Computational quantum experiments

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

Contents

1. **Measure tunnelling from a four-site spectrum**
   Four localised orbitals form a minimal double well. Raise the inner two site energies and ask why two states concentrated in the outer wells still acquire different energies. Measure the splitting in the spectrum, predict an oscillation in time and reproduce the levels independently with OneSimulation.

2. **Which part of a Bell correlation survives imperfect state preparation?**
   A source emits two spin-half systems in the singlet state. Two observers select measurement directions separately and record outcomes $+1$ or -1. Construct a Bell–CHSH statistic from four measurable correlations and determine how state preparation, angular alignment and finite sample size affect the prediction.

3. **Recover a transition frequency from two separated pulses**
   A two-level system is driven by two short pulses with a free evolution interval between them. Determine why a population measurement carries phase information. Use a scan at several delays to separate a constant detuning, intrinsic coherence loss and trial-to-trial frequency variation.

4. **When three perfectly legal spin measurements imply an impossible state**
   Prepare nominally identical spin-half systems and measure along three orthogonal axes, using a different subensemble for each axis. Reconstruct a density matrix, diagnose an unphysical linear estimate and see why the choice of estimator matters in finite data.

5. **Read wave-packet overlap from a coincidence dip**
   Send one photon into each of two input ports of a balanced lossless beam splitter. Two alternatives lead to one click at each output. Derive the coincidence probability when their wave packets overlap only partially, then determine how a timing scan separates pulse duration from distinguishability in another degree of freedom.

6. **Can an atom leave its initial state and still pass the final test?**
   A driven two-level atom undergoes repeated ideal measurements. Compare two laboratory records, one retaining only runs that passed every test and another retaining every run. The

distinction is visible even when every measurement is made with the same apparatus.

Computational quantum experiments · 1. Measure tunnelling from a four-site spectrum

Four localised orbitals form a minimal double well. Raise the inner two site energies and ask why two states concentrated in the outer wells still acquire different energies. Measure the splitting in the spectrum, predict an oscillation in time and reproduce the levels independently with OneSimulation.

Construct an operator before drawing an energy diagram

Label the sites $1, 2, 3, 4$ in order. The end orbitals have on-site energy zero, the middle two have energy $V > 0$ and every neighbouring pair has real hopping matrix element $-t$ with $t > 0$. The four occupation states form an orthonormal basis for one particle. Equal diagonal energies of the ends do not make their symmetric and antisymmetric combinations degenerate when there is a virtual path through the interior. There are no direct links between sites one and three or one and four. Keeping track of every allowed step will fix the power of t in the final answer.

$$H = \begin{pmatrix} 0 & -t & 0 & 0 \\ -t & V & -t & 0 \\ 0 & -t & V & -t \\ 0 & 0 & -t & 0 \end{pmatrix}.$$

Reflection maps $1 \leftrightarrow 4$ and $2 \leftrightarrow 3$, and commutes with H. In each parity sector use the normalised outer and inner basis vectors $(|1\rangle \pm |4\rangle)/\sqrt{2}$ and $(|2\rangle \pm |3\rangle)/\sqrt{2}$. The diagonal outer element vanishes in both sectors; the hopping inside the middle pair shifts its even and odd energies to $D_e = V - t$ and $D_o = V + t$. The outer–inner matrix element remains $-t$. This reduces the four-dimensional problem to two symmetric two-by-two matrices without neglecting any state.

Eliminate the barrier amplitudes and estimate the splitting

In a sector with interior energy D, the secular equation is

$$E(E - D) = t^2.$$

The lower root is $E_-(D) = [D - \sqrt{D^2 + 4t^2}]/2$, or the numerically better $-2t^2/(D + \sqrt{D^2 + 4t^2})$. For $D \gg t$ the shift is approximately $-t^2/D$. It is more negative for the even state because $D_e < D_o$. Subtract the two shifts before approximating their difference.

$$\Delta E = E_-(D_o) - E_-(D_e) \simeq t^2 \left(\frac{1}{V - t} - \frac{1}{V + t} \right) \simeq \frac{2t^3}{V^2}.$$

A single outward excursion produces an order-t^2/V shift in both end states. Only a three-hop path connects the different ends, so their splitting is smaller. Measuring an overall displacement of two spectral lines and measuring the distance between them test different physical mechanisms.

