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.

Contents

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 $\int p\,dV = Nk_BT\log(V_f/V_i)$, giving $Nk_B\log(V_f/V_i)$. 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 ah^3 per particle, the gas entropy would shift by $-Nk_B \log a$ 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

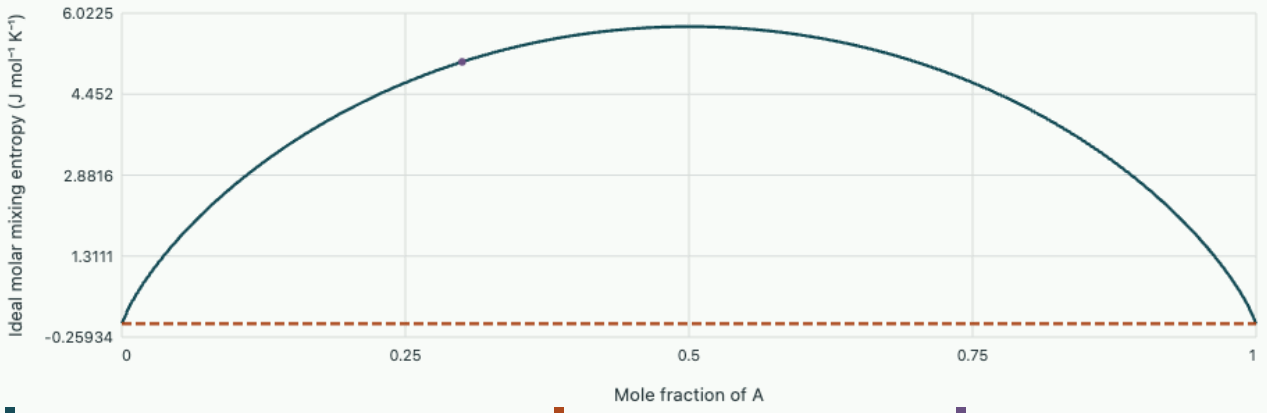
Vary the mixture composition and the reservoir temperature. Compare the entropy with reversible work for separately identifiable ideal components.

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 $\Delta/k_BT = 2$, $p = 0.1192$ 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 $-k_BT \log Z$ 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

$$U[q] - TS[q] - F = k_BT \sum_i q_i \log(q_i/p_i) \geq 0.$$

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

$d\Omega = -S dT - p dV - N d\mu$ for $\Omega = F - \mu N$. A potential's natural variables specify which gradient directly returns an observable. From $F(T, V, N)$ one obtains

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

$$S = -(\partial F/\partial T)_{V,N};$$

their cross derivatives agree if the potential is smooth. At phase boundaries use phase branches and appropriate coexistence conditions instead of differentiating through a kink as if it were smooth.

Let the reservoir exchange particles as well as energy

With both exchanges allowed, expand reservoir entropy in E and N. Its differential includes $-\mu_R dN_R/T$. A system microstate with E, N therefore acquires weight $e^{-\beta(E-\mu N)}$. Summing normalises the grand canonical distribution. At fixed temperature and volume, derivatives with respect to $\beta\mu$ give mean particle number and number variance. In a homogeneous equilibrium region with suitable thermodynamic limit, $\Omega = -pV$.

$$\Xi = \sum_N e^{\beta\mu N} Z_N,$$

$$\langle N \rangle = \partial_{\beta\mu} \log \Xi,$$

$$\mathrm{var}(N) = k_B T \left(\frac{\partial \langle N \rangle}{\partial \mu} \right)_{T,V}.$$

A surface of M independent adsorption sites, each empty or holding one particle with binding energy $-\epsilon$, has $\Xi = [1 + e^{\beta(\mu+\epsilon)}]^M$. Its occupancy fraction is logistic and its variance $Mf(1-f)$. At $M = 400$ and half coverage, the mean occupancy is 200 with standard deviation 10. The slope with respect to chemical potential is largest there. Near saturation, a small number variance and an apparently quiet detector indicate poor sensitivity to further increases of μ, not necessarily high precision in inferring μ.

Compare the adsorption response slope with the independently predicted number variance

The grand formula links number fluctuations to isothermal compressibility in a bulk one-component fluid. Write $n = \langle N \rangle / V$ and use $dp = n\, d\mu$ at constant T. Then $(\partial n/\partial \mu)_T = n^2 \kappa_T$ and $\mathrm{var}\, N / \langle N \rangle = n k_B T \kappa_T$. A dilute classical ideal gas gives unity. These fluctuations refer to a region exchanging particles. If an observer counts all particles of a strictly sealed closed box, the number variance is zero even though its compressibility is not. The boundary condition of the measurement is part of the formula.

Infer a temperature and explain when it is identifiable

The logarithm of a measured population ratio is linear in β for nondegenerate known levels. With enough independent samples one can estimate temperature either from this ratio or by matching

measured energy to its canonical expectation. If all observed states have nearly the same energy, changing β barely changes their probabilities. The available temperature information then becomes small, regardless of detector precision.

For independent canonical draws, differentiating a single log probability gives

$$\partial_\beta \log p_i = -(E_i - \langle E \rangle).$$

Its squared expectation is energy variance, the Fisher information per observation about inverse temperature. At n_s independent draws information adds. Equality between response and noise reappears as a statistical bound on temperature estimation. A long record with correlation time comparable to its length may contain few independent draws despite a large sample count. Recording more frequently without extending the observation window need not improve a temperature estimate.

A pressure reservoir selects enthalpy-weighted volume fluctuations

An isothermal piston in contact with a pressure reservoir supplies a fourth ensemble closely related to the first three. A configuration at volume V costs external work $p_{\mathrm{ext}} V$ in addition to its microscopic system energy. At fixed T, N, p_{ext}, integrate the canonical configurational partition over possible piston positions with weighting $e^{-\beta p_{\mathrm{ext}} V}$. For an ideal gas in a cylinder of fixed cross-sectional area the piston coordinate is proportional to V, so its geometric measure is a constant times dV. Mechanical kinetic factors independent of V cancel from its normalised distribution.

For $Z_N(V) \propto V^N$, the probability density for volume is proportional to $V^N e^{-\beta p_{\mathrm{ext}} V}$ on the positive real axis. Normalisation uses $\int_0^\infty V^N e^{-aV} dV = N!/a^{N+1}$. The maximum, mean and variance are

$$V_{\mathrm{mode}} = \frac{N}{\beta p_{\mathrm{ext}}},$$

$$\langle V \rangle = \frac{N+1}{\beta p_{\mathrm{ext}}},$$

$$\mathrm{var}(V) = \frac{N+1}{(\beta p_{\mathrm{ext}})^2}.$$

The finite-N offset between mode and mean comes from the skewed piston-volume distribution. The most likely volume obeys the usual ideal gas law evaluated using N, while the exact mean differs in this simple piston model. At $N = 99$ the difference is about one percent and relative volume fluctuations are ten percent. At thermodynamic sizes the fractional difference vanishes. When calibrating nanoscale pressure from a histogram, a mean, a peak and a mean of reciprocal volume should therefore not be treated as interchangeable.

Check the force equilibrium directly. The instantaneous gas pressure conditional on V is $N k_B T / V$. 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

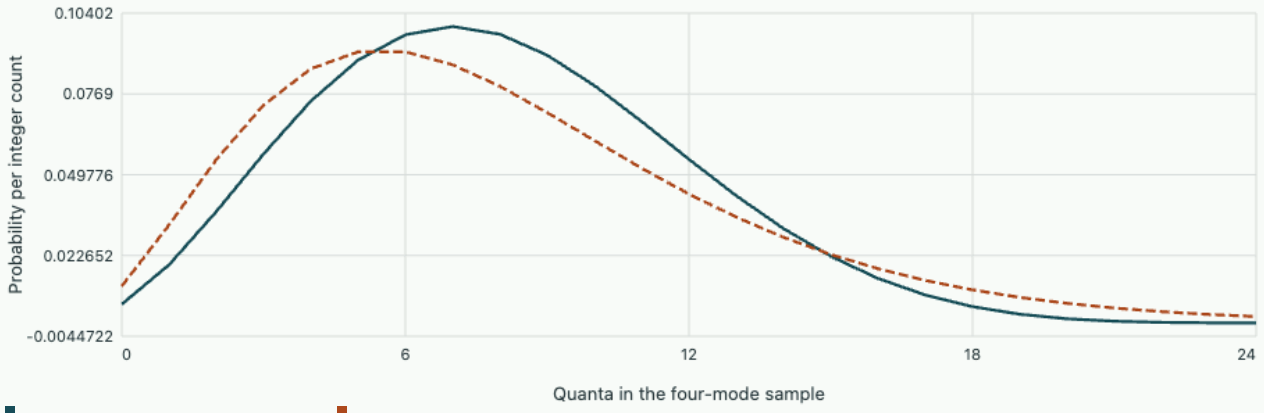
Change the reservoir size and compare the observed energy histogram with a canonical ensemble at the same mean.

