Counting, correlations and thermal measurements

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

Contents

1. **How many lattice sites remember the colour of a neighbour?**
   A chain of three-state spins gains energy when adjacent colours agree. A weak field favours one colour over the other two. Use a transfer matrix to calculate both free energy and the distance over which correlations survive. Recover the same numbers from OneSimulation eigenvalues without relying on symmetry reduction.
2. **Measure the heat capacity of a reservoir from the width of an energy histogram**
   Two small Einstein solids share a fixed number of excitation quanta. A calorimeter records how many are in the first solid. Count the allowed microstates, compare the exact distribution with a reservoir approximation and see how energy conservation correlates the readings.
3. **Recover a free-energy difference from a distribution of work**
   A two-state system starts in equilibrium, then a control jumps its energy gap. Individual work readings differ although every run uses the same control protocol. Use the elementary model to derive an exponential work identity, compare it with the average work and estimate how many runs it takes to see the events that dominate the estimate.
4. **Why a molecular beam and its source have different measured speeds**
   A tiny hole lets a dilute gas escape into a vacuum. A rotating chopper admits a short pulse and a downstream detector records arrival times. Derive the source speed density, the effusive speed density and the actual time-of-flight histogram. Each samples the same Maxwellian physics with a different weighting.
5. **Detect irreversible steady motion from transition counts**
   A three-state system appears equally likely to occupy each state. Is it at thermal equilibrium? Build two different stationary processes with identical populations, then determine what a time-resolved instrument must count to distinguish them.
6. **Why macroscopic measurements become reproducible**

The same collection of microscopic switches can teach us two different facts about a macroscopic sample. Its mean becomes predictable, and the fluctuations about that mean acquire a particular shape. Derive both statements, then find an experiment in which only the first prediction is reliable at the available sample size.

7. Weigh a colloid by counting where it spends its time

Brownian motion spreads a dilute suspension through a cell whilst buoyancy biases its vertical position. Derive the equilibrium concentration, predict counts in finite camera bins and estimate a buoyant mass without measuring how quickly any individual sphere falls.

Counting, correlations and thermal measurements · 1. How many lattice sites remember the colour of a neighbour?

A chain of three-state spins gains energy when adjacent colours agree. A weak field favours one colour over the other two. Use a transfer matrix to calculate both free energy and the distance over which correlations survive. Recover the same numbers from OneSimulation eigenvalues without relying on symmetry reduction.

Find a local weight whose product counts each energy once

Each site has colour $a = 0, 1, 2$. A matching pair contributes energy $-J$, with $J > 0$, and a site of colour zero earns an additional $-H$. In a periodic ring each site belongs to two adjacent bonds. Give each of those bonds half the local field energy so the product of their weights counts each field term exactly once. Define $K = J/(k_B T)$ and $h = H/J$.

$$M_{ab} = \exp\left[K\left(\delta_{ab} + \frac{h}{2}(\delta_{a0} + \delta_{b0})\right)\right],$$

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

To see the trace, write the partition function as a sum over $a_1, \ldots, a_N$ of the product $M_{a_1 a_2} \cdots M_{a_N a_1}$. Summing an intermediate index is matrix multiplication; closing the initial and final index is the trace. This is exact even for a short periodic chain defined by that bond count. The transfer matrix contains no time step: multiplying it advances spatial position.

Compare the leading eigenvalue with the next one

Without bias the matrix has diagonal e^K and all other entries one. The uniform vector has eigenvalue $e^K + 2$. Every vector with components summing to zero has eigenvalue $e^K - 1$, twice degenerate. Thus $Z_N = (e^K + 2)^N + 2(e^K - 1)^N$. The thermodynamic free energy per site is $-k_B T \ln(e^K + 2)$. To measure a connected spatial correlation, insert a diagonal colour observable at sites separated by r transfer operations. The uniform contribution cancels from the connected part; each nonuniform mode contributes a factor $(\lambda_j/\lambda_1)^r$ when its overlap is nonzero. The longest correlation length in sites is

$$\xi^{-1} = \ln(\lambda_1/|\lambda_2|).$$

At $k_B T = J$ and zero bias, $\lambda_1 = 4.71828,\ \lambda_2 = 1.71828$ and $\xi = 0.9900$ sites. At low temperature e^K is large and $\xi \simeq e^K/3$. A finite chain can then appear ordered simply because it is shorter than its correlation length. There is no finite positive temperature at which these two one-dimensional eigenvalues cross.

Keep track of a mode that a chosen detector cannot see

At $h \neq 0$, colours one and two remain interchangeable. Their difference vector is an eigenmode of eigenvalue $e^{K} - 1$. In the orthonormal basis comprising colour zero and the sum of one and two, the remaining block is

$$\begin{pmatrix} e^{K(1+h)} & \sqrt{2} e^{Kh/2} \\ \sqrt{2} e^{Kh/2} & e^{K} + 1 \end{pmatrix}.$$

Its eigenvalues follow from the trace and the square root of the squared diagonal difference plus four times the squared off-diagonal element. Sort all three eigenvalues. A detector sensitive only to whether the spin is zero is invariant under interchange of colours one and two. It has no overlap with their odd difference mode, even if that mode sets the mathematically longest length. A resolved measurement of colours one and two can see it. Quoting a correlation length from eigenvalues requires specifying which observation can actually couple to that eigenvector.

Change temperature and examine the spectral gap

Vary reduced temperature and reduced field H/J independently. The plot compares the longest colour-correlation length to the exact zero-field result. At a chosen point the downloadable model constructs the full three-by-three matrix and computes its modes using OneSimulation's contour eigensolver. It reports free energy per site and all three eigenvalues. In this small system the direct reduction above is available as an independent check. The same matrix construction generalises to more states even when no colour symmetry reduces the size.

Compare two chains at the same temperature. Set the colour bias to zero and apply a positive bias of 0.30 J. Use the transfer-matrix eigenvalue ratio to explain the change in the longest correlation length.

Open this experiment online

Try the calculation

Reduced temperature kBT/J **1**

Reduced colour field H/J **0.15**

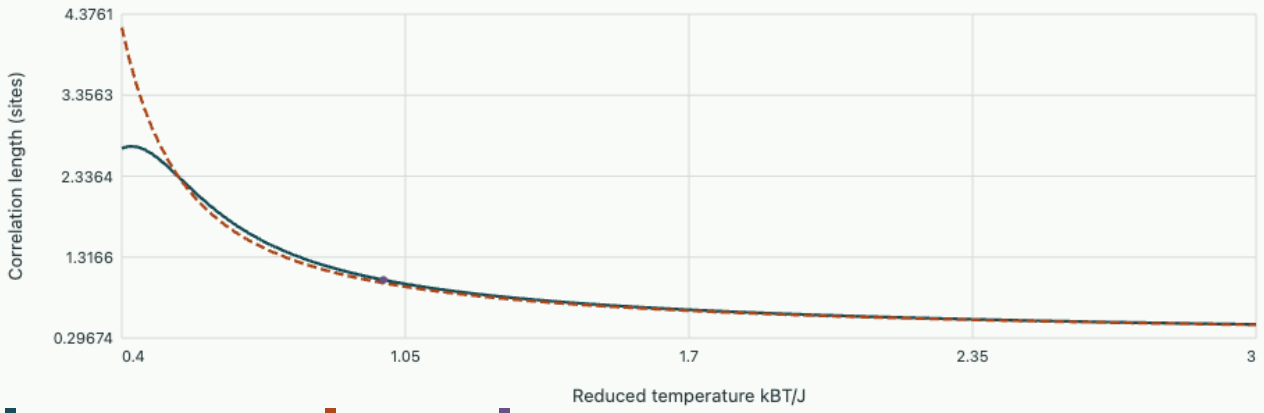

Selected colour bias Zero bias Selected temperature

Leading eigenvalue 4.9881, second 1.8884, longest length 1.0295 sites.