Predict a time trace and define what is measured

Denote the two low-energy eigenstates by $|e\rangle, |o\rangle$. Prepare $(|e\rangle + |o\rangle)/\sqrt{2}$. In the corresponding opposite superposition $(|e\rangle - |o\rangle)/\sqrt{2}$, the occupation at time τ equals $\sin^2(\Delta E\,\tau/2\hbar)$. Its first maximum is at $\pi\hbar/\Delta E$. These superpositions include interior weight, especially at low barrier. A detector on site four alone measures a different projector. Preparing bare site one excites both high and low doublets and can generate additional fast oscillations.

For $V/t = 4$, the two lower eigenvalues in units of t are $(3 - \sqrt{13})/2$ and $(5 - \sqrt{29})/2$, about −0.30278 and −0.19258. Their separation is about $0.11019t$. The high-barrier approximation gives $0.125t$. With $t = 1\,\mathrm{meV}$, $\hbar = 0.65821\,\mathrm{meV\,ps}$ and the first complete transfer in the dressed basis takes about 18.8 ps. If a measurement integrates over many such oscillations, it may return one half despite perfectly coherent evolution.

Compare spectra from three independent constructions

Use the parameter controls to vary the middle barrier and hopping. The plot shows the exact splitting and its leading virtual-hopping prediction over a range of V/t. The printed result gives an actual transfer time. The downloadable OneSimulation model builds the original four-site Hamiltonian and solves its eigenvalue problem by contour integration. Compare all four eigenvalues with the two parity blocks before interpreting their small difference. A wrong sign on one edge of this open chain can be removed by a local change of basis and may leave its spectrum unchanged, so also check any phase-sensitive observable used for preparation.

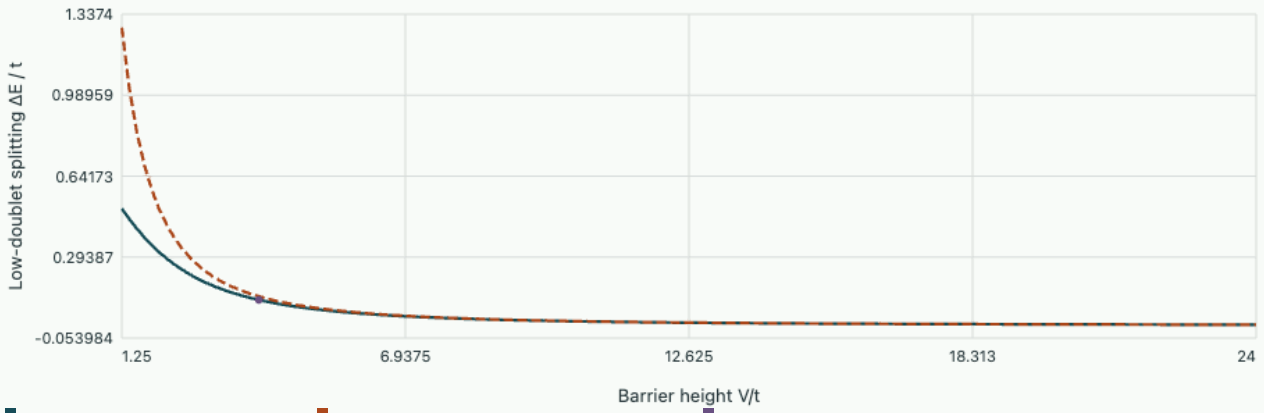
Keep the hopping fixed at 1 meV and compare two barriers. Start with a 4 meV barrier and raise it to 8 meV. Record both splittings and estimate the change in the first transfer time before reading off the prediction.

Open this experiment online

Try the calculation

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

Hopping t (meV) **1**



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

Download the OneSimulation model

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

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

Tutorial

Hold V fixed and halve t in the high-barrier limit. How do the common shift, splitting and transfer time change?

▼ Work through the solution

The leading common shift scales as t^2/V and falls by four. The splitting scales as t^3/V^2 and falls by eight. The transfer time grows eightfold. Halving both V and t instead keeps their ratio fixed and halves every eigenvalue.

