only up to r = 0.25.

What a periodic grid does to a decaying Fourier mode

Numerical partial differential equations

A periodic diffusion solver predicts the shape of a sampled wave while its amplitude decays too slowly. Fourier analysis separates spatial differentiation error, an unresolved physical mode and errors caused by the time integrator.

Diagonalise the differential operator on the periodic interval

Use a dimensionless coordinate $0 \le x < 2\pi$ and solve $u_t = \kappa u_{xx}$ with periodic values and periodic derivatives. The complex exponentials e^{imx}, with integer m, satisfy both conditions and are eigenfunctions of the second derivative. Put $u(x,t) = A_m(t)e^{imx}$ and cancel the common exponential. The scalar amplitude equation is

$$A_m' = -\kappa m^2 A_m.$$

Thus every cosine or sine mode decays by $\exp(-\kappa m^2 t)$, with no phase shift, while the m=0 mean is conserved.

On a uniform grid $x_j = 2\pi j/N$, the samples of m and $m + \ell N$ agree for every integer ℓ. A discrete Fourier transform returns coefficients in one representative band. For the real cosine test with even N we can fold a positive physical mode into $k = \min(r, N-r)$, $r = m \bmod N$. A Fourier spectral Laplacian multiplies that coefficient by $-k^2$. For a resolved mode it differentiates the trigonometric interpolant exactly, up to arithmetic. For unresolved data it differentiates the aliased representative; no operation using only those samples can identify the original m.

Apply the centred difference stencil to a single wave

The standard second difference is $(u_{j+1} - 2u_j + u_{j-1})/h^2$, $h = 2\pi/N$. Put the grid wave $u_j = e^{ikjh}$ into the stencil. Factor out the common phase and simplify with $e^{i\xi} + e^{-i\xi} = 2\cos\xi$. This gives the symbol and its small-h expansion:

$$\frac{e^{ikh} - 2 + e^{-ikh}}{h^2} = -\frac{4}{h^2}\sin^2(kh/2) = -k^2 + \frac{k^4 h^2}{12} + O(k^6 h^4).$$

The effective squared wavenumber is smaller than k^2 for every nonconstant mode in the resolved band, so this spatial approximation decays too slowly if the time equation is integrated exactly. The error is modest when $kh \ll 1$, but at the Nyquist mode $kh = \pi$ the ratio is

$$4/\pi^2 = 0.4053.$$

A stable time integrator cannot repair this spatial attenuation error without additional modeling. The experiment exponentiates both discretizations' symbols exactly, isolating space and sampling here.

Measure separate errors with the same sampled initial data

Choose $\kappa = 0.5,\ m = 3,\ N = 16$. The true squared wavenumber is 9 and the difference stencil gives 8.006. Starting at unit amplitude, at $t = 0.20$ exact and spectral amplitudes are 0.40657; finite differences give 0.44906. Doubling N to 32 changes the difference result to 0.41716. The absolute amplitude error falls from 0.04249 to 0.01059, a ratio of 4.01 consistent with the second-order expansion. This is a spatial-convergence study because time evolution introduces no truncation error.

Now take $m = 12,\ N = 16$. Every sampled cosine equals the low mode $k = 4$ at $t = 0$. At $t = 0.02$, the continuum amplitude for the true physical mode is $e^{-1.44} = 0.237$, whereas the spectral grid predicts

$$e^{-0.16} = 0.852.$$

A spectral computation can therefore agree exactly with its trigonometric interpolant and disagree strongly with the evolution of the intended signal. Sampling on a finer grid or obtaining prior bandwidth information resolves that ambiguity. The mean mode is particularly revealing when a nonconstant physical cosine aliases to a constant.

Start at mode 3 and successively double N from 8 to 32. Compute two amplitude-error ratios at the same time. Then use mode 12 on 16 points to observe an alias, and increase N until the Fourier grid represents the intended mode.

Open this experiment online

Try the calculation

Periodic grid points	16
Physical cosine mode	3
Diffusivity	0.5

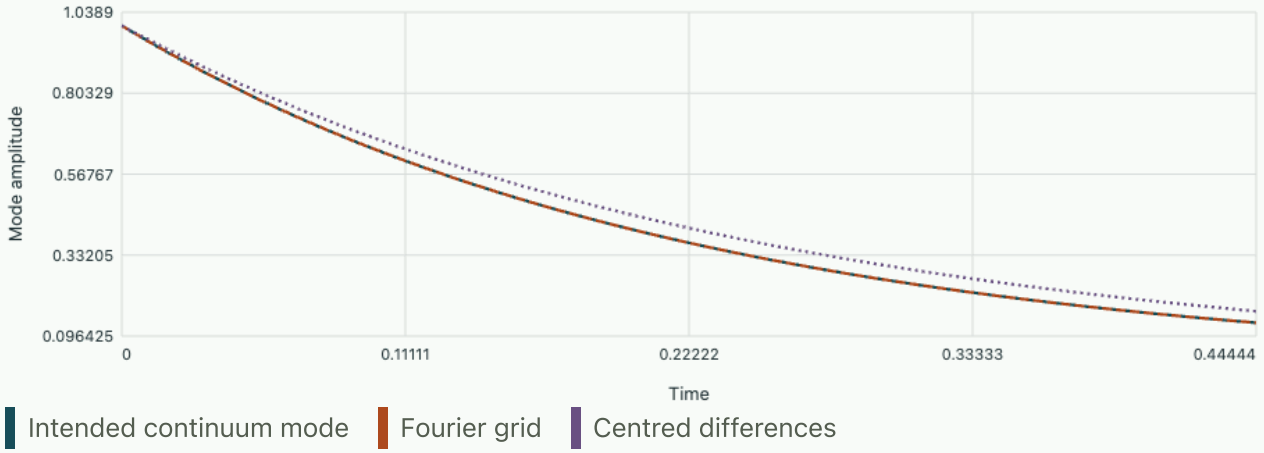

The grid represents mode 3 as 3. At t = 0.20, amplitudes are 0.40657 for the physical mode, 0.40657 for the Fourier grid and 0.44906 for the difference stencil.