Open this experiment online

Try the calculation

Reservoir modes **8**

Excitation quanta **24**



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

$$C_R = 3C_s = 180 k_B.$$

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 $V \int f(\mathbf{r}) d^3 r$ when boundary corrections are negligible. The logarithm of the configurational partition function retains connected contributions, a crucial step before taking the large-volume limit.

$$\frac{F}{N k_B T} = \log(n\lambda_T^3) - 1 + n B_2(T) + O(n^2),$$

$$B_2 = -\frac{1}{2} \int f(\mathbf{r})\, d^3 r.$$

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 $p = nk_BT(1 + nB_2)$ is consistent to this order. The factor two relative to the pressure correction arises because pressure contains n^2B_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:

$$p = \frac{Nk_BT}{L - Na},$$

$$\frac{p}{nk_BT} = \frac{1}{1 - na} = 1 + na + (na)^2 + \cdots,$$

$$n = N/L.$$

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

$p/(k_BT)$, 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 $x = e^{\epsilon/(400k_B)}$. 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

$$B_2^{\text{exchange}} = \mp \lambda_T^3/2^{5/2}.$$

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

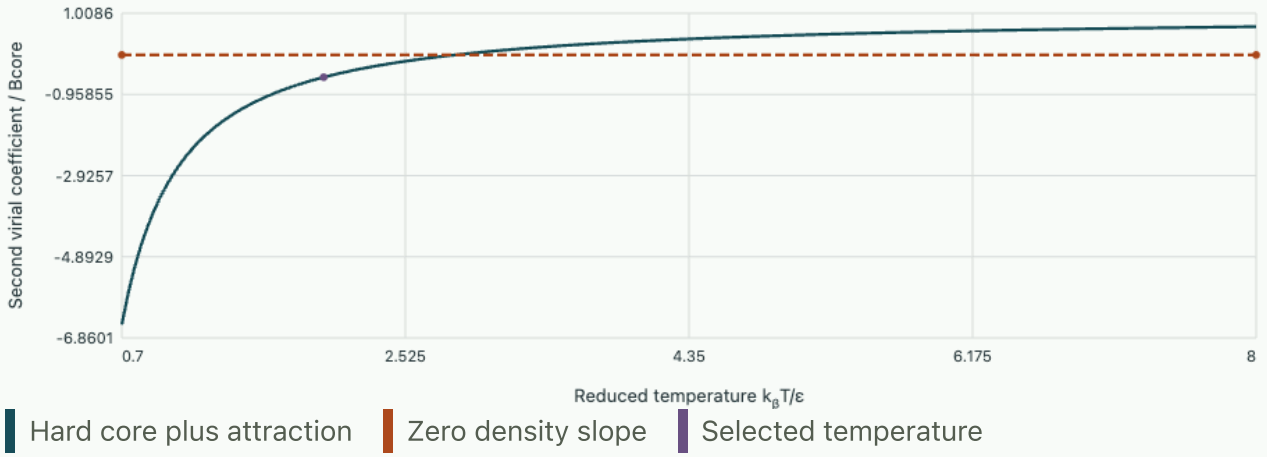
Vary the range and temperature of an attractive interaction. Identify where its attractive and excluded-volume contributions cancel.

Open this experiment online

Try the calculation

Outer radius / core diameter **1.5**

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



B_2/B_{core} = -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 $J_{\text{peak}} \simeq [\sqrt{2T/\theta_r} - 1]/2$, 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 $0.00454 k_B$, while at 1200 K it reaches $0.609 k_B$. 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 $2k_BT$, whereas the bulk mean is $3k_BT/2$. This selectivity occurs even when rotational and vibrational populations remain canonical.

To derive the difference, start with the bulk velocity density proportional to $e^{-mv^2/(2k_BT)}$. Write the outward normal speed as $v\cos\theta$, integrate over the outward hemisphere and include the velocity-space shell factor v^2dv. The resulting beam speed density is proportional to $v^3e^{-mv^2/(2k_BT)}dv$, 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

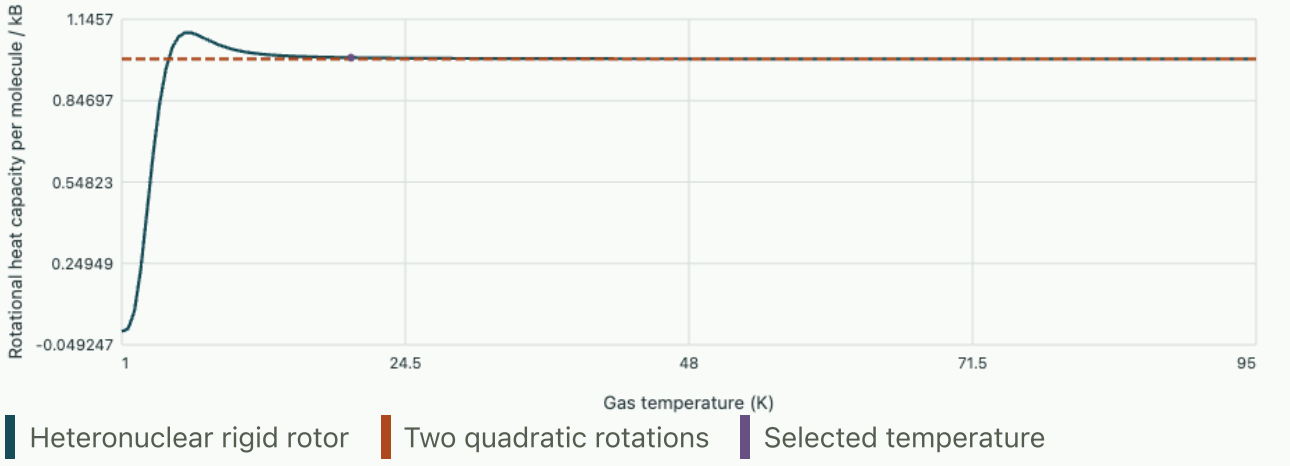
Sum the discrete levels of a linear rotor and compare the heat capacity with equipartition as temperature rises.

Open this experiment online

Try the calculation

Rotational temperature θrot (K) **8**

Chosen gas temperature (K) **20**



First excited rotational level has E_1/k_B = 16 K. At the selected temperature C_{rot}/k_B = 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 $0.70T_c$ the predicted condensate fraction is 0.414. Doubling density raises T_c by a factor $2^{2/3} = 1.587$. If imaging calibration rescales density but the reported temperature is independent, such a shift can move an inferred reduced temperature significantly. A finite trap or several spin components must be modelled before turning these example numbers into a laboratory calibration.

Why a uniform two-dimensional ideal gas is different

For free quadratic dispersion in d dimensions the density of states per volume scales as $\epsilon^{d/2-1}$. As $\mu \to 0^-$, occupancy behaves like $k_B T/\epsilon$ at low energy. The excited-population integral near its lower limit therefore behaves as $\int_0 \epsilon^{d/2-2} d\epsilon$. It converges only for $d > 2$. A uniform three-dimensional ideal gas saturates; in two dimensions the integral diverges logarithmically and in one dimension even more strongly. In the infinite uniform two-dimensional ideal gas at fixed density a finite positive temperature does not require a macroscopic ground occupation.

The two-dimensional integral can be evaluated explicitly. Per unit area $n_{2D} = -\lambda_T^{-2}\log(1-z)$. At phase-space density $n_{2D}\lambda_T^2 = 5$ this gives

$$z = 1 - e^{-5} = 0.99326.$$

In the formal grand description the mean ground occupation would be about 147 for this chemical potential, large on an atomic scale but still independent of total area at fixed density in the thermodynamic limit. Enlarging a system is how one distinguishes a finite prominent occupation from an extensive one.

A strictly finite box has discrete energies, so even a nominally two-dimensional sample can show a finite-size crossover with a narrow central state. Gapped transverse excitations may also make a nominally planar cloud effectively three-dimensional at sufficiently high temperature. Interacting two-

dimensional fluids can exhibit a vortex-mediated transition involving long-range phase correlations of a different kind. The ideal uniform density-of-states argument neither derives nor rules out that interacting transition.

Count a harmonic trap in phase space

In a three-dimensional harmonic trap with frequencies $\omega_x, \omega_y, \omega_z$, subtract the zero-point ground energy. Excited energies are $\hbar(\omega_x n_x + \omega_y n_y + \omega_z n_z)$ with nonnegative integers. When $k_B T$ exceeds all the level spacings by a large factor, approximate the number of states below energy ϵ by the volume of a positive octant tetrahedron in the three occupation-number coordinates. Its derivative yields

$$D_{\mathrm{trap}}(\epsilon) = \frac{\epsilon^2}{2\hbar^3 \omega_x \omega_y \omega_z},$$

$$N_{\mathrm{ex}} = \left(\frac{k_B T}{\hbar \bar{\omega}}\right)^3 g_3(z),$$

$$\bar{\omega} = (\omega_x \omega_y \omega_z)^{1/3}.$$

Condensation now gives

$$k_B T_c = \hbar \bar{\omega} [N/\zeta(3)]^{1/3},$$

$$N_0/N = 1 - (T/T_c)^3.$$

In a trap with geometric mean frequency 100 Hz and $N = 10^5$, the ideal estimate is 210 nK. At 0.70 of this trap-specific critical temperature the condensate fraction is 0.657, appreciably greater than the 0.414 box result at the same reduced temperature. Mixing a trap density of states with a homogeneous critical formula cannot be repaired by inserting one arbitrarily chosen cloud volume.

For a harmonic trap in d dimensions the low-energy density behaves as ϵ^{d-1}. Excited-state saturation in this semiclassical trapped model already occurs for $d > 1$. The changed criterion arises because a higher-energy particle also explores more physical space in the trap. The thermodynamic limit must be specified: holding all trap frequencies fixed while adding arbitrarily many particles changes the central density. A conventional trapped limit sends the trap frequencies to zero with a suitable combination such as $N\bar{\omega}^3$ fixed.

Connect a narrow peak to occupation and spatial coherence

After an ideal noninteracting cloud is released, a particle initially at $\mathbf{r}_0$ and momentum $\mathbf{p}$ reaches $\mathbf{r}_0 + \mathbf{p}t/m$. At long expansion times the position distribution maps approximately to momentum distribution after accounting for initial cloud size. A condensate occupying one trap orbital can therefore produce an extended but narrow component in a momentum image superposed on a broader thermal distribution. Ballistic expansion length, resolution and interaction-driven expansion all affect the observed width.

For the 87-unit-mass example at 200 nK a one-axis thermal velocity scale is

$$\sqrt{k_B T/m} = 4.37\,\mathrm{mm/s}.$$

After 20 ms its thermal contribution to one-axis rms width is about 87 micrometres. A point-spread width of ten micrometres is a small correction for this broad part but can dominate a genuinely much narrower central component. Fit integrated particle number and the spatial response of each component. A Gaussian peak height alone increases when the same number of atoms occupies a narrower instrument function.