Download the OneSimulation model

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

```
python spectra-and-orbits.py chain --temperature 1 --bias 0.15
```

Tutorial

At $k_B T = J,\ h = 0$ what fraction of the asymptotic leading contribution to a connected correlation survives after three lattice sites?

▼ Work through the solution

The factor is $(1.71828/4.71828)^3 \simeq 0.0483$. The correlation observable sets its initial amplitude; the eigenvalue ratio sets this fraction.

Why does squaring a single bond weight to represent the effect of two bonds give the wrong two-step transition amplitude?

▼ Work through the solution

Two steps can pass through any of the three intermediate colours. Matrix multiplication sums all possibilities $M_{ab}M_{bc}$ over b. Squaring the single number M_{ac} loses those intermediate configurations.

Suppose only the leading eigenvalue is reliably resolved. Which thermodynamic-limit prediction is possible and which remains undetermined?

▼ Work through the solution

One can obtain free energy per site, and with its parameter dependence the derivatives giving energy or magnetisation. A connected spatial length requires a subleading eigenvalue and overlap of the chosen observable. A precise leading free energy alone cannot give that numerical length.

Compare this transfer calculation with the two-state Ising chain experiment.

Counting, correlations and thermal measurements · 2. Measure the heat capacity of a reservoir from the width of an energy histogram

Two small Einstein solids share a fixed number of excitation quanta. A calorimeter records how many are in the first solid. Count the allowed microstates, compare the exact distribution with a reservoir approximation and see how energy conservation correlates the readings.

Count the ways to store an integer number of quanta

Let solid A contain a independent harmonic modes and solid B contain b, all with energy spacing ε. Omit their fixed zero-point energy. The total excitation count Q is fixed. In one solid a configuration is an ordered list of nonnegative mode occupations. An arrangement of q identical excitation symbols and $a-1$ separators specifies such a list. The multiplicity is $\binom{q+a-1}{a-1}$. The count for distinguishable packets would be different, because exchanging two excitations between identical occupations introduces no new configuration here.

$$P(q) = \frac{\binom{q+a-1}{a-1}\binom{Q-q+b-1}{b-1}}{\binom{Q+a+b-1}{a+b-1}},$$

$$q = 0, \ldots, Q.$$

The denominator counts all occupations of the $a+b$ modes; summing the numerator over every partition must reproduce it. As a simple check, set $a=2, b=1, Q=2$. The unnormalised weights for $q=0,1,2$ are 1, 2 and 3. The most probable count is 2, but the mean is $4/3$. "Thermal equilibrium" does not require every observed energy to coincide with the most likely energy.

Separate the mean from the fluctuations

By permutation symmetry each individual mode holds mean $Q/(a+b)$. Therefore $\langle q \rangle = Qa/(a+b)$. For the second moment, introduce an auxiliary variable x marking quanta in solid A. The coefficient of z^Q in $G(x,z) = (1-xz)^{-a}(1-z)^{-b}$ is the unnormalised probability generating function. Differentiating once in x weights by q, twice by $q(q-1)$. At $x=1$ the second derivative is $a(a+1)z^2(1-z)^{-(a+b+2)}$. Divide its coefficient by the total multiplicity and combine first and second moments. Writing $s = a+b$ gives

$$\operatorname{var}(q) = \frac{Qab(s+Q)}{s^2(s+1)}.$$

The second solid contains exactly $Q-q$. Its variance is identical, and the covariance of the two measured energies is minus their common variance in energy units. A pair of independently fitted canonical distributions cannot reproduce this anticorrelation. At $Q=24,\ a=4,\ b=8$ the mean number in A is 8 and the variance is about 14.77. These are fluctuations of a perfectly isolated combined system, not a leak in the calorimeter.

Background · Can two noisy calorimeters expose the true energy variance?

The two solids exchange energy while their sum is fixed. If the first energy fluctuates upward by δE, the other fluctuates downward by exactly the same amount. Let the two recorded energies be $X = E_1 + n_1$ and $Y = E_2 + n_2$. Suppose each readout noise is independent of the physical energies and of the other noise. The same-channel variance and the cross covariance then measure different combinations.

$$\operatorname{var}(X) = \operatorname{var}(E_1) + \operatorname{var}(n_1),$$
$$\operatorname{cov}(X, Y) = -\operatorname{var}(E_1).$$

A single histogram that is too broad can mimic a reservoir with different heat capacity. Simultaneous measurements in the two channels allow the exchanged-energy variance to be recovered from the negative covariance. A drifting bath or common electronic pickup violates the assumptions. The measured variance of $X + Y$ is a direct diagnostic for such extra fluctuations when the readout noise levels are known.

Compare with a canonical ensemble at the same mean

In a canonical description each mode has a geometric distribution. Match its mean to $\bar{n} = Q/s$. A single mode then has variance $\bar{n}(1 + \bar{n})$, and a independent modes have variance $a\bar{n}(1 + \bar{n})$. Dividing the exact microcanonical variance by this canonical variance leaves $b/(s + 1)$. At the example point this is $8/13$. Increasing a along with b at fixed fraction does not make a macroscopic subsystem independent of the reservoir; the ratio stays below one. Holding a fixed, taking b large and keeping Q/s fixed does recover the canonical width.

The negative-binomial canonical distribution has support on every nonnegative integer, even those above Q. This is another sign that it describes another experiment. Merely deleting the tail without changing its other weights does not recover the true finite-reservoir conditional probability.

Calculate the distribution without enormous factorials

The ratio of successive unnormalised weights is $(q + a)(Q - q)/[(q + 1)(Q - q + b - 1)]$. Start with a positive arbitrary initial weight, march through the allowed q and normalise once. The graph compares this exact probability to a canonical law at the same mean. Adjust reservoir size and total energy; use the displayed first two moments to distinguish a shift in the average from a change in the width. The plotted canonical values cover the physically accessible count interval, so their displayed partial sum can be less than one.