Tutorial · resolve a damping rate to a target tolerance

For a mode well inside the grid bandwidth, estimate the number of samples per wavelength needed for a relative error below 1% in the decay rate with centred differences. Test the estimate for $m = 2$. How could the time integrator contaminate a measurement of this spatial error?

▼ **Show the resolution estimate**

Writing p for samples per wavelength gives

$$kh = 2\pi/p.$$

The relative rate error is approximately $(kh)^2/12$. Requiring at most 0.01 gives $p \gtrsim 18.14$; try 19 samples per wavelength, or $N = 38$ for $m = 2$. The exact ratio $[\sin(\pi/19)/(\pi/19)]^2$ is about 0.991. A forward-Euler step would replace exponential propagation by $1 - \kappa k_{\mathrm{eff}}^2 \Delta t$. Its repeated-step error could partly mask the underdamping of the spatial stencil. Compare with the exponential of the discrete symbol before attributing a measured rate to space alone.

Further reading · Lloyd N. Trefethen, Spectral Methods in MATLAB · periodic grids, Fourier differentiation and downloadable programs

Check the discrete eigenvalue with OneSimulation

A spatial sine also diagonalises the five-point Dirichlet Laplacian on a square. Use the same modified wave-number argument in both coordinate directions. Run `python poisson-mode.py --mode 2 --cells 8 16 32` and compare the observed nodal error with the discrete eigenvalue prediction in the CSV. This is a steady forced problem whose spatial operator can be checked independently of any time integrator. Download the OneSimulation verification model.

Splitting a stiff relaxation from a slowly varying drive

Numerical ordinary differential equations

A sensor tries to follow a changing reference while a fast relaxation damps departures. An explicit method may oscillate wildly even though the underlying transient decays monotonically. Treating the restoring term implicitly can resolve the slow motion using steps longer than the transient time. Whether those steps measure the transient accurately is a different question.

Construct an exact benchmark with two time scales

Let $s(t) = \sin t$ and choose a relaxation time $\varepsilon > 0$. Consider $y' = s'(t) + (s(t) - y)/\varepsilon$, with $y(0) = 1$. Set $e = y - s$. Subtraction of s' gives $e' = -e/\varepsilon$, hence the reference solution $y(t) = \sin t + e^{-t/\varepsilon}$. The slow reference changes over times of order one, while the initial mismatch decays significantly over ε. We know both pieces exactly and can diagnose errors even when a stable result looks visually smooth.

Explicit Euler applied to a perturbation of the state multiplies that perturbation by $1 - r$, with $r = h/\varepsilon$. It damps monotonically only for $0 < r \le 1$, changes sign while damping for $1 < r < 2$, and grows for $r > 2$. These statements compare nearby solutions at equal drive. The full signal also contains the time variation of s. At the boundary $r = 2$ a state perturbation alternates without decaying. Refining the step to follow a slow input alone can therefore give the wrong choice of h.

Solve only the fast part implicitly

Split $f_{\text{slow}}(t) = \cos t$ from $f_{\text{fast}}(t, y) = (\sin t - y)/\varepsilon$. Evaluate the slow part at t_n and the fast part at t_{n+1}, y_{n+1}. This first-order additive implicit-explicit Euler step needs one scalar linear solve:

$$(1 + r)y_{n+1} = y_n + h\cos t_n + r\sin t_{n+1}.$$

The state perturbation now multiplies by $1/(1 + r)$. It is damped at every positive step and tends to zero as r grows. Subtract $\sin t_{n+1}$ from both sides after solving. The mismatch obeys $e_{n+1} = [e_n + h\cos t_n - (\sin t_{n+1} - \sin t_n)]/(1 + r)$. The bracketed forcing defect is $O(h^2)$ for smooth input. In the stiff limit at fixed h, the endpoint approaches the current reference. Moving the reference in the fast term to the old time would instead drag it one sample behind in that limit. The choice of what to treat implicitly includes the time at which a fast equilibrium is imposed.

Test early-time accuracy and late-time stability

For $\varepsilon = 0.02$, $h = 0.05$, $r = 2.5$. At the first step the exact value is

$$\sin(0.05) + e^{-2.5} = 0.1321.$$

Explicit Euler gives

$$1 + 0.05(1 - 50) = -1.45.$$

The implicit-explicit value is

$$[1 + 0.05 + 2.5\sin(0.05)]/3.5 = 0.3357.$$

The latter is stable but visibly overestimates the surviving transient. Its damping factor $1/3.5 = 0.286$ differs substantially from $e^{-2.5} = 0.0821$. A question about peak rate or initial settling requires shorter steps even if the late-time reference is recovered.

The experiment holds the final time at one while the number of steps changes. The vertical coordinate is the base-ten logarithm of absolute error, with values below 10^{-15} drawn at that floor. Negative heights therefore correspond to errors less than one. Track the maximum error separately from the final error as the transient dies out. Try an initially unstable explicit run, then double the number of steps. The stiffness criterion depends on the ratio h/ε, while accuracy on the reference also depends on h itself.