A more fundamental criterion uses the one-particle density matrix. Its eigenvalues are orbital occupations; a condensate has an eigenvalue proportional to the total number as the appropriate system size grows. Interference measurements of first-order coherence or occupation scaling over different atom numbers can test this picture. Condensation is not identical to zero viscosity or a critical flow speed. An ideal uniform quadratic gas has low-energy excitations with arbitrarily small energy per momentum and does not by itself predict the finite critical speed of an interacting superfluid.

Solve the discrete number equation before passing to a continuum

A small cold gas in a box is an excellent numerical test because the discrete energy levels, their degeneracies and their individual occupancies are known. For a periodic cube with ground energy zero, write its excited energies as $\epsilon_{\mathbf{j}} = \epsilon_L(j_x^2 + j_y^2 + j_z^2)$, where $\epsilon_L = 2\pi^2\hbar^2/(mL^2)$. The first shell has six states, the next twelve and the third eight. Replacing each shell by one level without its vector degeneracy artificially drives particles into the ground state.

For a specified finite number N, solve $N = \sum_{\mathbf{j}} [e^{\beta(\epsilon_{\mathbf{j}} - \mu)} - 1]^{-1}$ with a chemical potential strictly below zero. This is a monotone root problem: as μ goes from minus infinity toward zero the sum grows from zero without bound owing to the ground state. Use an energy cutoff well above the available thermal energy and estimate the omitted tail, increasing the cutoff to check convergence. Near zero, evaluate $e^x - 1$ with a numerically stable small-argument function. A naive evaluation of $e^x - 1$ can erase the tiny denominator representing a macroscopically occupied orbital.

A useful artificial example has $\beta\epsilon_L = 2$ and a mean ground occupation of 99. This fixes $z = 99/100 = 0.99$. Each of the six first-shell levels contains on average $1/[e^2/0.99 - 1] = 0.1548$ particles, contributing 0.929 in total. The twelve next-shell levels at $4k_BT$ contribute another 0.222 together. Their combined occupation is already greater than one even though every individually excited orbital is weakly occupied. A code can compare an explicit vector enumeration with a shell-degeneracy table and verify exact agreement for the same cutoff.

The finite grand-canonical solution fixes mean N, not the integer count in each run. A canonical simulation with exactly N can yield similar mean populations in an appropriate large system but very different ground-state number fluctuations. In the grand ensemble one isolated orbital has a geometric number distribution with variance $\langle n\rangle(1 + \langle n\rangle)$. A mean ground population of 99 then corresponds to

a standard deviation close to the mean. A sealed sample of exactly 100 atoms cannot physically realise this unbounded distribution. Comparing its observed condensate fluctuations to the independent grand-canonical orbital formula without enforcing number conservation would be misleading.

The heat capacity sees the excited cloud, not a hidden latent heat

Below the bulk transition in a uniform ideal gas the fugacity tends to one. The energy expression becomes

$$U = (3/2)Vk_BT\lambda_T^{-3}\zeta(5/2),$$

which grows as $VT^{5/2}$. At fixed box size and fixed total atom number, the excited fraction adjusts with temperature while the ground state supplies or accepts the remaining particles. Differentiating gives

$$C_V = (15/4)Vk_B\lambda_T^{-3}\zeta(5/2).$$

Dividing by the total particle count gives a heat capacity per particle proportional to $(T/T_c)^{3/2}$.

At the critical temperature just below the transition, the coefficient per particle is $15\zeta(5/2)/(4\zeta(3/2)) \simeq 1.93k_B$. At half that temperature it is about $0.681k_B$. The energy and entropy are continuous at the uniform ideal transition; saturation of the excited population is not the same as a first-order phase change with a latent heat between two bulk densities. In the harmonic trap the different density of states yields a different power and a different shape of the heat-capacity anomaly. Thermodynamic measurements can therefore help distinguish trap geometry when imaging alone is ambiguous.

A calorimetric estimate of temperature often assumes the gas can re-equilibrate after a controlled energy pulse. In a perfectly noninteracting isolated gas, occupation numbers of stationary orbitals do not redistribute by themselves. Weak collisions, a reservoir or an explicit thermalisation protocol must establish equilibrium. The ideal gas can be an excellent equation-of-state model while weak interactions remain necessary for reaching that state on laboratory timescales. If the waiting time after a pulse is shorter than the redistribution time, substituting the equilibrium heat capacity into $\Delta T = \Delta U/C$ may falsely imply anomalous heating.

One practical diagnostic is to vary the waiting time while monitoring both the broad-wing kinetic temperature and central fraction. Thermal equilibrium predicts a definite relation between those two quantities at fixed total number and geometry. A lagging central fraction together with apparently stable wings suggests a slow redistribution process or atom losses. Repeating with matched final number and with independently calibrated detection efficiency is more discriminating than fitting temperature from the central peak alone.

Detect bosonic bunching without assuming a condensate

Occupation can also be probed through intensity or number correlations. For a single thermal bosonic mode with geometric distribution, the normally ordered second-order correlation is

$$g^{(2)}(0) = \langle n(n-1)\rangle / \langle n\rangle^2 = 2.$$

A coherent-state Poisson number distribution gives one. A precisely fixed number n gives $1 - 1/n$. These statements concern the statistics of the selected mode and the preparation. They must not be read as a universal test in which every condensed cloud has the same value under every ensemble or detection geometry.

In a detector pixel collecting M independent identically populated thermal modes, the mean total number is $M\bar{n}$, variance $M\bar{n}(1+\bar{n})$. Substituting into the factorial moment gives

$$g^{(2)}(0) = 1 + 1/M.$$

Ten unresolved thermal modes therefore give 1.10 rather than 2, even though each separately bunches fully. Finite temporal averaging, velocity-bin width and spatial point spread can increase the effective number of collected modes. A near-unity measured correlation does not by itself demonstrate that the gas has developed a macroscopically coherent orbital.

Why use a factorial moment? For independent shot noise the ordinary second moment contains a self-count term $\langle n\rangle$. Subtracting it isolates pair coincidences between distinct particles. A detector that cannot distinguish simultaneous events may preferentially drop exactly those pairs, creating artificial suppression. Two independent detection channels after a calibrated splitter can estimate coincidences without a single-pixel dead-time bias, provided cross talk and timing jitter are characterised.

A useful exercise varies both integration width and total atom number. A correlation excess falling with broader detection windows supports collection of multiple unresolved thermal modes. A central mode whose occupancy remains a fixed fraction as the whole gas grows supplies different information. Neither test replaces a correctly modelled spatial coherence measurement, yet together they constrain interpretations that can fit one mean momentum image.

How a number solver reveals its own near-critical sensitivity

At fixed temperature, derivative of total ideal boson occupation with respect to chemical potential is $\beta \sum_j n_j(1+n_j)$. A highly occupied lowest orbital can dominate the derivative even if the measured energy is primarily carried by the excited states. Hence a small error in chemical potential can produce a large error in number while barely changing calculated energy. A solver that stops when energy converges may report the wrong condensate fraction.

For a mean ground occupation 10^4, its chemical-potential distance below the lowest level is $k_B T \log(1 + 10^{-4})$, approximately $10^{-4} k_B T$. A shift in that tiny gap by ten percent changes the occupancy by roughly ten percent while leaving the apparent fugacity close to one to four significant

digits. Track chemical potential relative to the exact ground energy and report particle-number residual directly. Conversely, a small absolute error in a reported fugacity near saturation is not by itself evidence for a large uncertainty in the empirically counted atom number.

Finally, compare continuum estimates on several finite boxes at the same density. The leading discrepancy often comes from low discrete modes rather than the high-energy cutoff. Refine the low-energy calculation and the tail independently: enumerate a set of low shells exactly and replace only the dense upper spectrum by an integral. Check for double counting at the matching cutoff. An apparently excellent bulk critical temperature at one size can otherwise result from two cancelling approximations with different size dependence.

Infer a box density from a harmonically confined thermal cloud

Far from degeneracy a harmonically confined gas has classical spatial density $n(\mathbf{r}) = n(0)e^{-V(\mathbf{r})/(k_B T)}$. For an isotropic trap $V = m\omega^2 r^2/2$, each coordinate is Gaussian with variance $s^2 = k_B T/(m\omega^2)$. Normalising its full three-dimensional density gives

$$n(0) = N/[(2\pi)^{3/2} s^3].$$

An absorption image generally measures a column integral, whose central value is $N/(2\pi s^2)$, with units per area. Substituting this column density into a volumetric condensation criterion is dimensionally incorrect.

For $N = 10^5$, mass 87 unified atomic units, ordinary trap frequency 100 Hz and temperature 400 nK, the classical one-axis size is about 9.85 micrometres. The inferred central volume density is 6.65×10^{18} per cubic metre. The phase-space density estimate $n(0)\lambda_T^3$ is about 0.17; the classical profile is therefore a useful leading reference rather than a precision model at all radii. A ten-percent systematic underestimate of width inflates the inferred central density by about 37 percent because width enters cubed.

Near the transition the local-density approximation replaces the constant chemical potential by $\mu - V(\mathbf{r})$ inside the uniform-gas density function. This preserves the global chemical potential while reducing local effective fugacity away from the centre. Integrating the resulting local density over position recovers the harmonic-trap density-of-states counting in the semiclassical regime. A central phase-space density approaching the uniform saturation value can therefore coincide with a trap-wide critical atom number governed by $\zeta(3)$, not $\zeta(3/2)$. These statements use different spatial integrations and are consistent.

A practical reconstruction fits widths along multiple axes and includes known imaging point spread. For a Gaussian image convolved with a Gaussian response, variances add. A measured width of 10.0 micrometres and known response width 3.0 micrometres imply true width 9.54 micrometres, changing a cubed-density estimate by roughly fifteen percent. A narrow condensate component often has a non-Gaussian profile or a size comparable with the resolution, so the simple Gaussian deconvolution should not be applied indiscriminately to both components.

Compare inferred atom counts over increasing fields of view. Truncating broad thermal wings undercounts the thermal population and makes the apparent condensed fraction too large, particularly near onset. Independent normalisation from the complete cloud and fits to held-out outer regions help separate this effect from an actual change in the fraction occupying the lowest orbital.