Explain why a small static asymmetry between the end wells can suppress almost complete transfer while increasing their measured frequency splitting.

▼ Work through the solution

In a two-state reduction, detuning δ and coupling J give separation $\sqrt{\delta^2 + 4J^2}$. Starting in one localised state, the maximum transfer probability is $4J^2/(\delta^2 + 4J^2)$. Both the frequency and the contrast are needed to infer the coupling.

How could one check that a contour solver has found both members of a closely spaced doublet rather than printing one twice?

▼ Work through the solution

Check the rank of the independent projected vectors and the residual norm for each vector. Then exploit reflection parity or solve in each symmetry sector independently. Repeating a low residual for two nearly parallel vectors does not establish a two-dimensional eigenspace.

The continuous barrier calculation provides a different model in which barrier width eventually produces an exponential dependence.

Computational quantum experiments · 2. Which part of a Bell correlation survives imperfect state preparation?

A source emits two spin-half systems in the singlet state. Two observers select measurement directions separately and record outcomes $+1$ or -1. Construct a Bell–CHSH statistic from four measurable correlations and determine how state preparation, angular alignment and finite sample size affect the prediction.

Relate a counted coincidence to an operator

Denote unit measurement directions by $\mathbf{a}, \mathbf{b}$. Outcome projectors are $(1 + s\,\boldsymbol{\sigma}\cdot\mathbf{a})/2$ and $(1 + t\,\boldsymbol{\sigma}\cdot\mathbf{b})/2$, for $s, t = \pm 1$. The singlet has zero individual spin expectation and tensor $\langle \sigma_i \otimes \sigma_j \rangle = -\delta_{ij}$. The diagonal entries follow because measurements along any common axis are exactly opposite and the state is rotationally invariant. Its joint probability is

$$p(s,t|\mathbf{a},\mathbf{b}) = \tfrac{1}{4}(1 - st\,\mathbf{a}\cdot\mathbf{b}),$$

$$E(\mathbf{a},\mathbf{b}) = -\mathbf{a}\cdot\mathbf{b}.$$

Summing over either outcome leaves one half for the other observer, independent of the remote angle. Covariance can change while the individual counting rates remain constant. At an angle of 60 degrees the probabilities of the two equal-outcome combinations sum to one quarter, and of opposite outcomes to three quarters. Their signed difference gives –0.5. The lab models spin-axis angles; linear polarisation analysers give a doubled angle in their effective two-state representation.

Derive a bound and choose a set of directions

If a trial had four predetermined local binary responses A, A', B, B', its combination $AB + AB' + A'B - A'B' = A(B + B') + A'(B - B')$ would have magnitude two, since exactly one of the parentheses vanishes. Averaging over a common distribution of such assignments preserves the bound. The quantum prediction combines expectation values measured on different subensembles with independently chosen setting pairs. Use coplanar angles $a = 0,\ a' = \pi/2,\ b = \theta,\ b' = -\theta$.

$$S = E(a,b) + E(a,b') + E(a',b) - E(a',b') = -2(\cos\theta + \sin\theta).$$

At $\theta = 45^\circ$ its magnitude is $2\sqrt{2}$. The derivative vanishes there, making small symmetric angular errors second order. If a fraction $1 - v$ of pairs is completely unpolarised, the density operator is $v|\Psi^-\rangle\langle\Psi^-| + (1 - v)I/4$. Every correlator gains a factor v. The best attainable magnitude exceeds two only when $v > 1/\sqrt{2}$. This is a statement about this measurement family and this noise model. A pair source and a scalar contrast alone do not define an experimental protocol.

Assess one set of finite counts

At $v = 0.80$, $\theta = 45^\circ$ the expected magnitude is about 2.263. Each trial product is a binary variable with variance $1 - E^2$. For N independent pairs per setting, the sample correlation variance is

$(1 - E^2)/N$. Assuming independent subensembles, variance of the signed sum is the sum of the four variances. Here every squared correlation is 0.32. With 1000 trials per setting the approximate standard error of the statistic is

$$\sqrt{4(1 - 0.32)/1000} = 0.0522.$$

The excess 0.263 is roughly five standard errors under this idealised independent-trial model. It is not a substitute for a finite-sample hypothesis test that accounts for an actual protocol's setting choices, losses, timing and possible memory.