Open this experiment online

Try the calculation

Relaxation time ε **0.02**

Steps to final time 1 **20**

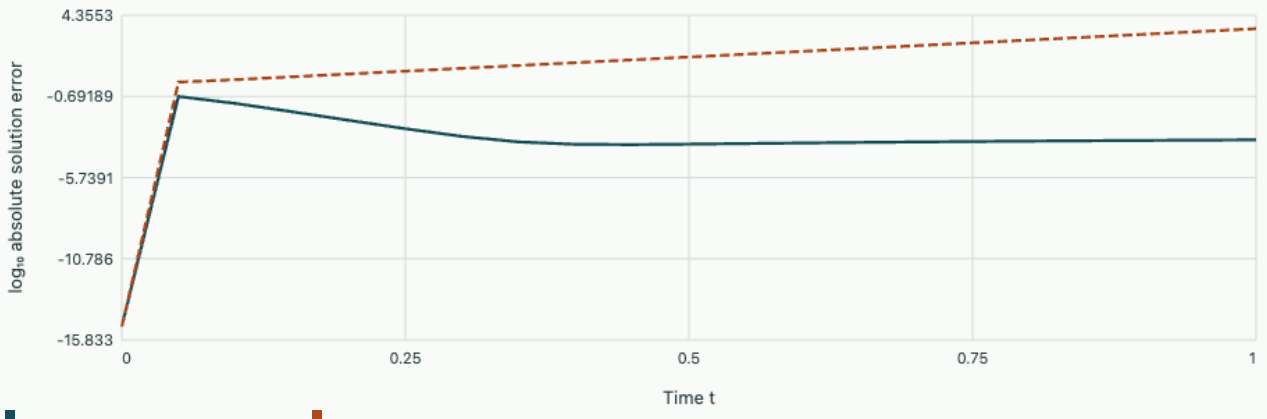

h/ε = 2.5. Explicit perturbation factor -1.5, implicit factor 0.2857. IMEX maximum error 0.2036, final error 0.0004054.

Tutorial on identifying which component is stiff

Keep $h = 0.04$ and reduce ε from 0.04 to 0.004. For each case compare exact, explicit and implicit amplification of a state perturbation. What quantity would force a reduction of h even with the implicit-explicit method?

▼ **Work through the solution**

At $r = 1$ the three factors are $e^{-1} = 0.368$, zero and $1/2$. At $r = 10$ they are 4.54×10^{-5}, -9 and $1/11 = 0.0909$. For a measurement of the initial decay rate the implicit discretisation substantially

misrepresents both exact factors. Bounding error during that transient requires smaller steps relative to ε, even though its numerical amplification is stable.

Further reading · Ascher, Ruuth and Wetton, implicit-explicit time discretisations.

Learning a drifting temperature from a sequence of readings

Inverse problems and least squares

Derive the scalar Kalman update by completing the square and see how process uncertainty controls memory of the past.

Propagate a distribution through time

A drifting offset follows $x_k = x_{k-1} + w_k$. A thermometer reports $z_k = x_k + v_k$. Assume independent, zero-mean Gaussian process increments of variance Q and sensor noise of variance R, independent at each time and of the initial state. Before the next reading, the mean stays at $m^- = m_{k-1}$ and the variance adds, $P^- = P_{k-1} + Q$.

The negative log posterior, up to constants, is $(x - m^-)^2/(2P^-) + (z_k - x)^2/(2R)$. Collect the coefficient of x^2 and its linear term. The result can be written in two equivalent forms:

$$P_k = (1/P^- + 1/R)^{-1},$$

$$m_k = m^- + K(z_k - m^-),$$

$$K = P^-/(P^- + R).$$

Thus the fraction of the innovation used is a model prediction. When sensors are noisy the gain falls. When the model permits rapid drift, uncertainty grows between readings, giving the new measurement a larger role. These are variances, so a stated standard deviation must be squared first.

A reading with units and a repeated reading

A prior says 20.0 ± 2.0 °C (one standard deviation), with negligible intervening drift. A sensor of standard deviation 1 °C reads 23 °C. The gain is $4/(4+1) = 0.8$, giving 22.4 °C and variance 0.8 °C². A second independent reading of 23 °C gives gain $0.8/1.8$ and mean 22.6667 °C. Treating a duplicate record as independent would falsely claim the same information gain.

In a long random walk with $Q = 0.25$ and $R = 1$ in squared units, a steady posterior variance satisfies

$$P = (P + Q)R/(P + Q + R).$$

Rearrangement gives $P^2 + QP - QR = 0$, so the physically positive root is about 0.3904. New uncertainty replenishes the information squeezed out by measurement.

Follow a change in the underlying temperature

For a reproducible illustration, a fixed synthetic thermometer trace contains a smooth periodic disturbance and a step in the hidden temperature. It is deliberately held fixed while the filter's assumed Q and R change. Compare the lag just after the step with the wiggle once the temperature settles.