There is a useful contrast with a single-occupancy adsorption site. Its occupation variance at mean fraction f is $f(1-f)$, whereas one thermal Bose orbital at mean occupation $\bar{n}$ exhibits $\bar{n}(1+\bar{n})$. At mean 0.50 these are respectively 0.25 and 0.75. The sign difference expresses whether multiple occupancy is excluded or enhanced. Vary the chemical potential in the single-site experiment and compare the response-to-noise identity before evaluating the Bose derivative on paper. Both obey the same grand-ensemble fluctuation formula, yet their number variances differ at the same mean because the allowed microscopic configurations differ. An observed sub-Poisson count can alternatively arise from detector dead time, and multiple unresolved thermal modes reduce bunching, so the ideal comparison is strongest for calibrated single-mode or single-site measurements.

Compare the measured single-site fluctuation law with the Bose orbital variance

Count the excited atoms in a trapped gas

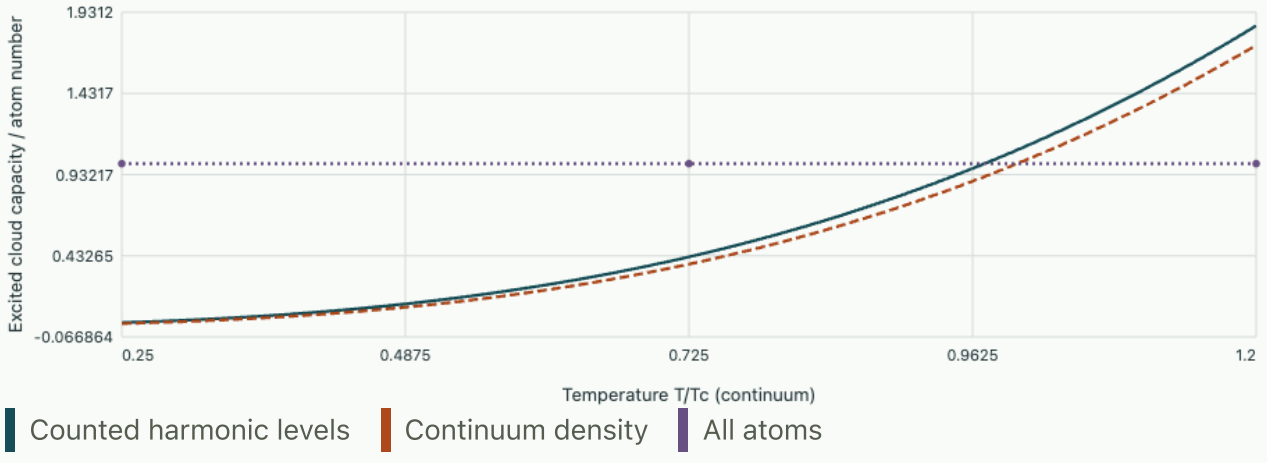
Change atom number and confinement frequency. Compare the discrete level sum with the continuum estimate and relate the difference to the density of states.

Open this experiment online

Try the calculation

Atom number **20000**

Isotropic trap frequency (Hz) **100**



Continuum Tc 122.52 nK. The discrete excited levels reach N at 119.17 nK.

Tutorial · Check what changes during an adiabatic box expansion

Below the transition a dilute ideal condensed box gas expands slowly and reversibly with fixed N. Determine how temperature and condensate fraction vary with volume. Would an isothermal expansion give the same outcome?

▼ **Use the entropy and critical-temperature scalings together**

Below transition the bulk chemical potential tends to zero. Since $U = 3pV/2$, $S = (U + pV - \mu N)/T$ is proportional to $VT^{3/2}$ for the saturated excited gas. At constant entropy $T \propto V^{-2/3}$. The box critical temperature at fixed N has exactly the same volume dependence, so T/T_c and the leading condensate fraction stay constant. In an isothermal expansion T_c falls while T stays fixed, decreasing condensed fraction and eventually exhausting it. The result rests on ideal-gas equilibrium and reversible expansion; a sudden opening into a second container is not that adiabatic path.

Tutorial · Determine whether a trap can be treated as three-dimensional

An anisotropic harmonic trap has ordinary frequencies 40, 40 and 8000 Hz at a measured $T = 100$ nK. Compute the spacings expressed as temperatures and assess using the three-dimensional semiclassical trap formula. Propose a consistent calculation near the dimensional crossover.

▼ **Identify which quantum numbers are actually sampled**

Using $h/k_B = 4.799 \times 10^{-11}$ K s, the level spacings are 1.92, 1.92 and 384 nK. Many transverse levels in the soft directions are populated, whereas a single quantum along the tight axis costs nearly four thermal energies. A three-dimensional integral over all three quantum numbers would overcount accessible low-energy states. Keep the tight-axis sum discrete while approximating the two soft directions semiclassically. Separate the exact global ground orbital from that hybrid excited-state count before solving for chemical potential at fixed total number.

At a few hundred nanokelvin the next tight-axis branches start to contribute appreciably. Observe widths or occupation across a temperature scan instead of announcing a fixed dimensionality based solely on trap shape. The model should recover a two-dimensional low-temperature limit and cross over smoothly toward three-dimensional counting when all spacings become small relative to thermal energy.

Tutorial · When is a ground-state peak extensive?

Three otherwise comparable measurements at fixed estimated areal density and temperature use planar boxes of areas in the ratio one, two and four. They yield approximately the same central-state

population of 150 atoms but total numbers 3000, 6000 and 12000. Compare with the ideal planar fugacity estimate above and explain what conclusion follows.

▼ Examine scaling and the resolution of a single mode

The fractions are 0.050, 0.025 and 0.0125, falling inversely with area. This is consistent with a highly occupied but nonextensive single level at fixed fugacity near 0.993. A homogeneous three-dimensional bulk condensate at fixed reduced temperature would instead have approximately constant fraction under the appropriate volume scaling. Absolute occupation larger than one is insufficient to distinguish the two regimes.

Confirm that the same experimentally defined central window isolates a single momentum mode as the box size changes. Allowed wavevector spacing shrinks with system size, so a fixed physical momentum aperture can sum an increasing number of levels. Correct for finite resolution, compare the complete momentum distributions and control interactions before treating the simple size dependence as a proof about the thermodynamic limit.

Cambridge Statistical Physics · ideal quantum gases and Bose condensation

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 $\mp\mu_B B$ 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 $D(E_F)\mu_B B$. 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^2D 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 $2k_B T \log 9$, approximately $4.39 k_B T$.

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 $m = \mu \tanh(\mu B/k_B T)$. Its small-field derivative gives volume susceptibility $\mu_0 n_i \mu^2/(k_B T)$. 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 $C k_B/(\mu_0 \mu_B^2) \simeq 1.28 \times 10^{25}$ 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.

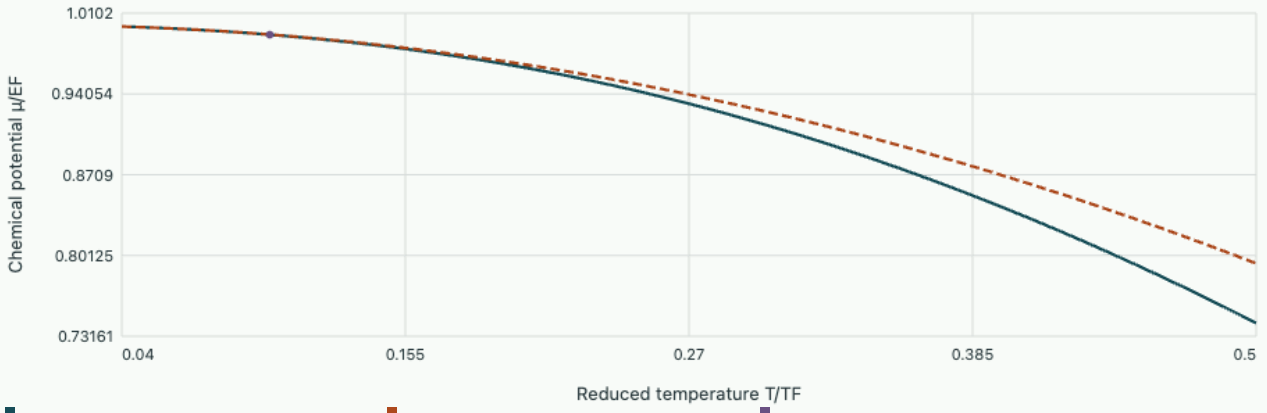
Decide which electronic states contribute to heat capacity

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 $2k_BT$, and put a transition configuration $10k_BT$ above the lower state. With a common attempt rate of 10^6 per second, the upward rate is 45.4 per second and the downward rate 335 per second. The stationary upper fraction is approximately 0.119 and the relaxation time 2.63 ms. Lower the barrier by $2k_BT$ without changing the two state energies. Both rates increase by $e^2 = 7.39$, while occupancy is unchanged. A detector recording only the equilibrium level histogram could miss the intervention entirely.

For a temperature scan, a straight line in the log rate against inverse temperature is evidence for an approximately constant activation energy and prefactor over that interval. A free-energy barrier includes entropic as well as enthalpic contributions; its numerical value can vary with temperature even if the microscopic potential does not. A change of dominant pathway can bend the apparent Arrhenius plot. Measure forward and reverse rates independently, then compare their ratio with an independently determined equilibrium population. A barrier fit explaining one rate but violating the equilibrium ratio usually indicates inconsistent state definitions, missed short visits or an unmodelled pathway.

A useful control is a weak static bias that lowers one well relative to the other. The equilibrium log-odds shift measures the difference between the state couplings to the bias. The individual log-rate shifts

additionally report the related dividing-surface response. If the state-occupation ratio moves as expected but only one rate changes appreciably, the effective transition configuration may be much closer in bias response to the other state. Such statements depend on the activated single-path interpretation; the response slopes themselves are more directly observed than an inferred barrier shape.