Change settings and preparation independently

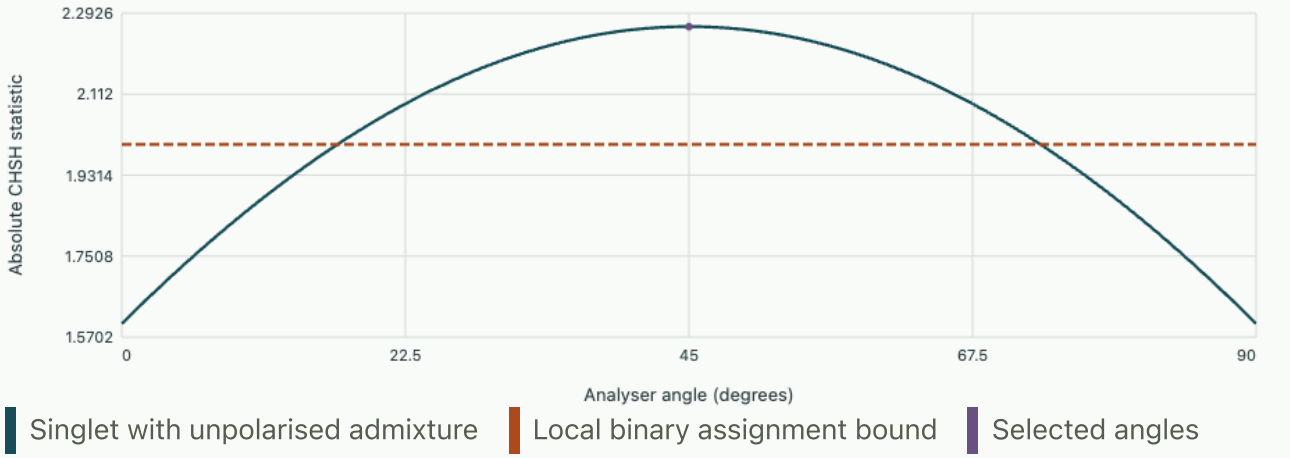
Inspect the predicted absolute statistic and the local assignment bound. Find the angles for which the chosen source can exceed two. Then vary visibility at a fixed angle. Small individual spin polarisation drifts could alter the marginal counts; unequal channel efficiencies could change a coincidence-selected sample. Monitoring all four joint count tables and both single-observer marginal distributions allows more checks than retaining only their final sum.

Open this experiment online

Try the calculation

Preparation visibility **0.8**

Symmetric analyser angle (degrees) **45**



Predicted |S| = 2.2627. With these angles the visibility required to reach two is 0.70711.

Tutorial

At visibility 0.70, can increasing the number of trials eventually make the expected maximum exceed two?

▼ **Work through the solution**

No. The largest mean value is $2\sqrt{2}(0.70) = 1.980$. More trials narrow the sampling distribution around that value. Changing the preparation or measurement family addresses a different question.

What is the predicted fractional loss at an angle δ away from the optimum, and how large is it at five degrees?

▼ **Work through the solution**

Write $\cos\theta + \sin\theta = \sqrt{2}\cos(\theta - \pi/4)$. The fractional loss is $1 - \cos\delta$, about 0.00381 at five degrees. Asymmetric offsets on different analyser axes need not enter the same way.

A paper gives one large value of $|S|$ but discards all pairs for which one detector was silent. Can the ideal bound derived above automatically be applied to these retained counts?

▼ **Work through the solution**

The derivation averaged all setting pairs over a common underlying ensemble. Setting- or state-dependent retention can alter those ensembles. One needs the trial definition, data handling and efficiency assumptions to interpret the reported statistic.

For state preparation and independent measurement on tensor factors, revisit the composite-state chapter.

Computational quantum experiments · 3. Recover a transition frequency from two separated pulses

A two-level system is driven by two short pulses with a free evolution interval between them. Determine why a population measurement carries phase information. Use a scan at several delays to separate a constant detuning, intrinsic coherence loss and trial-to-trial frequency variation.