Open this experiment online

Try the calculation

Assumed process variance Q (°C²) 0.2

Assumed sensor variance R (°C²) 1

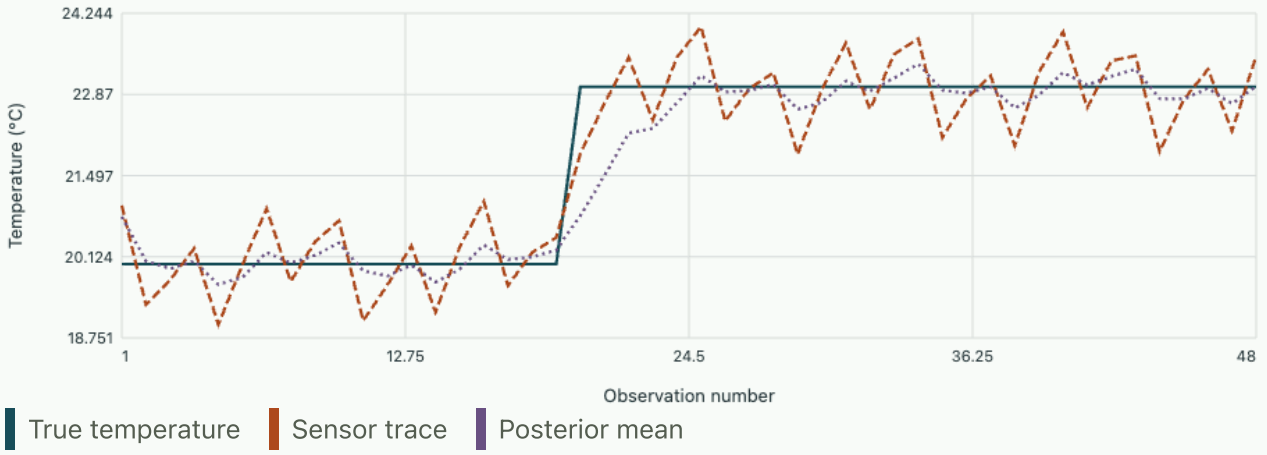

Last posterior mean 23.005 °C, standard deviation 0.59855 °C. Gain 0.35826.

Tutorial on noisy data over time

If $Q = 0$, show how n independent readings with the same variance affect precision.

▼ **Worked solution**

Each update adds $1/R$ to precision $1/P$. Induction gives

$$1/P_n = 1/P_0 + n/R.$$

This is precisely the precision of a batch Gaussian update.

What pattern in the innovations would call the random-walk assumptions into question?

▼ **Worked solution**

A long run of positive innovations after a systematic ramp suggests bias in the prediction. Serial correlation suggests missing dynamics or correlated observation errors. Check against predicted innovation variance $P^- + R$ and examine signs and lags as well as their sizes.

Read the university treatment

Extend the scalar precision update to several state variables and identify the matrix made invertible by a positive observation variance. Stephen Boyd · Stanford EE363, estimation and the Kalman filter

Predicting a new measurement when the fitted parameters are correlated

Inverse problems and least squares

Propagate the full posterior covariance to a quantity of interest, distinguish latent and measured uncertainty and select the observation that reduces the uncertainty of the actual prediction.

One model supplies many possible targets

For parameters θ with posterior mean m and covariance C, a linear target $s = h^T\theta$ has mean $h^T m$. Subtract the mean and expand the quadratic product to obtain

$$\operatorname{Var}(s) = h^T C h.$$

An independent sensor with variance R changes the variance of a future measured value to $h^T C h + R$. Repeated fits of noisy datasets and repeated physical measurements are distinct sources of variation.

A candidate new observation $z = a^T\theta + \epsilon$ with variance R_a has covariance $h^T C a$ with the target. For a Gaussian posterior the conditional covariance update is $C_+ = C - Caa^T C/(a^T C a + R_a)$, found either by completing the square or projecting out the optimal multiple of the innovation. The target variance decreases by

$$\Delta \operatorname{Var}(s) = \frac{(h^T C a)^2}{a^T C a + R_a}.$$

Background · Posterior predictive

$$p(y_*|y) = \int p(y_*|m)p(m|y)\,dm$$

A distribution over parameters induces uncertainty in a new observation. Measurement noise and uncertainty of the prediction mean are distinct contributions.

A precise sum and an uncertain difference

Let C have diagonal entries 1 and correlation -0.9. A target $h = (1, 1)$ has variance $2 - 1.8 = 0.2$. The difference $h = (1, -1)$ has variance 3.8. Ignoring correlations would predict 2 for both. With sensor variance 0.25, a newly measured sum has variance 0.45 although the latent sum variance remains 0.2.

For the target sum, measuring the first coordinate with variance 0.25 reduces variance by $(1 - 0.9)^2/(1 + 0.25) = 0.008$. Measuring the sum itself at the same variance reduces it by $0.2^2/(0.2 + 0.25) = 0.0889$. These two sensors have differently normalised sensitivity vectors. Their stated noise levels refer to their own output units and must be calibrated before a cost comparison.

Choose the measurement direction

For a fairer geometric comparison, restrict both measurement and prediction directions to unit length. Fix the target along the normalised sum. Sweep candidate observation angles and read the remaining prediction variance. Change the parameter correlation: when does a coordinate sensor improve the target appreciably?