Use uniformisation to propagate a generator without negative probabilities

The formal finite-time propagator of a time-independent jump process is $P(t) = e^{Lt}$. A forward Euler update $p(t + \Delta t) = (I + L\Delta t)p(t)$ conserves total probability if the generator is assembled correctly, but a large step can yield negative diagonal probabilities. An exact algebraic representation starts with a rate q at least as large as every total exit rate. Then $K = I + L/q$ is a stochastic matrix in the same column convention. Write $L = q(K - I)$; since K and the identity commute,

$$e^{Lt} = e^{-qt} \sum_{n=0}^{\infty} \frac{(qt)^n}{n!} K^n.$$

This can be read as drawing a Poisson number of opportunities at rate q, followed by n transitions of K. At a lightly connected state some opportunities become self transitions. The construction leaves real waiting-time statistics correct while eliminating negative transition probabilities. Its truncation error in total probability is bounded by the omitted Poisson tail because every matrix power is stochastic.

For a two-state model with rates two and six per second choose $q = 6$ per second. Starting in the lower state, the first opportunity after some waiting time moves upward with probability one third and otherwise does nothing. In the upper state the next opportunity moves downward with probability one. At 0.25 s the Poisson mean is 1.5. Retaining terms through order eight leaves a tail below 3×10^{-5}. Compare the resulting upper probability against $(2/8)[1 - e^{-8(0.25)}] = 0.2162$. A small discrepancy due solely to truncation should shrink with the computable tail probability.

The representation also suggests a measured-data sanity check. A fitted generator with negative off-diagonal rates cannot yield a valid K at any positive q. A snapshot transition matrix can itself be nonnegative and stochastic but need not arise from a time-homogeneous continuous-time generator. In the two-state case its nontrivial eigenvalue must be $e^{-(a+b)t} > 0$. A measured negative eigenvalue rules out that simple generator model under perfect sampling, though finite counts and hidden periodic driving can distort empirical estimates. Taking a matrix logarithm blindly may return complex or negative rates rather than a physical kinetic explanation.

Population relaxation and a first-arrival experiment answer different questions

A first-passage measurement stops its clock when a target is first reached, even if that target would subsequently be left many times. To calculate its mean, treat the target as absorbing. For any other

state i, wait on average $1/r_i$, then add the mean remaining time from the next state weighted by its conditional transition probability. The backward equations are

$$-1 = \sum_j w_{ij}(t_j - t_i),$$

with $t_{\mathrm{target}} = 0$. Their solution depends on possible returns before the first successful passage. It generally differs from an inverse equilibrium occupation or from a slowest relaxation time.

Consider a three-site linear chain $0 \leftrightarrow 1 \leftrightarrow 2$ with equal rate k on each adjacent directed edge, starting at zero and stopping at two. From zero one must reach one after mean $1/k$. From one the mean wait is $1/(2k)$, then there is probability one half of returning to zero and one half of finishing. Solving $t_0 = 1/k + t_1$, $t_1 = 1/(2k) + t_0/2$ gives

$$t_0 = 3/k.$$

Ignoring returns would yield $1.5/k$, exactly half the correct first-passage mean.

The unmodified closed chain has equal stationary weights, but that fact alone says nothing about the numerical value of $3/k$. Speeding every link by ten leaves occupancies unchanged and reduces passage times tenfold. A detector that registers only visits longer than a threshold makes a brief first arrival invisible and can report even longer apparent passage times. Calibrate detector dead time and distinguish entrance to the target from persistence within it.

For a larger network, solve the sparse backward linear system and compare against many event-driven paths. Check boundary initial conditions by starting simulations exactly at the absorbing target: the first-passage time must then be zero under that convention. Some experiments instead ask for a first return after leaving, which requires a different initial condition. A zero-lag event in one definition is a physically different measurement in the other.

Estimate a cycle affinity with finite counting statistics

A clockwise ring trajectory contains gross counts in both directions on each edge. Estimating each conditional rate as count divided by origin-state exposure, the product of clockwise-to-anticlockwise rates around a complete closed loop cancels the exposures. The estimated cycle affinity is the sum of log forward counts minus log reverse counts, assuming every required count is nonzero and rates were stationary during the record. This is a useful cross-check independent of how the total time was partitioned among states.

If all three edges register roughly 120 clockwise and 80 anticlockwise events, the estimated affinity is

$$3\log(1.5) = 1.216.$$

For large independent Poisson-like gross counts the approximate variance of this log product is $3(1/120 + 1/80) = 0.0625$, giving standard error 0.25 as a first estimate. Actual path constraints couple counts; uncertainty can be assessed from independent records or a complete stochastic

likelihood. The point estimate alone does not tell whether an observed directionality is statistically distinguishable from zero.

One missing reverse event makes the unregularised logarithm infinite, but a finite record cannot establish an infinite physical affinity. Report a confidence bound based on the exposure and upper limit on the unseen reverse rate. Extend the acquisition or change sampling speed to exclude missed short-lived reverse states. Numerical addition of an arbitrary count can help regularise a plotting routine but is not a substitute for a specified prior or measurement model when claiming physical entropy production.

A driven occupation loop can carry charge without changing concentrations

Place three sites in a cycle and drive their transition rates by two controls, for example energies and barriers. Even if each instantaneous frozen generator satisfies detailed balance with a well-defined equilibrium distribution, the physical occupations need not equal that moving equilibrium distribution at each moment. A slow lag can produce a small edge current at every instant and a finite integrated current over a cycle. The current through an edge is time-integrated transfer; it can be nonzero while all occupations return to their initial values after a complete drive period.

A simple piecewise illustration avoids subtleties of adiabatic geometric pumping. In stage one, connect only sites zero and one while strongly favouring one. In stage two connect only one and two while favouring two. In stage three connect only two and zero while favouring zero. Begin with nearly all probability at zero and wait long enough for each enabled pair to relax. Almost the entire population travels once around the loop per cycle. No site accumulates probability over repeated complete cycles, but each edge registers about one forward crossing net per cycle. External control changes the energy and barrier landscape and supplies the required work.

If each programmed transfer succeeds with probability 0.95, it is tempting to announce a current of 0.95^3 per cycle. This is the probability of one fully successful trajectory starting on the prescribed site in a simplified no-return preparation, not generally the long-run edge current of a periodically driven ensemble. Failed transfers leave probability on other sites in later stages. The correct periodic distribution is the unit-eigenvalue vector of the product of the three finite-time propagators. Compute edge currents during each stage using its evolving occupations and integrate, then compare a complete cycle after transients have decayed.

This construction supplies a statistical check on experimental counting. Over a complete cycle the integrated net inflow at a site equals its change in occupancy probability. In a periodically stationary state that change is zero, so all three consistently oriented integrated edge currents must agree. A disagreement between measured edges can reveal a missed transition channel, a start-up transient or incorrect gating. Agreement does not prove absence of drive or detailed balance. A torque-free stationary ring and a three-stage external pump can share nearly the same time-averaged site histogram yet differ clearly in time-resolved edge currents.

Finally compare the work and heat bookkeeping with the master-equation entropy-production expression. During an external change of a site's energy without a transition, the state's energy changes by work. During a transition at frozen controls, assigning heat requires the energy difference and the relevant bath. Lumping an entire driven cycle into a single time-averaged generator can conceal those distinctions and underestimate dissipation.

The ordered protocol can also be run backwards. If the three stages are reversed while preparations are equilibrated in each corresponding periodic state, the preferential current direction should reverse in the ideal transfer picture. A fixed instrumental preference for counting neighbouring sites in one direction would not generally reverse. This distinguishes controlled transport from a direction-dependent detection efficiency. Do not confuse this observed current reversal with exact microscopic time reversal of each individual stochastic trajectory: the reversed experiment has its own ensemble of random jumps. Comparing path probabilities, rather than pairing individual paths as though their noise histories were reproducible, is the appropriate starting point for entropy-production and fluctuation-relation tests. The distinction becomes particularly visible when a stage transfers almost all the probability in one direction while a few rare trajectories travel against the applied protocol.

Distinguish stationarity from detailed balance

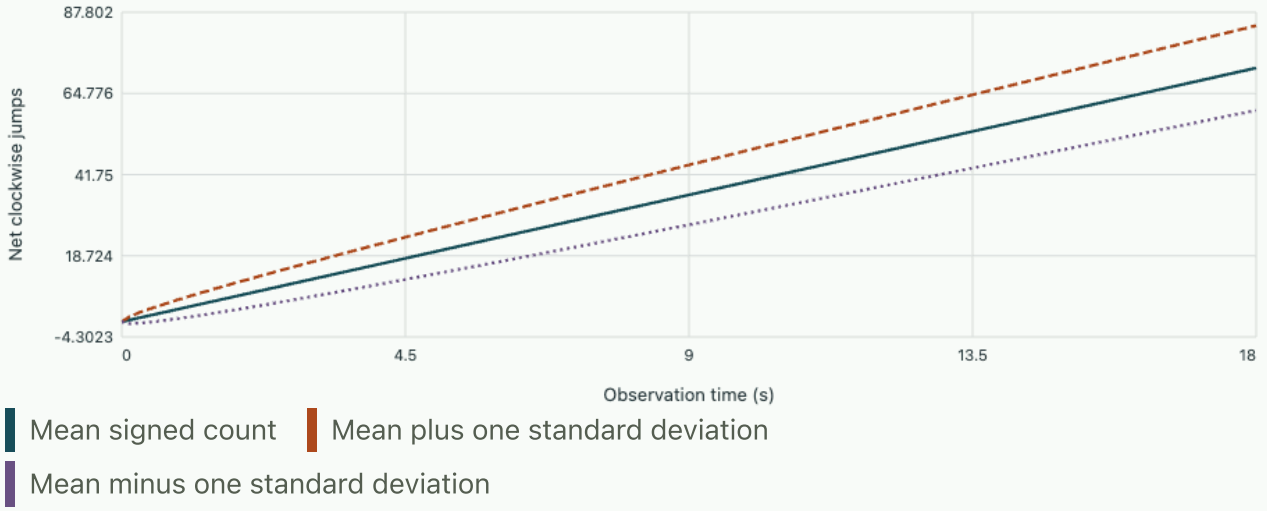
Keep the three populations identical while varying the directed transition rates. Identify the observable that reveals time-reversal asymmetry.