Follow the amplitudes through both pulses

Use basis $|g\rangle, |e\rangle$. Let the first pulse prepare $(|g\rangle + |e\rangle)/\sqrt{2}$ in the frame of the applied oscillator. A detuning Δ in radians per second accumulates a relative phase $\Delta\tau$ in free time τ. Immediately before the second pulse each state has population one half; measuring population here cannot reveal detuning. Choose a second pulse equivalent to projecting onto $(|g\rangle + e^{-i\phi}|e\rangle)/\sqrt{2}$. Then one output population is

$$P(\tau, \phi) = \tfrac{1}{2}[1 + \cos(\Delta\tau - \phi)].$$

The sign of the phase follows our definitions of detuning and the reference pulse. Shifting the reference phase makes it possible to check that sign. A measurement using only $\phi = 0$ cannot distinguish Δ from $-\Delta$. Coherence is the off-diagonal density-matrix element that the second pulse converts into an observable population difference.

Derive two envelopes with different origins

A homogeneous decay time T_2 multiplies the off-diagonal element by $e^{-\tau/T_2}$. If instead the detuning of each experimental shot is constant during that shot and drawn from a Gaussian distribution of mean $\bar{\Delta}$ and standard deviation σ_Δ, average the complex phase, not its absolute value. Completing the square in the Gaussian integral gives

$$\langle e^{i\Delta\tau} \rangle = e^{i\bar{\Delta}\tau} e^{-\sigma_\Delta^2 \tau^2/2}.$$

With both independent mechanisms the measured contrast is $C(\tau) = \exp[-\tau/T_2 - \sigma_\Delta^2 \tau^2/2]$. One mechanism gives an initial logarithmic slope, the other adds curvature. A single low contrast measured at just one delay cannot establish which one caused it.

Infer parameters from quadratures and a frequency scan

Record the $\phi = 0, \pi/2$ outputs, then form $X = 2P_0 - 1 = C\cos(\bar{\Delta}\tau)$ and $Y = 2P_{\pi/2} - 1 = C\sin(\bar{\Delta}\tau)$. Their length estimates contrast and their angle estimates accumulated phase modulo 2π. Suppose at $0.50\,\mathrm{ms}$ the readings are 0.500 and 0.850. Then $X = 0,\ Y = 0.700$, and the phase is $\pi/2$ modulo 2π. With an independent bound $|\bar{\Delta}/2\pi| < 1\,\mathrm{kHz}$, the implied cyclic detuning is $+500\,\mathrm{Hz}$. Longer delays resolve fine frequency differences at the cost of phase ambiguity and declining contrast. Use a shorter-delay record to unwrap a longer one.

At a contrast of 0.700 the relation $-\ln(0.700) = \tau/T_2 + \sigma_\Delta^2 \tau^2/2$ supplies one equation for two unknowns. Several delays, with independent error bars, are needed to identify the linear and quadratic terms. At very low signal the length $\sqrt{X^2 + Y^2}$ has a positive noise bias. Fitting the two directly measured quadratures can be preferable to fitting the length.

Explore two mechanisms

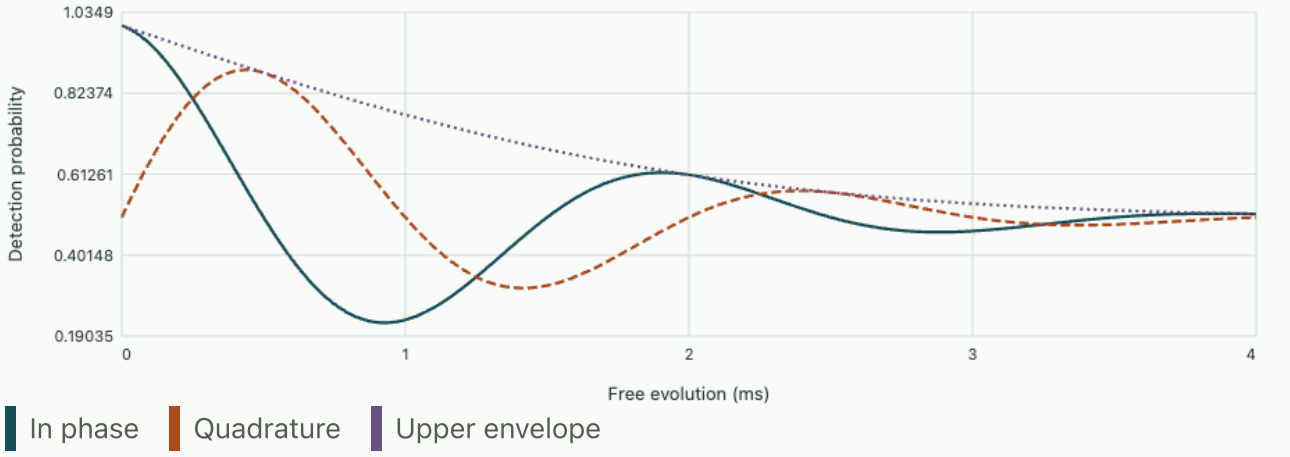
The graph holds the homogeneous T_2 at 2 ms. It lets you set signed cyclic detuning and shot-to-shot cyclic frequency spread. Compare both quadrature populations with the upper contrast envelope. First reverse the sign of the detuning, then increase the spread without changing the mean. Which effect interchanges peaks of one quadrature and which damps both?