Open this experiment online

Try the calculation

Reservoir modes 8

Excitation quanta 24

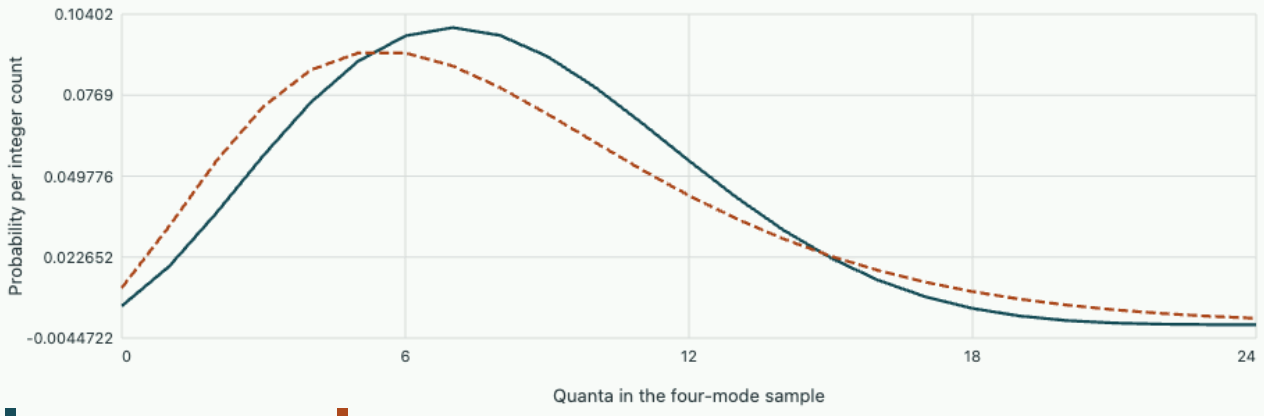

Mean 8 quanta; closed-pair variance 14.769. Canonical variance 24.

Tutorial

If A and B contain identical mode counts, is $Q/2$ necessarily the most likely observed value?

▼ **Work through the solution**

The distribution is symmetric. For odd Q there is no integer value at exactly $Q/2$; two equally likely central counts are possible. With one mode in each solid, every admissible value is equally likely. Symmetry fixes the mean exactly even when it does not identify a unique most probable observation.

For $a = 4, b = 8, Q = 24$ what variance would an unconditioned canonical fit predict?

▼ **Work through the solution**

Here $\bar{n} = 2$. The canonical count variance is $4(2)(3) = 24$. Multiplying by $8/13$ returns 14.769 for the closed pair. Count-to-energy conversion multiplies variances by ε^2.

An experiment finds a noisy sum of the two reported energies. What can one infer before relaxing the energy-conservation model?

▼ **Work through the solution**

In the ideal model $E_A + E_B$ is constant. Check both readouts' measurement noise and energy calibration. Independently fluctuating detector errors add variance in the sum even for a closed

sample. Record a reference configuration or estimate detector noise from repeated measurements of the same state.

The finite-reservoir chapter connects this exact count to the curvature of the entropy.

Counting, correlations and thermal measurements · 3. Recover a free-energy difference from a distribution of work

A two-state system starts in equilibrium, then a control jumps its energy gap. Individual work readings differ although every run uses the same control protocol. Use the elementary model to derive an exponential work identity, compare it with the average work and estimate how many runs it takes to see the events that dominate the estimate.

Specify how work is measured

Set the lower energy to zero and the upper energy initially to ϵ_i. Prepare a canonical ensemble at inverse temperature β, isolate it and suddenly change the gap to ϵ_f without changing the energy eigenvectors. A run that started low stays low and requires zero work. A run that started high stays high and requires work $\epsilon_f - \epsilon_i$, positive if the external control gives energy to the sample. No heat flows during the jump. Returning the system to equilibrium afterwards can exchange heat, but it is a separate part of the protocol.

The initial high-state probability is

$$p_i = e^{-\beta\epsilon_i}/(1 + e^{-\beta\epsilon_i}).$$

Therefore

$$\langle e^{-\beta W}\rangle = (1 - p_i) + p_i e^{-\beta(\epsilon_f - \epsilon_i)} = \frac{1 + e^{-\beta\epsilon_f}}{1 + e^{-\beta\epsilon_i}} = \frac{Z_f}{Z_i}.$$

Taking a negative logarithm and multiplying by $1/\beta$ gives the equilibrium free-energy difference of two ensembles at the same temperature. The actual final population after the jump is still p_i, which generally differs from the final equilibrium occupation. This distinction explains why an equilibrium state function can be recovered from a nonequilibrium final distribution.

Compare the two averages

The mean work is $p_i(\epsilon_f - \epsilon_i)$. Convexity of the exponential gives

$$\langle W\rangle \geq \Delta F.$$

In units of $k_B T$, take an initial gap of 4 and a final gap of 0. The high state occurs with probability 0.01799 and gives work −4 when it occurs. Mean work is about −0.07194, yet $\Delta F = -\ln[2/(1 + e^{-4})] \simeq -0.6750$. Both numbers are negative. The greater number is the one closer to zero. Checking only the sign of average work would not reveal the large difference.

Each rare upper-state event contributes e^4 to the exponential average. Its probability times its weight is nearly one, comparable to the entire contribution of the common lower state. This is why replacing an

average of exponentials by an exponential of an average loses essential information. Repeating the same protocol more slowly with contact to a bath would describe a different path through state space.

Ask what a finite dataset can support

For N independent preparations, the probability of seeing no initial upper state is $(1 - p_i)^N$. With 20 runs in the example this is about 0.696. In such a dataset every measured work is zero and the empirical free-energy difference is zero. This is fully compatible with the model. To have at least a 95 percent chance of seeing even one upper-state event requires $N \geq \ln(0.05) / \ln(1 - p_i)$, or about 166 runs. Seeing only one such event still gives a noisy estimate of its frequency.

The graph holds the final gap at zero. Change the initial gap and ensemble size. It shows the probability of missing the rare branch at each sample size, together with the five-percent level. The reported exact free energy and mean work use all initial states. These predictions can be checked by explicitly enumerating a two-outcome histogram, without drawing a random sample and silently treating its fluctuations as a new law.

Open this experiment online

Try the calculation

Initial gap (kBT) 4

Independent preparations 20

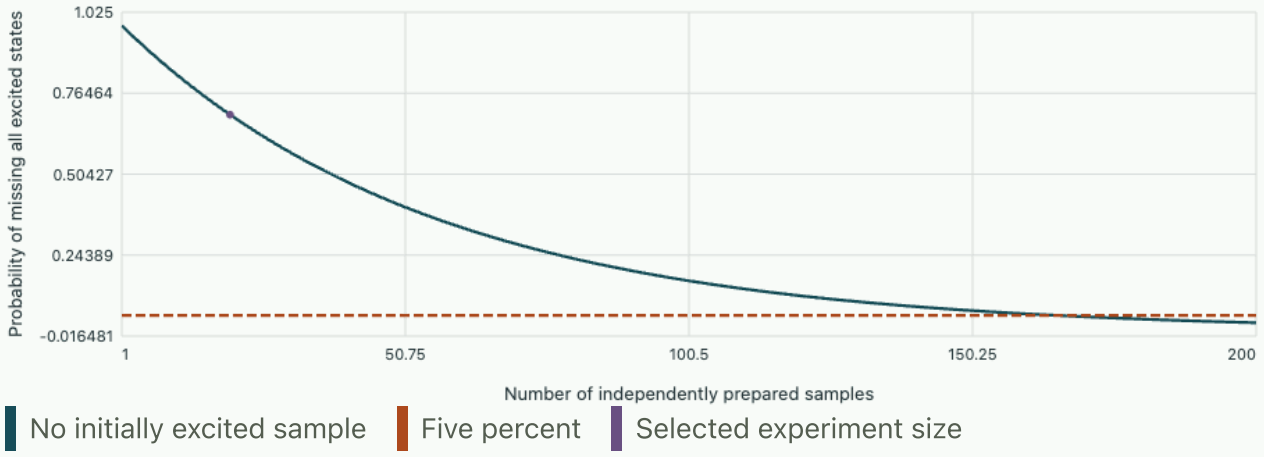

P(no rare work sample) = 0.69559. Mean work / kBT = -0.071945, ΔF/kBT = -0.675.