Open this experiment online

Try the calculation

Clockwise bias **0.5**

Total jump rate (s^{-1}) **8**



Every stationary population is 1/3. Net clockwise edge current 1.3333 s^{-1}. Entropy production 4.3944 kB/s.

Tutorial · Recover a generator from snapshots rather than full events

At sampling interval 0.10 s, a two-state record yields $P_{01} = 0.12$ and $P_{10} = 0.28$. Infer the continuous-time rates, and compare with directly dividing each transition probability by 0.10 s.

▼ **Use the exact finite-interval propagator**

The exact transition probabilities are $P_{01} = p_*[1 - e^{-(a+b)\Delta t}]$ and $P_{10} = (1 - p_*)[1 - e^{-(a+b)\Delta t}]$. Their sum is 0.40, so $a + b = -\log(0.60)/0.10 = 5.108$ per second. The stationary upper fraction is

$$0.12/0.40 = 0.30.$$

Thus $a = 1.532$ and $b = 3.576$ per second. Dividing by the interval gives only 1.2 and 2.8 because return transitions inside a sampling interval can erase evidence of the intermediate visit.

Tutorial · Can measured rates arise from one equilibrium potential?

For a three-state triangle suppose clockwise to counterclockwise conditional rate ratios are 2, 3 and 0.25 around the three distinct edges. Could these be equilibrium weights in a common bath? If a single uncertain ratio were corrected, what would it have to be on the third link?

▼ **Apply the closed-loop consistency condition**

The product is $2 \cdot 3 \cdot 0.25 = 1.5$, giving nonzero loop affinity $\log 1.5 = 0.405$. No consistent detailed-balance set of three state weights exists for these exact ratios. The third ratio would have to be $1/6$, giving a telescoping unit product. If counts are noisy, propagate uncertainty in the log ratios or compare a constrained detailed-balance likelihood against an unconstrained model. A value barely inconsistent with unity within one sparse record is not enough to establish driving.

Tutorial · Detect a hidden state with the same observed level

A detector reports two visible levels, but its lower level combines two unresolved substates. The measured lower-state survival probability is $0.5e^{-t} + 0.5e^{-9t}$ with time in seconds. Calculate the initial hazard and its long-time limit. Can a two-state Markov fit reproduce both?

▼ **Differentiate survival before interpreting the tail**

The hazard is minus the derivative of $\log S$. At zero time it is the initial weighted mean of rates, 5 per second. At long times the rapidly departing component has almost disappeared and the hazard tends

to 1 per second. A single Markov exit rate would give a constant hazard for all dwell ages. The observed model cannot reproduce both limits with one rate.

A changing apparent hazard may also result from limited detector temporal resolution or variable preparation of the substates. Repeat with shorter acquisition intervals and condition the dwell distribution on the preceding visible transition or an independent preparation marker. A fitted sum of exponentials identifies a need for a richer observation or state model; it does not uniquely identify microscopic structures absent in the measurements.

Cambridge Kinetic Theory · stochastic processes and approach to equilibrium

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 $\log f = a + \mathbf{b} \cdot \mathbf{v} + c v^2$, 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 $m v_i$ 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 $P_{ij} = n k_B T \delta_{ij}$. 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:

$$P_{xy} = -\tau \frac{m^2}{k_B T} \partial_y u_x \int c_x^2 c_y^2 f_M d^3v = -n k_B T \tau \, \partial_y u_x .$$

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 $mv^2/(2k_BT) = 1$, 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

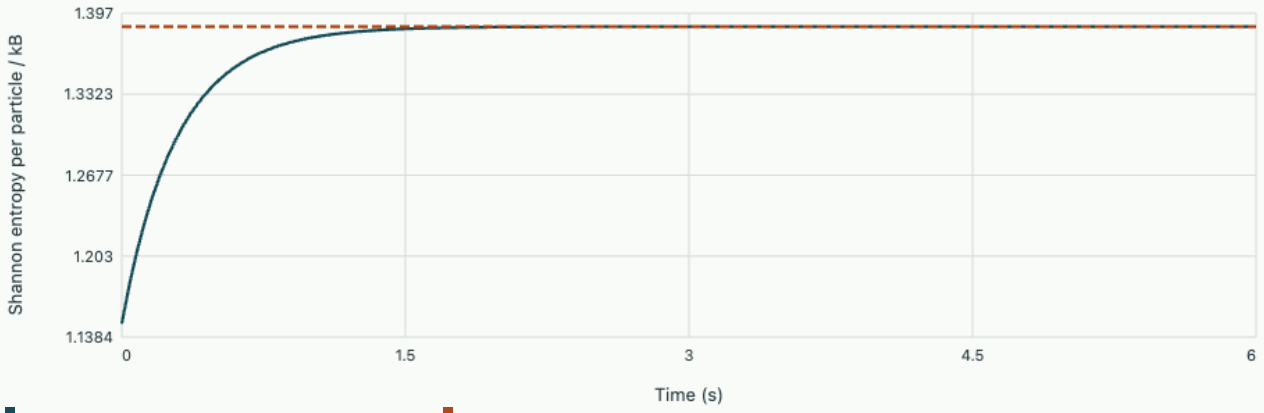
Prepare unequal populations in two velocity pairs. Count forward and reverse collision rates and compare the sign of entropy change with the population imbalance.

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 $2s_x^2$ 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

$$D_{xy}^2 \le D_{xx}D_{yy}.$$

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

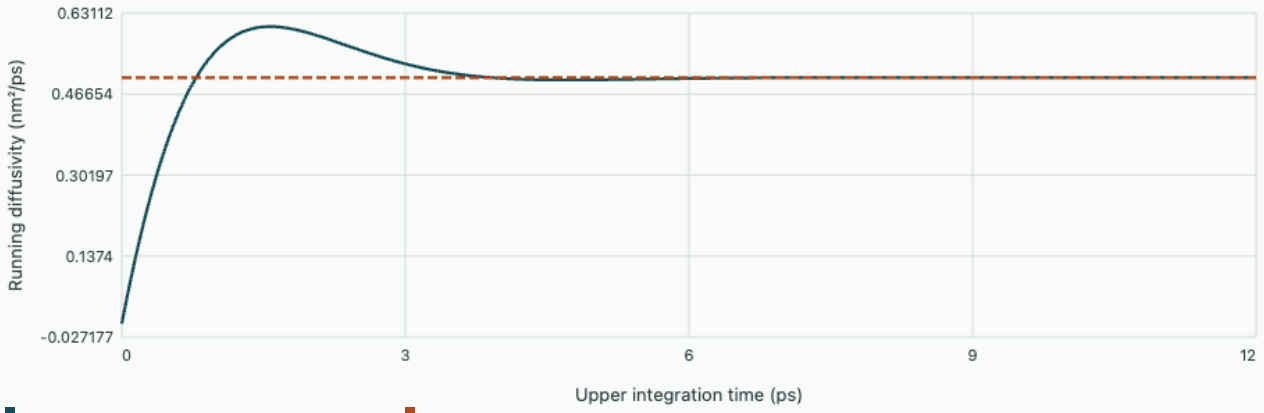
Integrate an oscillatory velocity correlation over a finite record. Compare a short apparent plateau with the independently computed infinite-time result.

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 $\beta N(\langle m^2\rangle - \langle|m|\rangle^2)$. 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 $k_B T/J$. 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 $T_c = 2.27 J/k_B$. 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 $K = J/(k_B T)$. 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

$$\partial_\beta \langle A \rangle = -\mathrm{cov}(A, E).$$

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

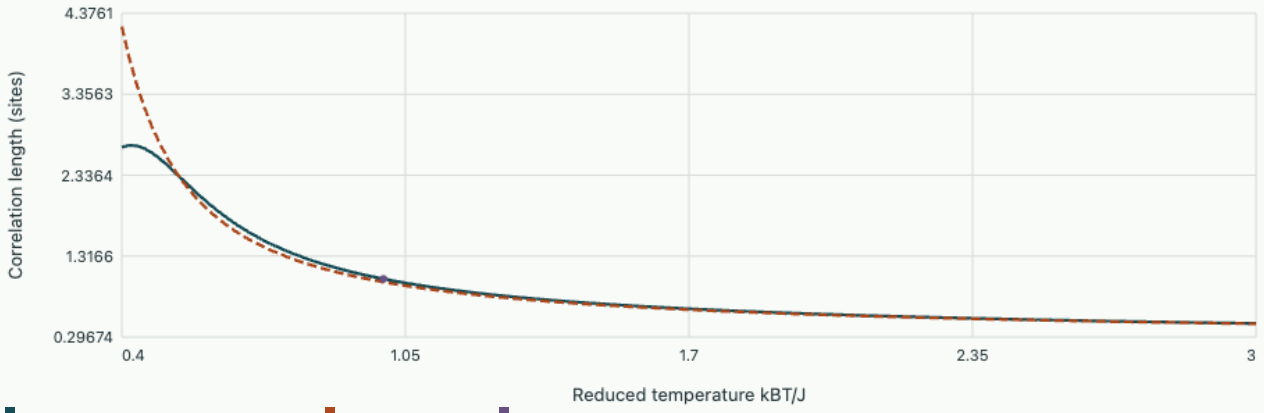
Determine how the subleading transfer-matrix eigenvalue controls the length seen in a colour measurement. A finite chain much shorter than this length can look ordered.