Open this experiment online

Try the calculation

Parameter correlation **-0.9**

Candidate sensor variance **0.25**

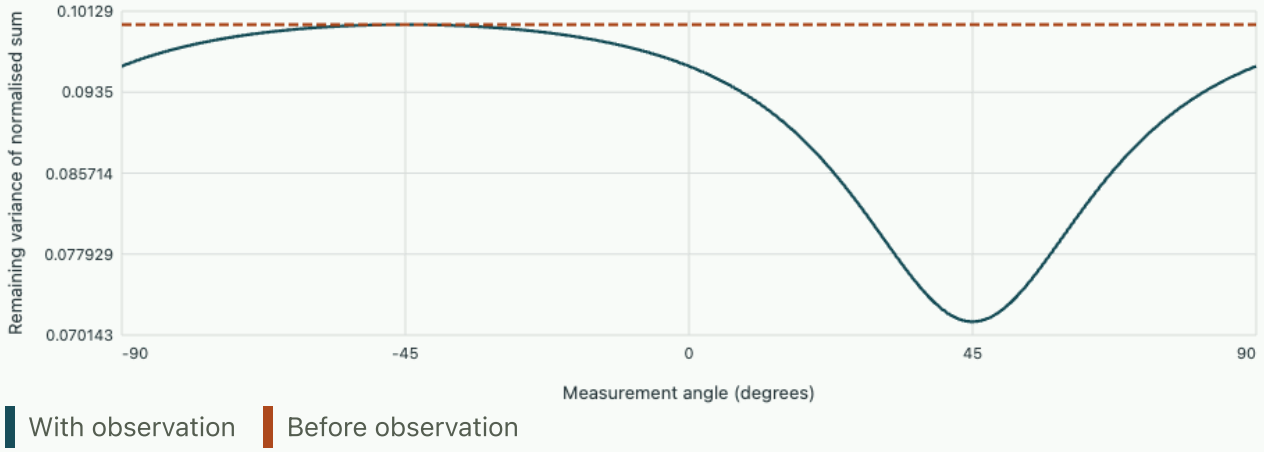

Target variance 0.1 initially, 0.096 with a first-coordinate sensor, 0.071429 with a sensor along the target.

Tutorial on prediction and parameter recovery

If $h^T C a = 0$, can a noise-free measurement of $a^T \theta$ improve the linear prediction under these Gaussian assumptions?

▼ Worked solution

The covariance reduction is exactly zero. The measurement may determine another direction precisely while carrying no linear information about the target.

The posterior on θ is strongly non-Gaussian. Which conclusions still hold exactly?

▼ Worked solution

Mean and variance of a fixed linear prediction still follow from its first two moments, and independent sensor variance still adds. The Gaussian conditional update and an interval based only on two standard deviations need additional assumptions.

Read the university treatment

Compare the projection of a vector measurement with the covariance of a scalar target. Check the direction of each matrix product. Stephen Boyd · Stanford EE363, estimation

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

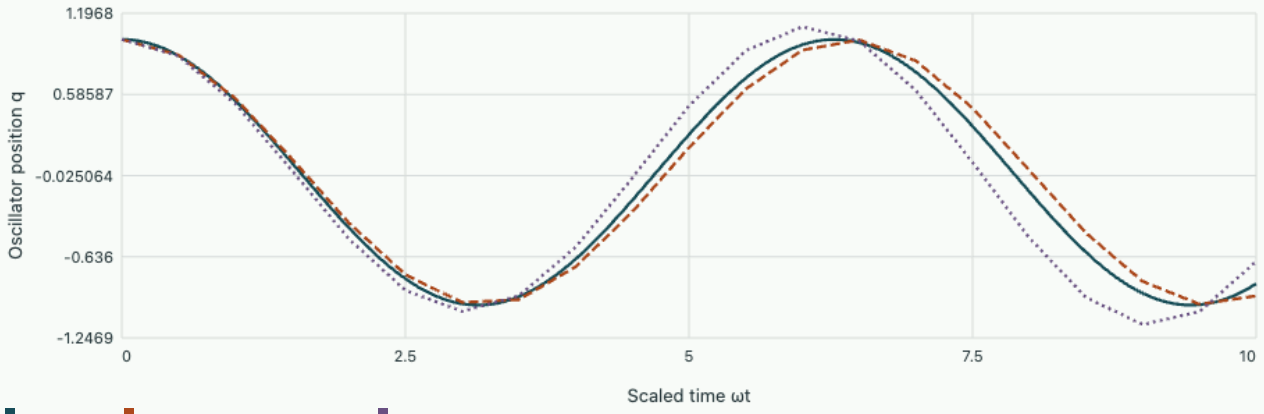

Implicit midpoint phase lag 0.2009 rad. Forward-backward distance from the start is 8.951e-16 for implicit midpoint and 0.3635 for explicit midpoint.

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

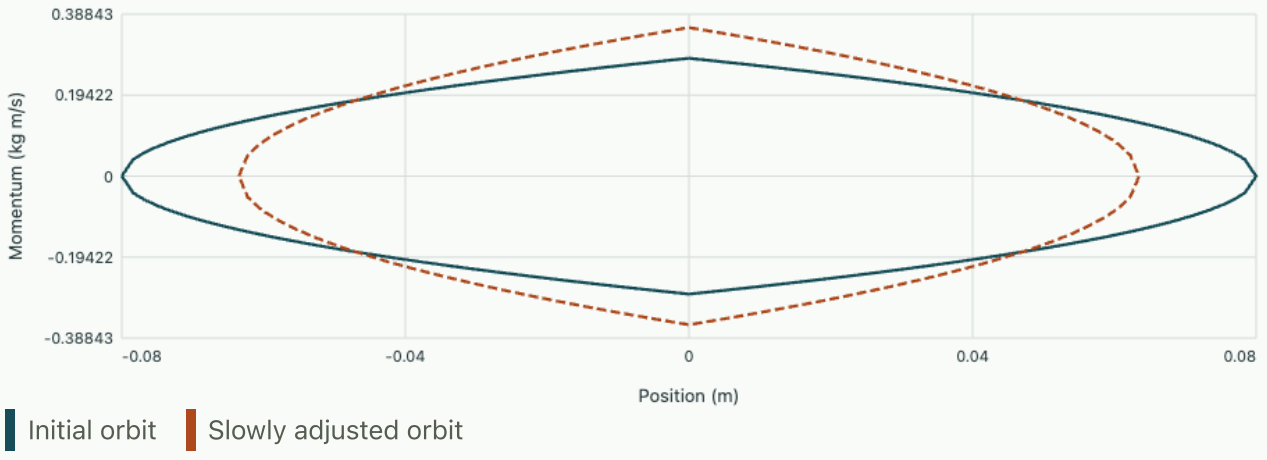

Final amplitude 0.0635 m, energy 0.254 J and period 0.3564 s.