Tutorial

For an initial gap zero and a final gap $2k_BT$, calculate $\langle W \rangle$ and ΔF.

▼ **Work through the solution**

Initially $p_i = 1/2$. The average work is $k_B T$. The free-energy difference is $k_B T \ln[2/(1+e^{-2})] \simeq 0.5662 k_B T$. Some energy supplied by the control is unavailable as recoverable equilibrium free energy at this temperature.

A preparation procedure always starts the sample in its lower state independent of the specified initial gap. Can the same equality be used by changing only the sample count?

▼ **Work through the solution**

No. The canonical initial probabilities were used algebraically to turn the weighted sum into Z_f/Z_i. With deterministic preparation the exponential average is one in this protocol. Repeated measurements approach that different expectation.

At fixed initial and final gaps, what does reducing temperature do to the number of runs required to see an initial excitation?

▼ **Work through the solution**

The excitation probability is approximately $e^{-\beta\epsilon_i}$ for a large positive reduced gap. Required sample count grows approximately as $1/p_i$ for any fixed detection probability. Controlling preparation and understanding tails become more important as temperature falls.

Connect this model to the partition function and to the role of free energy in a reversible isothermal process.

Counting, correlations and thermal measurements · 4. Why a molecular beam and its source have different measured speeds

A tiny hole lets a dilute gas escape into a vacuum. A rotating chopper admits a short pulse and a downstream detector records arrival times. Derive the source speed density, the effusive speed density and the actual time-of-flight histogram. Each samples the same Maxwellian physics with a different weighting.

State what crosses a surface

For a gas with particle mass m and temperature T, write $a = m/(2k_BT)$. Cartesian velocity components have density proportional to $e^{-a|\mathbf{v}|^2}$. Integrating over a velocity shell gives an equilibrium speed density proportional to $v^2e^{-av^2}$. Sampling gas particles uniformly in a volume measures this density. For particles crossing a plane aperture in unit time, a particle with normal speed $v_z > 0$ can reach the plane from a depth $v_z\,dt$. Multiply by this flux factor and integrate over directions in the outward hemisphere. In polar coordinates about the aperture normal the angular integral of $\cos\theta\,\sin\theta$ is independent of speed and an extra factor v remains.

$$p_{\text{gas}}(v) = \frac{4a^{3/2}}{\sqrt{\pi}} v^2 e^{-av^2},$$

$$p_{\text{beam}}(v) = 2a^2 v^3 e^{-av^2}, \qquad v > 0.$$

This model assumes a hole small compared with the molecular mean free path, negligible collisions between aperture and detector, and acceptance that selects a fixed narrow angular cone with speed-independent efficiency. A nozzle that produces an expanding collisional jet needs a different velocity distribution.

Transform speed into arrival time

Let L be the flight distance and $t = L/v$ the transit time along the accepted direction. Probability conservation for a histogram bin requires

$$p_t(t)\,dt = p_{\text{beam}}(v)|dv|.$$

The derivative $|dv/dt| = L/t^2$ is essential:

$$p_t(t) = 2a^2L^4t^{-5}e^{-aL^2/t^2}.$$

The mode of the equilibrium speeds is $1/\sqrt{a}$, that of beam speeds is $\sqrt{3/(2a)}$, and the mode of arrival times is $L\sqrt{2a/5}$. The speed calculated as L/t_{mode} is $\sqrt{5/(2a)}$; it is neither of the speed modes. Peaks of densities are not invariant under nonlinear changes of variable. Integrated probabilities for corresponding intervals are invariant. Changing the horizontal labels on a histogram while retaining its original bin heights gives the wrong likelihood.

Fit a temperature from counted arrivals

Define thermal speed $v_T = \sqrt{2k_BT/m}$ and time $t_0 = L/v_T$. The dimensionless variable $s = t/t_0$ has density $2s^{-5}e^{-1/s^2}$. An argon atom of mass 39.95 atomic mass units at 300 K has $v_T \simeq 353\,\mathrm{m\,s^{-1}}$. For a flight of 0.50 m the most common arrival time is about 0.895 ms; interpreting its reciprocal as the most probable gas speed would overestimate the temperature by a factor of 2.5. Finite gate width convolves the ideal distribution with the opening-time distribution. Record the gate profile or vary L to separate instrument timing from thermal spread.

Rather than fitting one maximum, consider ideal independent arrival times t_i. The a-dependent part of the log likelihood is $2N\ln a - aL^2\sum_i t_i^{-2}$. Differentiation gives

$$\hat{a} = 2N/[L^2\sum_i t_i^{-2}].$$

Consequently $\hat{T} = mL^2\sum_i t_i^{-2}/(4Nk_B)$. This estimator is just a beam-weighted mean squared speed divided by the coefficient appropriate to an effusive beam. A velocity-dependent detector efficiency would alter that coefficient.

Compare density shapes in the measurement coordinate

Change source temperature and flight length. The graph compares arrival times for particles sampled from a volume with arrivals from the effusive flux, both transformed with their correct Jacobians to equal-width time bins. Compare the most probable times and widths. Diluting an ideal gas at fixed temperature changes the signal rate, while the normalised ideal shapes stay the same. A shift of shape with pressure would suggest collisions or a change in how the source is sampled.

Open this experiment online

Try the calculation

Argon temperature (K) **300**

Flight length (m) **0.5**

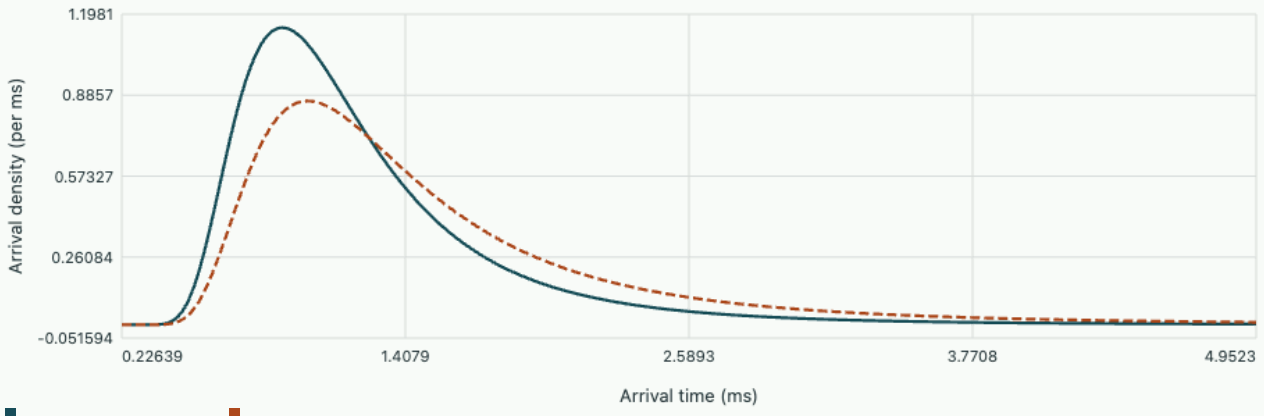

Thermal speed 353.37 m/s. Most probable effusive arrival 0.89488 ms. Interpreting L/t as a gas speed would give 558.73 m/s.