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 $J=J_0e^{-\Delta G_*/(k_BT)}$ 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 $P_0(t+dt)=P_0(t)(1-JVdt)$: $P_0=e^{-JVt}$. Mean wait is $1/(JV)$, median $\log 2/(JV)$. A 1.0 microlitre sample with $J=2.0\times10^8$ 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_vV+J_sA. 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 $P_0(t)=\exp[-\int_0^t J(T(s))Vds]$ 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 $\Delta g(v) = \Delta g_0 - av$, 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 $\Delta G/(k_B T) = 4v^{2/3} - 4v + 0.25v^2$, with a specified arbitrary reference volume absorbed into v. Its derivative is $(8/3)v^{-1/3} - 4 + 0.5v$. 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 $h^{N_e} e^{-hT_{\mathrm{exposure}}}$. 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 $5.0k_BT$ and variance $4.0(k_BT)^2$, the equality predicts $\Delta F = 3.0k_BT$ if the Gaussian law including tails is justified. With equal-variance reverse statistics, the corresponding mean of $-W_R$ is $1.0k_BT$. The two densities cross at $3.0k_BT$. 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_BT 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 $\operatorname{var} W = 2k_B T\langle W\rangle$. 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 $\Delta F = -0.566k_BT$. Mean work is only $-0.238k_BT$, 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 $6k_BT$ 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 $0.5k_BT$, this Gaussian-noise effect is $-0.125k_BT$. 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

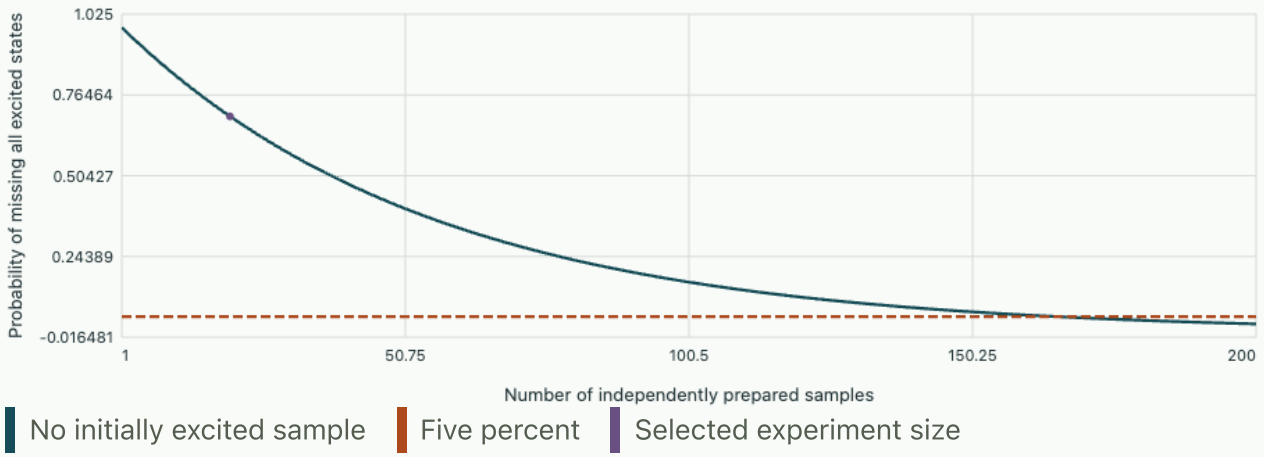
Change the preparation gap and the number of repetitions. Compare the average work, the exact free-energy change and the chance of missing the rare branch entirely.

Open this experiment online

Try the calculation

Initial gap (kBT) **4**

Independent preparations **20**



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 $4.0 k_BT$. 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 $\sqrt{0.2^2 + 0.3^2}/2 = 0.18 k_BT$ 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,

$$\beta_{n+1/2}\varepsilon = \frac{S_{n+1} - S_n}{k_B} = \ln\frac{N-n}{n+1},$$

$$\beta = \frac{1}{k_B T}.$$

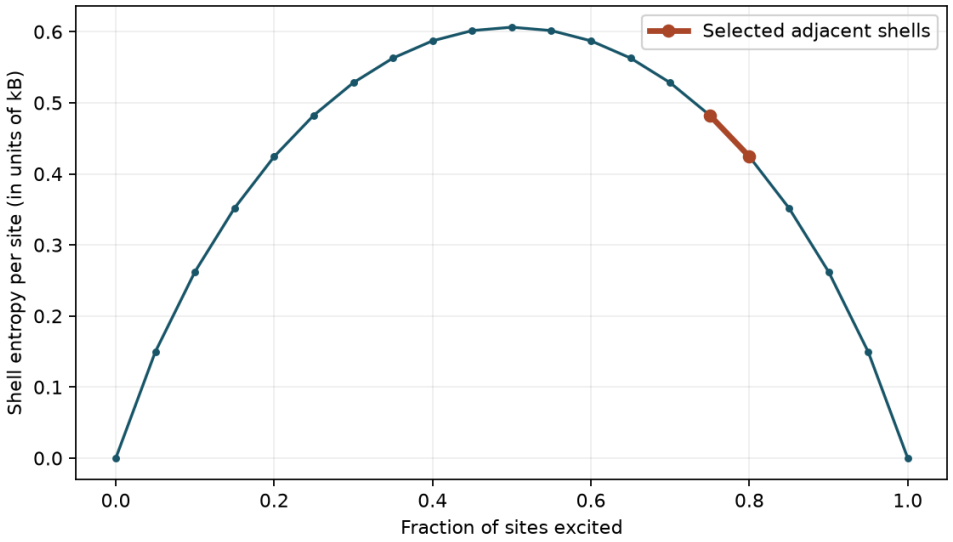
Adding an excitation near the bottom initially opens more arrangements. Beyond the middle, it closes them. The decreasing entropy on this upper branch assigns a negative inverse temperature to the population. In the limit of large N at fixed excitation fraction p, the slope becomes $\ln[(1-p)/p]$. The 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.

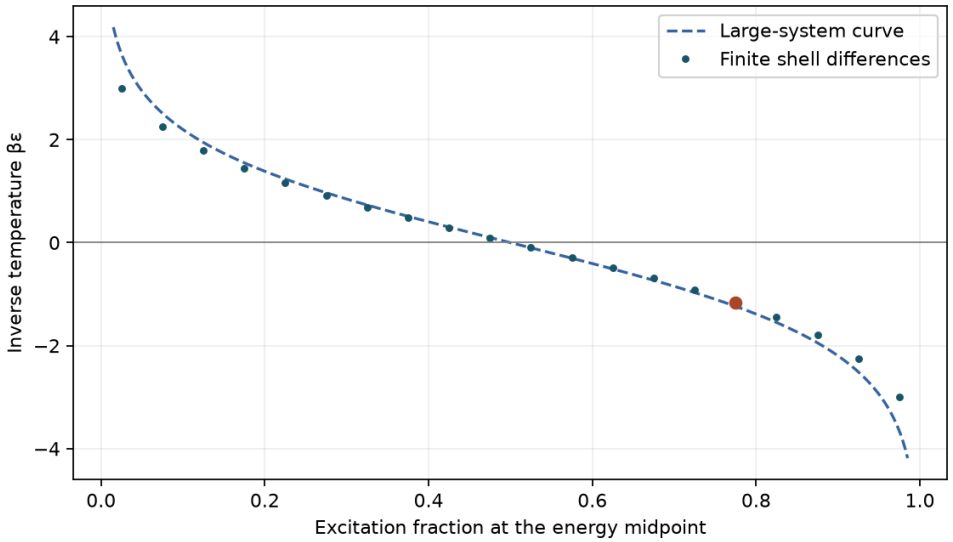
Predict the direction of an exchange

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

$$100 p\varepsilon = 75\varepsilon.$$

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

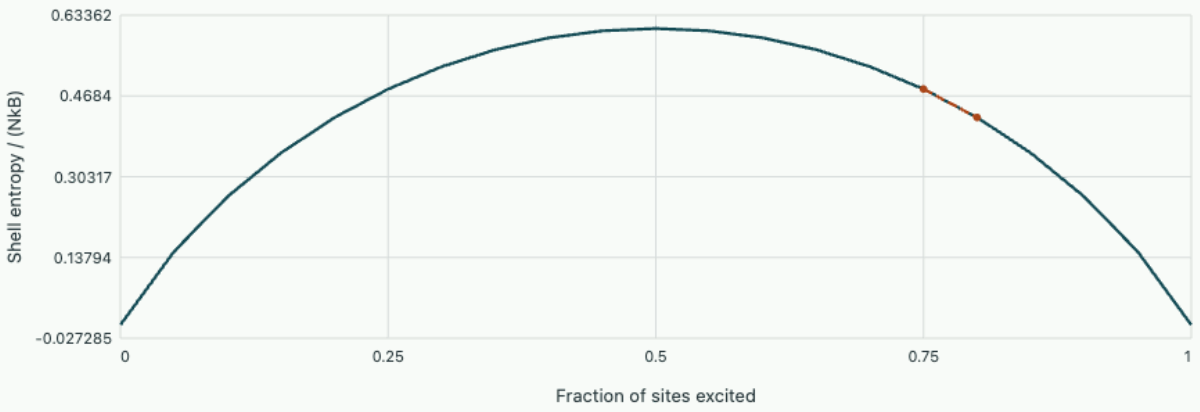
Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

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

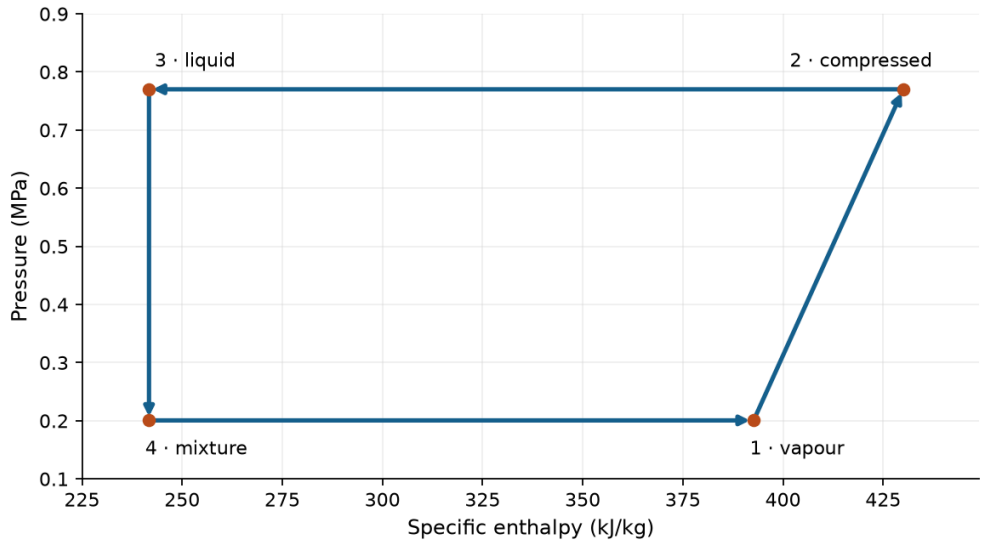
For R-134a take saturated vapour leaving at −10 °C and saturated liquid leaving the condenser at 30 °C. Rounded ASHRAE saturated-fluid data give $p_1 = 0.20060$ MPa, $h_1 = 392.66$ kJ/kg and $h_{f,1} = 186.70$ kJ/kg. At 30 °C the pressure is 0.77020 MPa and $h_3 = 241.72$ kJ/kg. Suppose the measured 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