Open this experiment online

Try the calculation

Mean cyclic detuning (Hz) **500**

Shot-to-shot rms spread (Hz) **80**



Contrast after 1 ms = 0.53455. The two probabilities at 0.50 ms are 0.5 and 0.87729.

Tutorial

What is the fringe spacing in a scan of cyclic detuning at fixed delay of 0.80 ms?

▼ **Work through the solution**

One full phase cycle needs $\delta f \, \tau = 1$. The spacing is

$$1/0.00080 = 1250\,\mathrm{Hz}.$$

The angular-frequency spacing is 2π times larger.

A fixed delay gives

$$P_0 = P_{\pi/2} = 0.5.$$

Does that show a phase of 45 degrees?

▼ Work through the solution

Both centred quadratures are zero, so the contrast is unresolved and the angle is undefined. A phase of 45 degrees with finite contrast gives equal, nonzero centred quadratures.

Explain how a refocusing pulse halfway through the evolution helps identify quasistatic detuning noise.

▼ Work through the solution

The pulse exchanges the phase-accumulating states, reversing phase winding for a detuning that stays constant through both halves. Random variations during the interval need not cancel. Recovering contrast in the refocused sequence supports a slowly varying contribution to the original envelope.

Compare a continuous drive in the transition chapter with the role of the second pulse here.

Computational quantum experiments · 4. When three perfectly legal spin measurements imply an impossible state

Prepare nominally identical spin-half systems and measure along three orthogonal axes, using a different subensemble for each axis. Reconstruct a density matrix, diagnose an unphysical linear estimate and see why the choice of estimator matters in finite data.

Translate detector counts into three components

Write $\rho = (I + \mathbf{r}\cdot\boldsymbol{\sigma})/2$. For measurement along a unit direction $\mathbf{n}$, the probability of spin up is

$$p_+ = (1 + \mathbf{n}\cdot\mathbf{r})/2.$$

With axes x, y, z, count k_j positive outcomes in N_j trials. Linear inversion gives

$$r_j = 2k_j/N_j - 1.$$

The trace condition is already built into this parametrisation. Positivity remains a separate physical requirement. Using the Pauli product identity, $(\mathbf{r}\cdot\boldsymbol{\sigma})^2 = |\mathbf{r}|^2 I$, the eigenvalues are $(1 \pm |\mathbf{r}|)/2$. They form legitimate probabilities only for $|\mathbf{r}| \le 1$.

Each axis can report a frequency in $[0, 1]$ even when the three reconstructed components have length greater than one. This need not imply a broken detector. The axes are measured on different sets of particles and can fluctuate in mutually inconsistent directions. A state that slowly drifts between the measurements, however, can produce deviations that persist as more particles are counted. Interleaving the settings can test that explanation.

Reconstruct a specific estimate

In 25 shots per axis an experiment obtains positive counts $23, 19, 19$. The reconstructed components are $0.84, 0.52, 0.52$, with norm about 1.1165. The smaller eigenvalue is approximately −0.058. A fourth analyser set opposite the inferred Bloch direction would therefore receive a predicted negative probability from this linear model. That prediction exposes the issue independently of any particular plotting convention.