Tutorial

A detector records twice the flight time after doubling the path length while all dimensionless histogram quantiles coincide. Which interpretation is consistent with the data?

▼ **Work through the solution**

Flight is ballistic at unchanged speed distribution. Under this model t_0 scales linearly with L. If a delay were dominated by an electronic timing offset it would not scale this way.

Suppose the detector efficiency is inversely proportional to speed across the relevant band. How does this affect the effusive weighting?

▼ **Work through the solution**

The source flux multiplies the gas density by v; the detector multiplies the result by $1/v$. After normalisation the detected speed density has the same shape as the gas density. The time histogram still requires the Jacobian L/t^2.

Why is averaging the directly transformed speeds L/t_i different from dividing flight length by the mean arrival time?

▼ **Work through the solution**

The reciprocal is nonlinear. The latter is related to a harmonic mean of speeds. A broad distribution enhances the difference, even with perfect timing.

Compare a phase-space sample with a flux measurement in the kinetic theory chapter.

Counting, correlations and thermal measurements · 5. Detect irreversible steady motion from transition counts

A three-state system appears equally likely to occupy each state. Is it at thermal equilibrium? Build two different stationary processes with identical populations, then determine what a time-resolved instrument must count to distinguish them.

Construct a stationary distribution without assuming equilibrium

Arrange states $0, 1, 2$ on a directed ring. From each state the clockwise rate is k_+, the anticlockwise rate is k_-, both positive. For example $\dot{p}_0 = k_+ p_2 + k_- p_1 - (k_+ + k_-) p_0$, with the other equations obtained by cycling the indices. Uniform $p_i = 1/3$ makes every derivative vanish for any values of the two rates. The normalised stationary populations contain no information on their ratio.

Detailed balance would require $p_i k_{i\to j} = p_j k_{j\to i}$ across each individual edge. Here the stationary net clockwise current on each edge is $J = (k_+ - k_-)/3$. Summing over all edges gives mean signed jump count per time $k_+ - k_-$. One complete rotation contains three net clockwise jumps. A stationary histogram does not determine whether those currents exist because a divergence-free current leaves each state's population unchanged.

Read the direction of time from a trajectory

At stationarity compare the probability of one clockwise jump across an edge with its reverse. Their logarithmic ratio is $\ln(k_+/k_-)$. Multiplying this increment by k_B and the net flux, then summing the three oriented edges, gives the irreversible entropy-production rate for the observed Markov model.

$$\frac{\dot{S}_{\mathrm{prod}}}{k_B} = (k_+ - k_-)\ln(k_+/k_-) \geq 0.$$

Both factors change sign together if the circulation is reversed. There is no $1/2$ here because each undirected edge was included once with its chosen orientation. Counting all ordered pairs instead would require compensating for double counting. An interpretation in terms of actual heat delivered to a reservoir requires thermodynamic information such as local detailed balance. The transition record alone measures time-reversal asymmetry of this model.

At $k_+ = 6\,\mathrm{s}^{-1}$, $k_- = 2\,\mathrm{s}^{-1}$, each edge carries $4/3$ net jumps per second. The entropy-production rate is

$$4\ln 3 = 4.394\,k_B/\mathrm{s}.$$

An unbiased process with $k_+ = k_- = 4\,\mathrm{s}^{-1}$ has the same uniform occupations and the same total jump rate of eight per second, yet zero net current and zero entropy production. Counting only total events is insufficient too.

Resolve what a slow camera misses

Because every state has the same outgoing rates, the number of clockwise and anticlockwise jumps in duration t can be constructed from independent Poisson processes of means k_+t, k_-t. Their difference has mean $(k_+ - k_-)t$ and variance $(k_+ + k_-)t$. At the example rates a 10 s record has expected signed count 40 and standard deviation $\sqrt{80} \simeq 8.94$. If the electronics wraps angles modulo three, a sequence of three consecutive clockwise jumps between frames looks like zero displacement. High temporal resolution and a sufficiently long record solve different experimental problems.

Parameterise the rates by total activity $r = k_+ + k_-$ and bias $b = (k_+ - k_-)/r$. Then $\dot{S}_{\mathrm{prod}}/k_B = rb\ln[(1+b)/(1-b)]$. Close to equilibrium this grows as $2rb^2$, while current grows linearly in b. A weak current may be detectable in a signed record when estimates of an entropy-production rate remain noisy. The logarithmic estimate from raw counts is particularly unstable if a short record has no detected backward jumps.

Try equal stationary histograms with different histories

Change the bias and total transition rate. The graph gives the expected unwrapped signed count and one-standard-deviation band. The predicted occupation of each state stays fixed at one third. Reverse the bias and compare the mean drift, fluctuation width and entropy production. This experiment plots analytical ensemble statistics, so repeated changes of the controls give reproducible values.

Open this experiment online

Try the calculation

Clockwise bias **0.5**

Total jump rate (s⁻¹) **8**

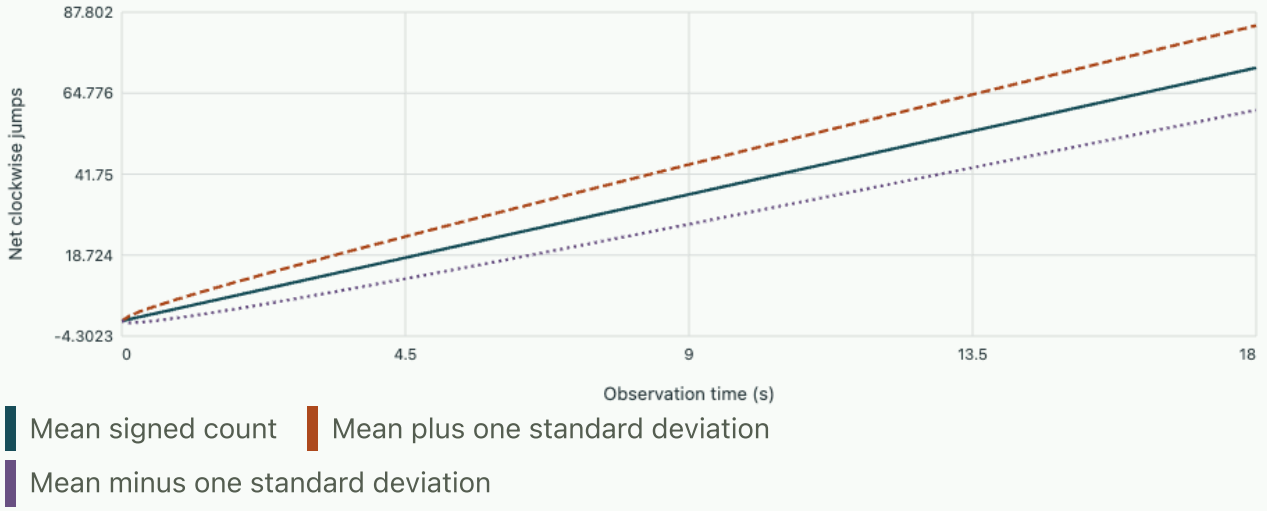

Every stationary population is 1/3. Net clockwise edge current 1.3333 s⁻¹. Entropy production 4.3944 kB/s.

Tutorial

A stationary 100 s record contains 480 clockwise and 320 anticlockwise jumps across all edges. Estimate the two rates and the entropy production under the stated homogeneous model.