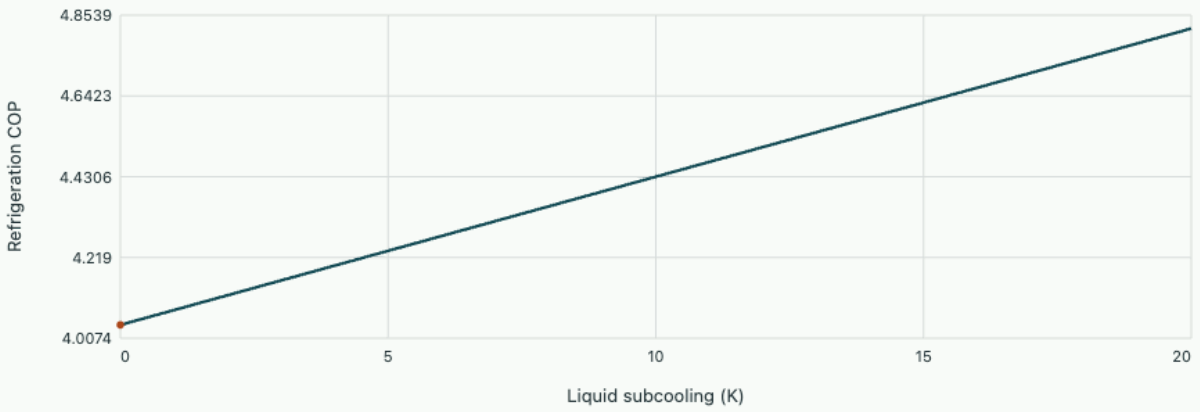
Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

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.

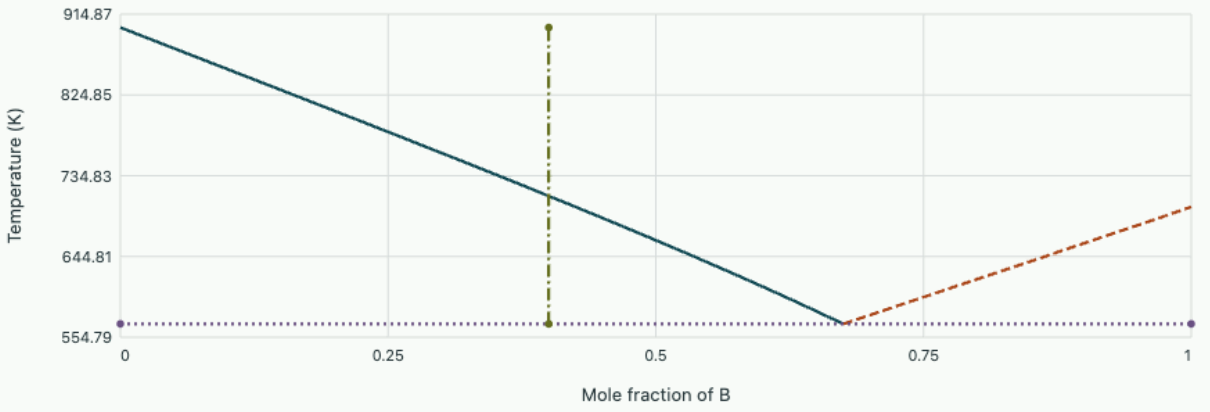
Open this experiment online

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

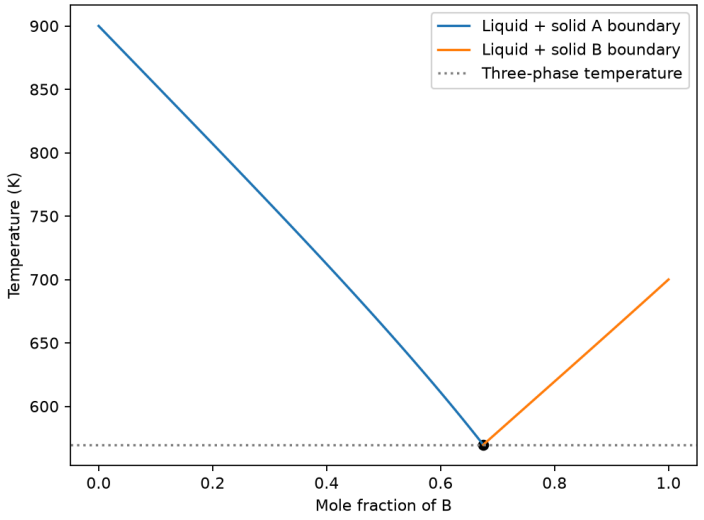
def liquid_fraction(t, melting, heat):
    return np.exp(-heat / gas_constant * (1/t - 1/melting))

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

$$T_f - T_{\text{freeze}} \simeq \frac{RT_f^2}{\Delta h_{\text{fus}}} x_s.$$

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 $\mu_g^* - \mu_l^* \simeq -(\Delta h_{\text{vap}}/T_b)(T - T_b)$ gives the opposite temperature shift.

$$T_{\text{boil}} - T_b \simeq \frac{RT_b^2}{\Delta h_{\text{vap}}} x_s.$$

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, $\Delta h_{\text{fus}} = 6.01$ kJ/mol. A 0.10 mol/kg nonelectrolyte gives a freezing-point depression near 0.186 K. Using $T_b = 373.15$ K and $\Delta h_{\text{vap}} = 40.7$ 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.

$$E = \frac{RT}{F} \ln \frac{a_{\mathrm{high}}}{a_{\mathrm{low}}}.$$

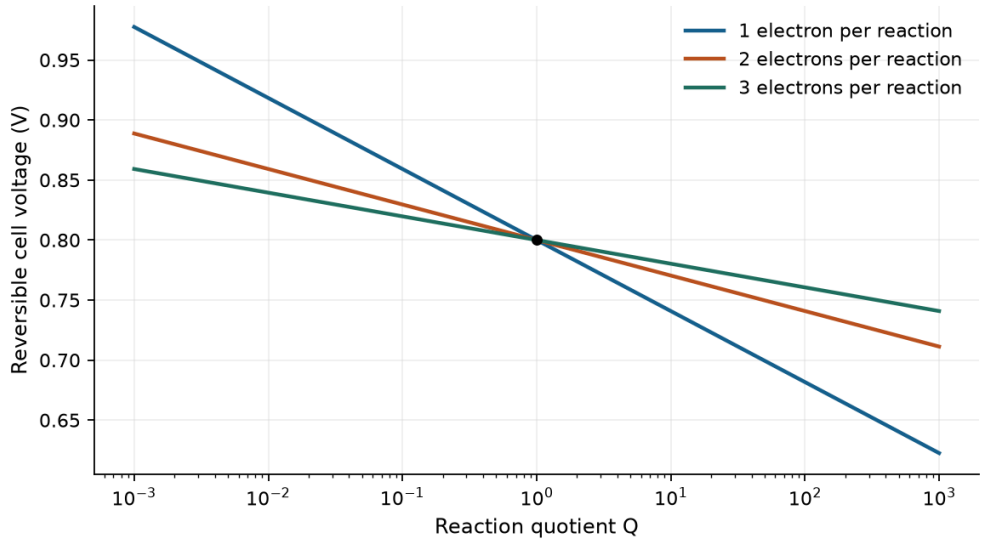
At 298.15 K a tenfold ratio gives 59.16 mV. Doubling both activities leaves the voltage unchanged for ideal activities. In a nonideal solution measure activity coefficients or use a calibrated reference. A voltmeter that draws significant current changes the surface concentrations and thus the very chemical 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.

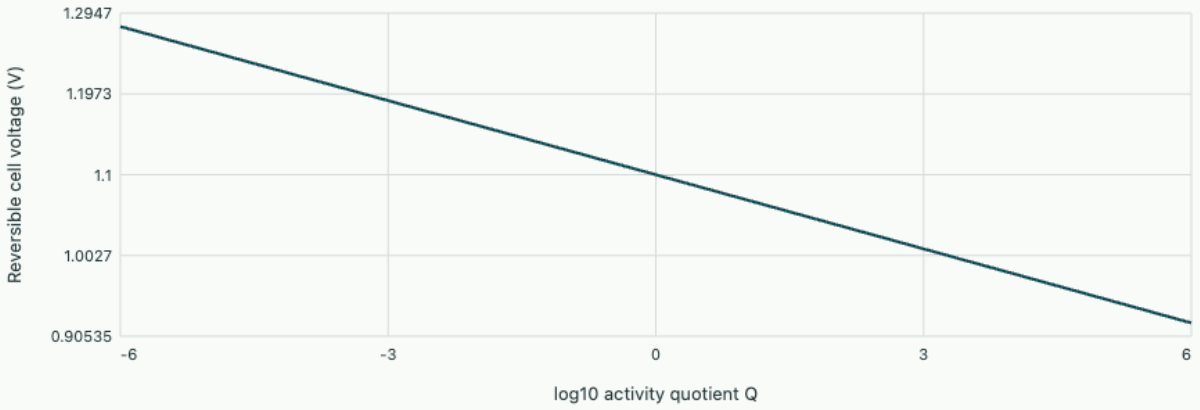
Try it · Nernst

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.