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.

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**

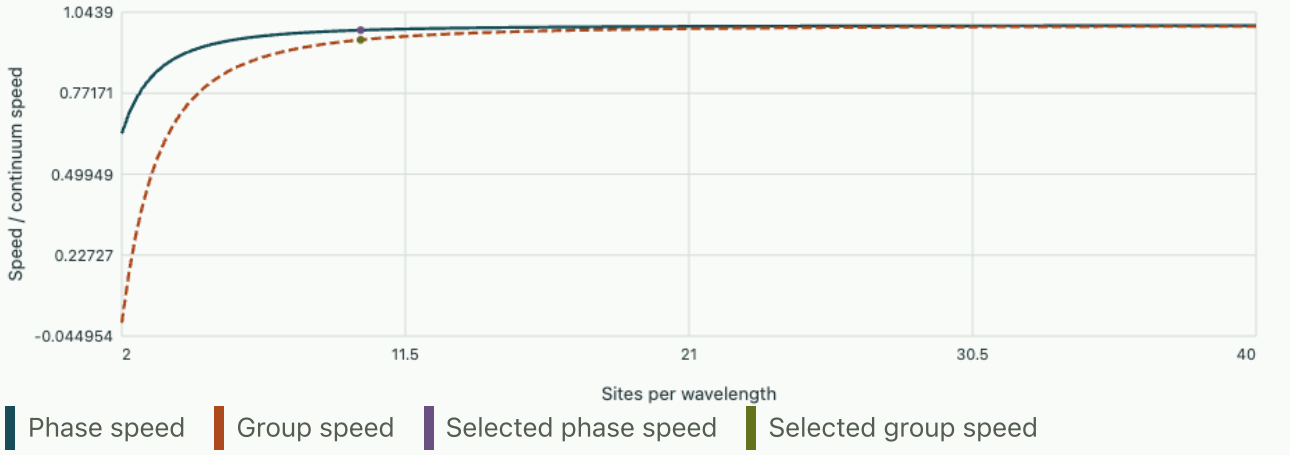

Frequency and phase speed are 98.36% of the continuum prediction. Group speed is 95.11%.

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.

Applied study

Connect heat exchange and thermal equilibrium to energy conservation, entropy and microstates.

Which temperature will two bodies reach?

Touch two metal blocks together inside an insulated container. Their final temperature is constrained by conservation of energy. That law alone also permits the hot block to heat up at the expense of the cold one. A second idea, entropy, selects the direction we actually observe. A third, thermal conductance, sets the time it takes.

Assume that each block is internally close to equilibrium and has a constant total heat capacity C_i in J/K. The blocks exchange heat but do no external work. A transfer dQ from block 1 to block 2 produces $C_1 dT_1 = -dQ$ and $C_2 dT_2 = dQ$. Adding and integrating gives

$$C_1 T_1 + C_2 T_2 = \text{constant},$$

$$T_* = \frac{C_1 T_{1,0} + C_2 T_{2,0}}{C_1 + C_2}.$$

The second expression assumes the bodies have reached a common temperature. A larger heat capacity puts more weight on that block's starting temperature.

Count the ways to share energy

At fixed total energy U, suppose block 1 has energy U_1. Weakly coupled blocks have $\Omega_1(U_1)\Omega_2(U - U_1)$ accessible combinations of microscopic states. For an isolated equilibrium system with equally likely microstates, this multiplicity determines the probability of an energy partition. Taking $S = k_B \ln \Omega$ turns the product into a sum. At its maximum

$$\frac{dS_{\text{total}}}{dU_1} = \frac{1}{T_1} - \frac{1}{T_2} = 0,$$

$$\frac{1}{T_i} = \frac{\partial S_i}{\partial U_i}.$$

Positive heat capacities make the curvature negative, since $\partial^2 S_i / \partial U_i^2 = -1/(C_i T_i^2)$. Equal temperature thus maximises total entropy. Before equilibrium a small transfer from hot to cold increases entropy by $dQ(1/T_{\text{cold}} - 1/T_{\text{hot}}) > 0$. Temperatures here must be absolute.

From an isolated pair to a canonical ensemble

When a small system contacts a much larger reservoir, its temperature is effectively set by the reservoir. The chance of observing a particular microstate of energy E_j is proportional to the number of reservoir states at energy $U - E_j$. Expand the reservoir entropy about U:

$$S_R(U - E_j) \simeq S_R(U) - E_j / T.$$

Since multiplicity is exponential in entropy, the probability becomes

$$p_j = e^{-\beta E_j}/Z,$$

with $\beta = 1/(k_B T)$ and $Z = \sum_j e^{-\beta E_j}$.

For one unit with levels 0 and Δ,

$$Z = 1 + e^{-\beta \Delta}.$$

If $\Delta = k_B T \ln 3$, the excited state has probability $1/4$ and mean energy is $\Delta/4$. Raising temperature at fixed gap drives the two probabilities toward one half. This microscopic account relates heat capacity to the response of mean energy to temperature.

Same change of state, different heat and work

For a closed system the first law reads $dU = \delta Q - \delta W$, with work positive when the system does it. Internal energy and entropy are state functions. Heat and work depend on the path. To see this, compare two ways for one mole of ideal gas initially at 300 K to expand to twice its volume.