▼ **Work through the solution**

The rate estimates are 4.8 and 3.2 per second. The inferred rate divided by k_B is $1.6\ln(1.5) = 0.649\,\mathrm{s}^{-1}$. Estimating separate edge rates for an inhomogeneous ring would also require time spent in each source state.

Can an observed two-state continuous-time chain sustain a nonzero stationary edge current with one reversible edge?

▼ **Work through the solution**

Stationarity for either state directly sets the net edge current to zero. Hidden states or multiple unresolved transition channels can still dissipate. Absence of a current in a coarse record does not determine every microscopic flux.

How long must the idealised process with rates 6 and 2 per second be observed for the expected drift to be five times its standard deviation?

▼ **Work through the solution**

Solve $4t = 5\sqrt{8t}$, giving $t = 12.5$ s. This signal-to-noise criterion compares an ensemble mean to its spread and does not guarantee a particular outcome in one record.

Relate this inference to equilibrium reversibility in the master-equation chapter.

Counting, correlations and thermal measurements · 6. Why macroscopic measurements become reproducible

The same collection of microscopic switches can teach us two different facts about a macroscopic sample. Its mean becomes predictable, and the fluctuations about that mean acquire a particular shape. Derive both statements, then find an experiment in which only the first prediction is reliable at the available sample size.

Start with an experiment that has only two outcomes

A detector records whether one two-level site is excited. Write $X_i = 1$ for an excitation at site i and $X_i = 0$ for its ground state. For now take the sites to be equivalent and independent, with fixed excitation probability $0 < p < 1$. Sites with identical level spacings in a heat bath can share a bias different from $p = 1/2$. For a nondegenerate gap ϵ in thermal equilibrium, the ratio $p/(1-p) = \exp(-\epsilon/k_B T)$ fixes the bias. Even when every site has the same preparation, a single record can contain any count from zero to N.

To obtain exactly k excitations, first specify one pattern with k occupied sites. Independence gives its probability $p^k(1-p)^{N-k}$. There are $\binom{N}{k}$ ways to select the occupied positions. Thus

$$\mathbb{P}(K = k) = \binom{N}{k} p^k (1-p)^{N-k},$$

$$K = \sum_{i=1}^{N} X_i.$$

Distinguish this probability from the multiplicity. With four sites and $p = 0.2$, the probability of exactly two excitations is

$$6(0.2)^2(0.8)^2 = 0.1536.$$

Six is a count of patterns. The factor multiplying it carries information about the physical preparation. A change of temperature changes that factor even though the available patterns stay the same.

Derive a law of large numbers from one inequality

Linearity of expectation gives

$$\mathbb{E}K = Np,$$

with no requirement of independence. For the variance, expand the square of the deviation as the sum of all pair products. Distinct sites give zero covariance under the present assumptions. Each diagonal term contributes $p(1-p)$, because $X_i^2 = X_i$. Consequently the measured excited fraction has variance $p(1-p)/N$. The width of the count grows with N. The width of the fraction shrinks.

Choose an absolute tolerance $a > 0$. On the event that the measured fraction differs from p by at least a, its squared deviation is at least a^2. Keeping only that event in the expectation of the nonnegative square gives

$$\mathbb{P}(|K/N - p| \geq a) \leq \frac{\mathbb{E}(K/N - p)^2}{a^2} = \frac{p(1-p)}{Na^2}.$$

At fixed a the upper bound tends to zero. This proves convergence in probability for the measured fraction. It gives a conservative guarantee even if little else is known about its distribution. At $p = 1/2$, $a = 0.05$, $N = 1000$, the bound is 0.1. It is a limit on a probability, and substituting it as the probability of failure would discard information we actually possess about these sites.

Find the shape hidden by the concentration bound

For the count itself, magnify deviations on their natural scale. Set $\sigma = \sqrt{p(1-p)}$, $Y_i = (X_i - p)/\sigma$ and $Z_N = N^{-1/2} \sum Y_i$. Now $\mathbb{E}Y_i = 0$ and $\mathbb{E}Y_i^2 = 1$. The characteristic function $\phi(t) = \mathbb{E}e^{itY_i}$ encodes moments in its Taylor series. In this bounded two-state example its first three terms yield

$$\phi(t) = 1 - t^2/2 + O(|t|^3).$$

The missing linear term follows from centring. At $p = 0$ or $p = 1$ the count is already fixed and its zero standard deviation makes this rescaling undefined.

Independence turns the characteristic function of the sum into a product. Each argument gets divided by $\sqrt{N}$, while each of the N factors is identical. Taking logarithms turns this product into

$$\ln \mathbb{E}e^{itZ_N} = N \ln \phi(t/\sqrt{N}) = -t^2/2 + O(|t|^3/\sqrt{N}) \quad \text{for fixed } t.$$

The limit $e^{-t^2/2}$ is the characteristic function of the unit normal density. The continuity theorem for characteristic functions therefore gives convergence in distribution to a standard normal. This statement concerns the probabilities of scaled intervals. A normal random variable assigns zero probability to any single value, whereas a finite count assigns positive probability to integer bins. To compare them in an experiment, integrate the normal density over each bin from $k - 1/2$ to $k + 1/2$ in the original count coordinate. Probability outside the physical range is a visible defect of the approximation for small samples.

The first correction predicts which experiment will look asymmetric

For one Bernoulli site the centred third moment equals $p(1-p)(1-2p)$. In the sum, terms with any singly occurring independent centred factor vanish. Thus the third centred moment of the count grows as N, while its standard deviation cubed grows as $N^{3/2}$. Their ratio is

$$\gamma_1 = \frac{1-2p}{\sqrt{Np(1-p)}}.$$

A rare excitation produces a longer upper tail. For $N = 20, p = 0.05$, even an entirely empty detector has probability $0.95^{20} \simeq 0.3585$. Here a symmetric curve extends into a physically impossible negative count and may miss the leftmost bin. At $N = 20, p = 1/2$ the skewness vanishes exactly by reflection symmetry. Finite support and discreteness still matter, particularly in the far tails. Increasing N while reducing the excitation probability to $p_N = 1/N$ describes a different limit. The mean count stays at one and the empty-bin probability $(1 - 1/N)^N$ tends to e^{-1}. A Poisson count of mean one describes this rare-event limit.

Open this experiment online

Try the calculation

Independent sites N **20**

Excitation probability p **0.2**

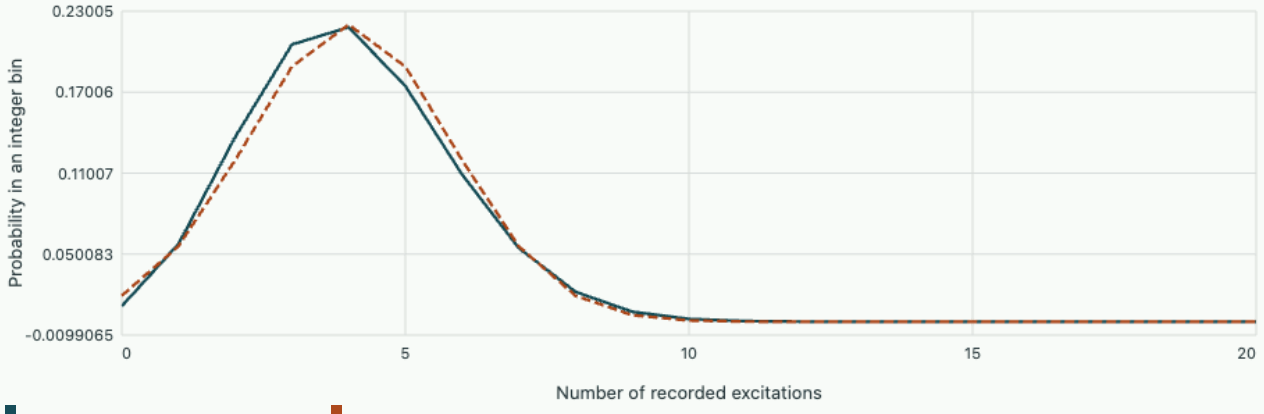

Mean 4, standard deviation 1.7889. Largest absolute bin error 0.017167 at count 2. Skewness 0.33541.

Compare a detector with one expected excitation and a detector with twenty. The displayed normal-bin probabilities use the same mean and variance as the exact count. Identify the bin with the largest absolute discrepancy. A small central discrepancy alone gives little information about a very small tail probability.

Changing the physics can change the scaling

Now imagine observing the same slowly fluctuating site at evenly spaced times. Centring removes its mean without removing memory. If the sequence is stationary with lag-h covariance $C(h)$, then expanding the variance of the sample mean once more gives

$$\operatorname{var}(\bar{X}) = \frac{1}{N}\left[C(0) + 2\sum_{h=1}^{N-1}(1 - h/N)C(h)\right].$$

There are $N - h$ pairs at separation h in either order. Positive correlations inflate the variance compared with independent observations. For $C(h) = C(0)\rho^h$, $0 \le \rho < 1$, and a record long compared with the correlation time, the bracket tends to $C(0)(1 + \rho)/(1 - \rho)$. With $\rho = 0.9$, roughly nineteen records contain the information of one independent record for estimating the mean. Dividing the same observation interval into more samples need not buy nineteen times the information.

A different failure comes from the single-site law itself. A standard Cauchy variable has characteristic function $e^{-|t|}$ and has no finite variance or mean. The characteristic function of the arithmetic average of N independent Cauchy readings is $(e^{-|t|/N})^N = e^{-|t|}$. The average has exactly the original spread at every sample size. Identify the physical observable and its tail behaviour before assigning the familiar square-root error bar.

Tutorial on counting, fluctuations and experimental design

1. Size an equilibrium sample. For sites prepared with $p = 0.2$, the measured excited fraction should have standard deviation no greater than 0.01. An instrument can instead measure repeated copies with the same gap. How many independent readings suffice, and how does this differ from a guarantee that every report is within 0.01?

▼ **Calculate the required number**

The variance criterion is $0.2(1 - 0.2)/N \le 10^{-4}$, giving $N \ge 1600$. It supplies one standard deviation. It leaves a finite probability of larger deviations. Chebyshev's inequality alone would require 32,000 readings to guarantee that the probability of a deviation of at least 0.01 is at most 0.05. If 0.2 is only a pilot estimate, a guarantee valid for any preparation uses $p(1 - p) \le 1/4$, giving 2,500 and 50,000 readings for the two requirements.

---

2. Test a tail. At $p = 0.05, N = 20$, approximate the probability of no excitations in two ways. First calculate it exactly. Then use the normal probability contained between –0.5 and 0.5.

▼ **Separate the empty bin from a negative count**

The exact value is 0.35849. The mean is 1 and the standard deviation is

$$\sqrt{0.95} = 0.97468.$$

The standardised boundaries are about –1.539 and –0.513, yielding approximately 0.242. The same discrete event can be written as $K \le 0$. Applying a one-sided continuity correction to that cumulative probability gives $\Phi((0.5 - 1)/0.97468) \simeq 0.3040$. The difference of 0.0619 is the Gaussian

probability below −0.5. The exact detector has no bins there. The plot integrates over individual finite bins, so its empty-bin value is 0.2421.

---

3. Identify shared noise. One hundred identical thermometers each have independent zero-mean noise of standard deviation 0.20 K. All share a random calibration offset with zero mean and standard deviation 0.10 K. What is the standard deviation of the mean reading under repeated calibration?

▼ **Keep the common contribution**

Write $X_i = T + B + \eta_i$, with B independent of the individual noises. The average contains B exactly once and the mean of all η_i. Its variance is $0.10^2 + 0.20^2/100 = 0.0104\,\mathrm{K}^2$, so the standard deviation is about 0.102 K. Collecting more readings under the same calibration can only shrink the second term.

---

4. Keep the units when converting a fluctuation. Fifty independent sites have an energy gap of 3 meV. At one temperature, one quarter of the sites are typically excited. Find the mean excitation energy of the whole sample and its standard deviation.

▼ **Convert the count to energy**

The energy relative to the all-ground state is $E = \epsilon K$. Hence

$$\langle E \rangle = 3(50)(0.25) = 37.5\,\mathrm{meV},$$

$$\sigma_E = 3\sqrt{50(0.25)(0.75)} = 9.19\,\mathrm{meV}.$$

The standard deviation of the sample energy per site is 0.184 meV. Which of these the detector reports determines whether the measured width grows or shrinks as additional sites are included.

Fließbach places mathematical statistics before the construction of the statistical ensembles. The publisher's contents for *Statistische Physik* identify this starting point. Continue to the finite-reservoir experiment to see how conservation couples the recorded counts.

Counting, correlations and thermal measurements · 7. Weigh a colloid by counting where it spends its time

Brownian motion spreads a dilute suspension through a cell whilst buoyancy biases its vertical position. Derive the equilibrium concentration, predict counts in finite camera bins and estimate a buoyant mass without measuring how quickly any individual sphere falls.

Balance drift against diffusion

Measure z upwards from the bottom of a closed cell. A sphere of volume V displaces fluid of density ρ_f. The downward gravitational force on the particle exceeds the upward buoyancy force by $m_b g$, where $m_b = V(\rho_p - \rho_f)$. Signed buoyant mass is what a concentration profile can reveal. Let μ denote mobility, the ratio of speed to applied force. For dilute, independently moving particles at uniform temperature, the total vertical number flux is

$$J = -D\frac{dn}{dz} - \mu m_b g n.$$

At steady state the flux is constant, and impermeable cell ends set that constant to zero. The Einstein relation $D = \mu k_B T$ eliminates the mobility. Integrating $d\ln n/dz = -a$ with $a = m_b g/(k_B T)$ gives

$$n(z) = n(0)e^{-az}.$$

The signed slope changes sign for buoyant particles and vanishes at exact density matching. Viscosity affects the time to reach equilibrium.

Suppose the camera counts centres in equal-width horizontal slices of a column with uniform cross-section. A bin centred at z_i has expected count proportional to $e^{-az_i}\int_{-h/2}^{h/2} e^{-au}du$. The integral is common to all bins, hence normalising the exponential values at their centres yields the exact bin probabilities for this model. Wall exclusion sets the interval accessible to centres. Define the specified cell height as that accessible interval.

Infer a mass with its sign

Use a radius of 0.50 µm, particle excess density 50 kg/m³ and temperature 293 K. The volume is 5.236×10^{-19} m³, the buoyant mass is 2.618×10^{-17} kg and $k_B T/(m_b g) = 15.75$ µm. Counts in two equal slices separated by 10 µm should have lower-to-upper ratio $e^{10/15.75} \simeq 1.89$. Doubling the radius at fixed density ratio multiplies the mass by eight and divides the gravitational length by eight.