During reversible isothermal expansion, $p = RT/V$ and $W = \int p\, dV = RT \ln 2 = 1.73$ kJ. The ideal gas's internal energy depends only on temperature, so $Q = W$. Free expansion into an evacuated chamber inside an insulated shell instead gives

$$W = Q = 0.$$

The final temperature is again 300 K. Both paths give the same gas entropy increase $R \ln 2 = 5.76$ J/K. In the reversible case the reservoir loses this much entropy, while in free expansion the insulated gas produces it internally. In the two-block problem the entropy production likewise measures irreversibility inside the chosen boundary.

Give the contact a conductance

Model a short thermal bridge by a constant conductance G in W/K with negligible heat storage. Its heat flow is $\dot{Q} = G(T_1 - T_2)$. Subtract the two heat balances to obtain

$$d(T_1 - T_2)/dt = -G(1/C_1 + 1/C_2)(T_1 - T_2).$$

The difference decays exponentially with $\tau = C_1 C_2/[G(C_1 + C_2)]$. Meanwhile the entropy production rate is $G(T_1 - T_2)^2/(T_1 T_2)$, which falls to zero at equilibrium.

The equilibrium temperature is independent of the bridge conductance, provided heat transfer can continue. The same energy balance for a body in contact with a reservoir gives a different time constant because the reservoir temperature stays fixed. Cooling a body in a room.

A hot block and a larger cold block

Take $C_1 = 500$ J/K at 360 K, $C_2 = 1500$ J/K at 300 K and $G = 2.0$ W/K. The common equilibrium temperature is 315 K. Block 1 eventually releases $500(360 - 315) = 22.5$ kJ and block 2 absorbs $1500(315 - 300) = 22.5$ kJ. The initial heat current is 120 W and $\tau = 187.5$ s.

To find the entropy change between the initial and final states, integrate $dS = C\,dT/T$ for each block along a reversible reference path. The sum is

$$500 \ln(315/360) + 1500 \ln(315/300) = 6.42\,\mathrm{J/K}.$$

Energy transfer cancels between the two bodies. Their entropy changes have different magnitudes because they receive or lose heat at different temperatures.

Exercises on microscopic and macroscopic equilibrium

1. Count a small system by hand

Two groups each contain four distinguishable two-level units. Distribute four identical excitation quanta over the eight units, at most one quantum per unit, with all microstates equally likely. Which partition is most probable? Does it occur every time?

▼ **Count the microscopic arrangements**

If group 1 has j quanta, the multiplicity is $\binom{4}{j}\binom{4}{4-j}$. For $j = 0, 1, 2, 3, 4$ this gives $1, 16, 36, 16, 1$, summing to 70. Equal sharing is the most probable partition, with probability $36/70$. A small system shows substantial fluctuations around this partition. In macroscopic systems the relative spread of the energy share is typically much smaller.

---

2. Diagnose an energy-conserving prediction

From the worked initial state a proposed model raises the hot block by 3 K and lowers the colder block by 1 K. Test its energy and entropy changes for an isolated, passive pair.

▼ **Apply both thermodynamic laws**

The net energy change is

$$500(3) + 1500(-1) = 0.$$

But $500 \ln(363/360) + 1500 \ln(299/300) = -0.86$ J/K to two decimals. The macroscopic change contradicts the second law for a passive isolated pair. An external device moving heat uphill would require a larger system boundary to account for its operation.

ONESIMULATION LABORATORY

Separate the equilibrium temperature from the relaxation time

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 notebook couples two thermal capacities through their temperature inputs. Run its parameter, solve and comparison cells for each setting. The log-slope of $T_1 - T_2$ provides a measured decay rate. Its inverse is the measured relaxation time. Use the capacity-weighted mean at every stored time to estimate T_*, since the two temperatures still differ at the final stored sample. Two-body thermal exchange model.

1. **Baseline.** Start with heat capacities of 500 and 1500 J/K, initial temperatures of 360 and 300 K, and a 2 W/K bridge. Compare the time for the initial temperature difference to fall by a factor e with 187.5 s. Check entropy by calculating $C_1 \ln(T_1/T_{1,0}) + C_2 \ln(T_2/T_{2,0})$ across the stored history.
2. **Improve the bridge.** Change `conductance` from 2 to 4 W/K and rerun. Predict which of the weighted mean and decay rate will double.
3. **Change the hot body.** Restore the conductance and set `heat_capacity_1 = 1500.0` J/K. Examine how the fitted rate and weighted mean both change. For each run compare the fitted rate with $1/\tau$ within 0.1%, and check that the weighted mean drifts by less than 0.001 K.

Stored temperature histories from the coupled model

Thermal contact	Weighted mean (K)	Fitted decay rate (s⁻¹)	Relaxation time (s)
500 J/K hot block, 2 W/K bridge	315.00	0.005333	187.50
500 J/K hot block, 4 W/K bridge	315.00	0.010667	93.75
1500 J/K hot block, 2 W/K bridge	330.00	0.002667	375.00