Sampling is important when using the upper part of a steep profile. Suppose bins separated by 10 µm contain 120 and 60 approximately independent particle observations. The estimated slope is $\ln(2)/10 = 0.0693$ µm⁻¹. For independent Poisson counts the standard error from propagating the two count variances is $\sqrt{1/120 + 1/60}/10 = 0.0158$ µm⁻¹. Repeatedly imaging the same particle at closely spaced times introduces correlations. Counting all those frames as independent would make the quoted sampling precision too optimistic.

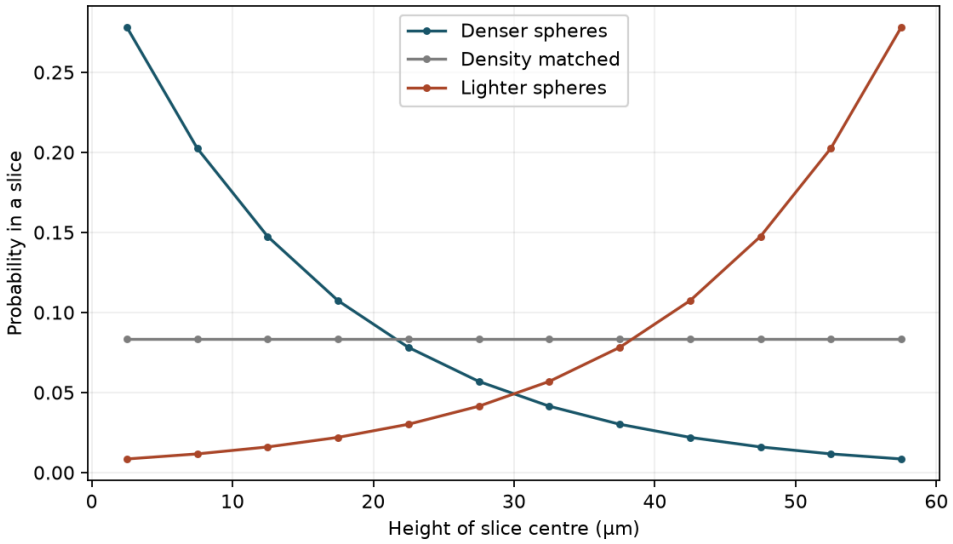

The total probability in the cell is one for each curve. Matching the fluid density removes the gravity-induced concentration gradient.

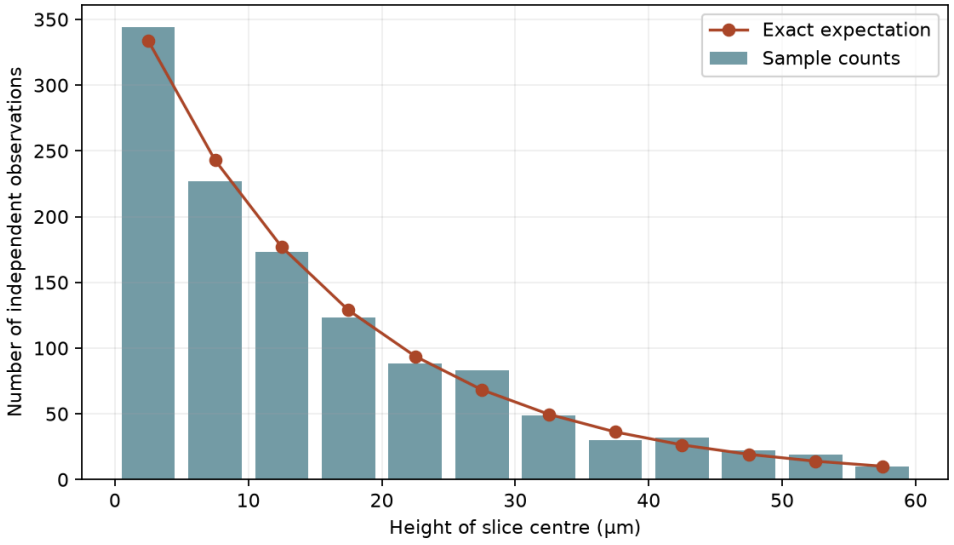

Reduce the number of independent particle observations and change the random seed. Noise is particularly large relative to the expected number in sparsely occupied slices.

Exercise · determine the size of a density-matched particle

A uniform profile persists in a well-equilibrated, isothermal cell. Can gravity alone distinguish two different radii if the particles have the same density as the fluid?

▼ **Work through the solution**

At exact density matching, $m_b = V(\rho_p - \rho_f) = 0$ for either radius, so both profiles are uniform. Their diffusion times can still differ because hydrodynamic mobility depends on size. A time-resolved measurement carries information absent from the equilibrium profile.

Exercise · reverse the measured gradient

Two equally wide slices separated vertically by 12 µm contain 80 and 160 observations, in order of increasing height. Estimate m_b at 293 K.

▼ **Work through the solution**

The signed slope is $\ln(80/160)/(12\,\mu\mathrm{m}) = -0.05776$ µm⁻¹. Converting to metres and multiplying by $k_B T/g$ gives $m_b \simeq -2.38 \times 10^{-17}$ kg. The sign indicates that the displaced fluid is heavier than the particle.

Try it · Sedimentation

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

Open this live calculation

Try the calculation

Particle radius (µm)	0.5
Particle minus fluid density (kg/m³)	50
Independent observations	1200

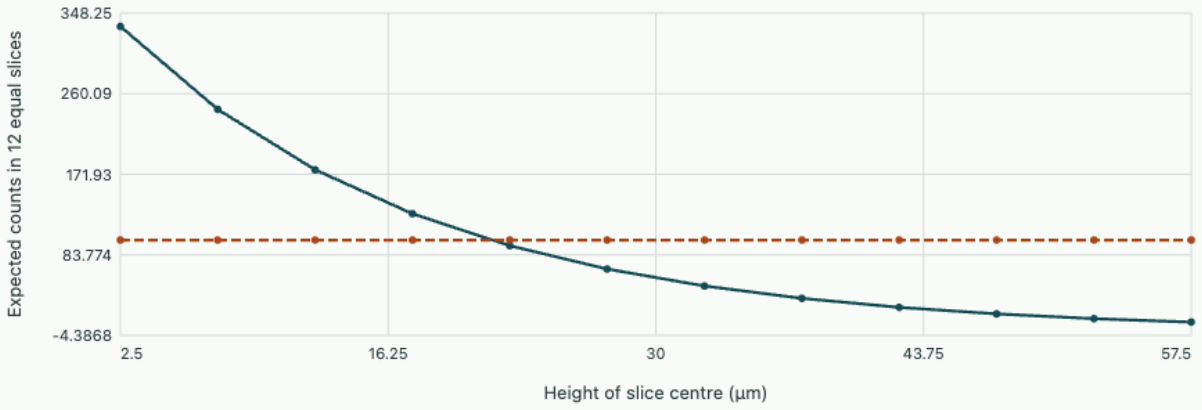

Temperature 293 K; signed gravitational slope 0.063466 µm⁻¹. Counts sum to 1200. The notebook adds seeded multinomial sampling to these expectations.

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

Open editable notebook

Download .ipynb