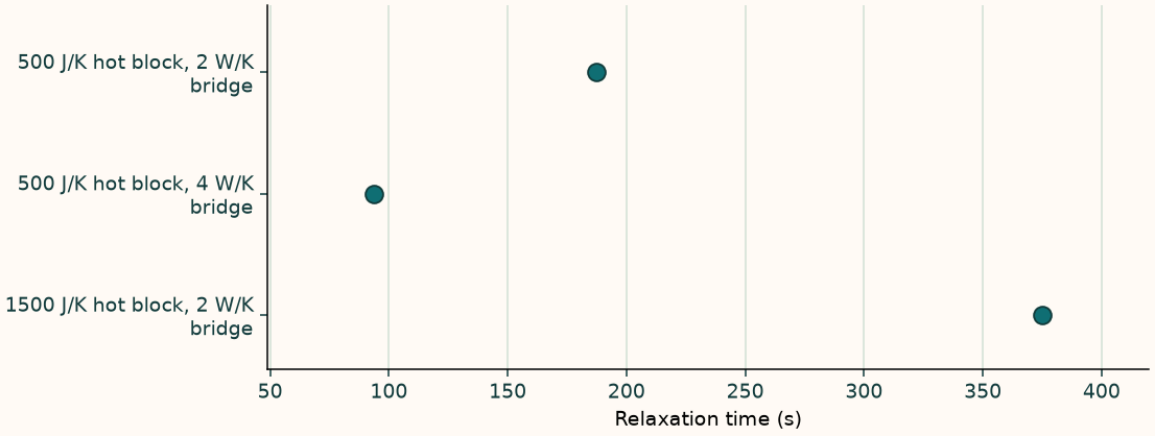

Doubling the bridge conductance cuts the relaxation time from 187.5 to 93.75 s. Increasing the hot block's heat capacity instead doubles the time to 375 s.

Download experiment files (ZIP)

Extract the files together and open `two_mass_thermal_exchange.ipynb` in a local OneSimulation environment.

Try it · Two blocks approaching thermal equilibrium

The initial temperatures are the notebook's 360 K and 300 K, with heat capacities 500 and 1500 J/K and conductance 2 W/K. Double the conductance to halve the relaxation time; restore it and set the first capacity to 1500 J/K to move the equilibrium to 330 K. These curves solve the same constant-property ODE analytically.

Open the corresponding OneSimulation notebook, code and output

Open this live calculation

Try the calculation

Heat capacity C_1 (J/K)	500
Heat capacity C_2 (J/K)	1500
Thermal conductance G (W/K)	2

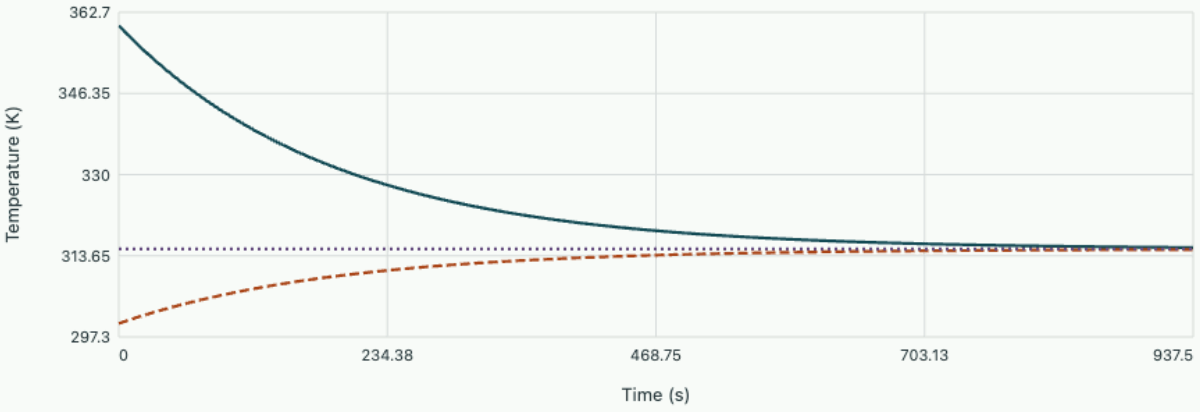

Equilibrium 315 K; relaxation time 187.5 s; decay rate 0.0053333 s^{-1}; total entropy increase 6.4195 J/K. Exact solution of the notebook's isolated two-block ODE, starting at 360 K and 300 K with constant heat capacities and conductance.

Problem sheet 6 · Identify a heat capacity and a bridge conductance

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.

Chemical potential

$$\mu = (\partial F/\partial N)_{T,V}$$

This change in free energy per added particle controls exchange between subsystems. In a finite small system an actual finite difference may matter.

Derivation · Chemical equilibrium and reaction extent

Connected correlation

$$C(r) = \langle s_i s_{i+r} \rangle - \langle s_i \rangle \langle s_{i+r} \rangle$$

Subtract the independent contribution before identifying a decay length. A periodic sample introduces a second propagation distance around the ring.

Derivation · Extracting a length from thermal spin fluctuations

Heat capacity

$$C_V = \mathrm{var}(E)/(k_B T^2)$$

This fluctuation formula uses a canonical equilibrium distribution at fixed volume and a temperature-independent energy operator. A finite reservoir suppresses subsystem fluctuations.

Derivation · How a finite heat bath changes a fluctuation

Helmholtz free energy

$$F = -k_B T \ln Z$$

At fixed particle number and temperature, differentiating with respect to volume yields minus the pressure. The ensemble specifies which variables to hold fixed.

Derivation · Thermodynamic potentials and phase boundaries

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 · Partition functions and thermal fluctuations

Multiplicity and entropy

$$S = k_B \ln \Omega$$

Count microstates compatible with the macroscopic constraints. Distinguishable particles, indistinguishable occupations and a changed energy window need distinct counts.

Derivation · Counting states and defining temperature

Further reading

- Derive the ideal gas in the canonical ensemble, then study the failure of the independent-particle picture near a phase transition. Track which variables are held fixed in each derivative. David Tong · Statistical Physics

Compare specific topics with the university sources