A simple constrained estimate minimises Euclidean distance from the measured vector to the closed unit ball. If the vector lies inside, the minimiser is the measured vector. Outside, a Lagrange multiplier or elementary projection shows the minimiser is the vector divided by its length. One can now predict valid probabilities for every analyser direction. The result minimises a specified quadratic loss; it need not maximise the likelihood of the original counts.

Why a likelihood fit weights measurements differently

For independent binomial trials, the log likelihood up to constants is $\sum_j [k_j \ln(1+r_j) + (N_j - k_j)\ln(1-r_j)]$. The positive physical state constraint couples the components. Differentiating each term gives a score $[2k_j - N_j - N_j r_j]/(1-r_j^2)$. Thus the cost of moving a component away from its unconstrained estimate depends on both its count and its location. A radial Euclidean projection ignores these different weights. Doubling only N_z while leaving its relative count unchanged preserves the inversion vector but changes the evidence for retaining its z component. A fit using likelihood can account for that evidence.

The shape of the confidence region is a further question. An unconstrained Gaussian ellipse may extend beyond physical states near the surface. A physical point estimate does not by itself specify an uncertainty interval; a measurement report should include both the estimator and the count data or a stated statistical model.

Evaluate a new measurement using the two estimates

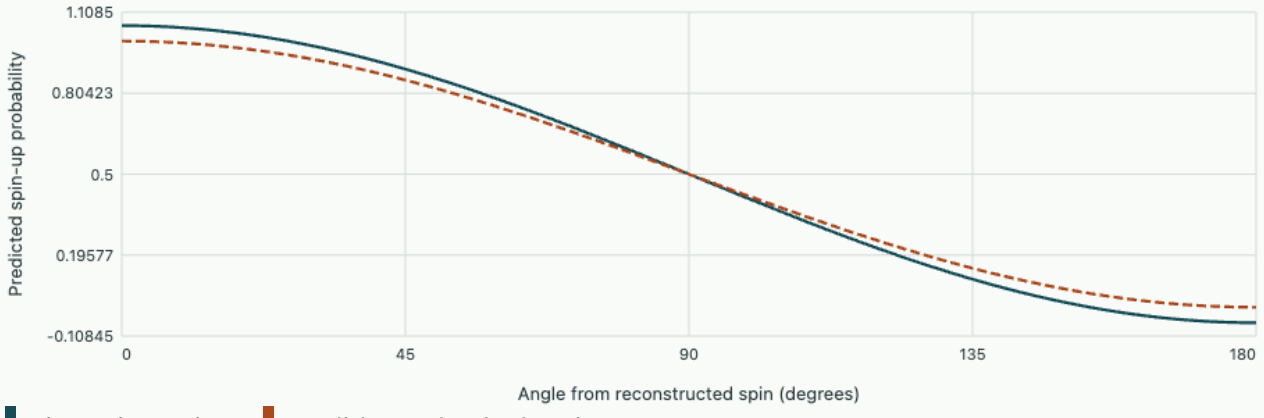
Fix the first count at 23 of 25. Vary the positive counts on y, z, then scan a test analyser in the plane through the estimated direction and an orthogonal axis. The graph plots probability versus angle from the raw vector. Radial projection preserves that direction. Predict which orientation first gives an impossible probability when the estimate moves outside the ball. A truly independent fourth measurement near that direction is more informative about the failure than a repeated plot of the three fitted counts.

Open this experiment online

Try the calculation

Up counts along y (out of 25) **19**

Up counts along z (out of 25) **19**



Reconstructed Bloch radius 1.1164. Eigenvalues 1.0582 and -0.058211.

Tutorial

All three directions are measured on 100 independent systems and the counts are 50, 50 and 50. Does this prove that every specimen was prepared in a mixed state?

▼ Work through the solution

The estimate is $I/2$ as a description of the measured ensemble. Many preparations have the same density operator, including equally likely pure states pointing along opposite directions. The counts do not identify one particular preparation history.

A data analysis repairs a negative eigenvalue by setting it to zero while retaining the original positive eigenvalue unchanged. Which density-matrix condition now fails?

▼ Work through the solution

The trace was initially one. Removing a negative number increases that trace above one. A valid construction must enforce positivity and normalisation together. Projection onto the state ball automatically gives both for a two-level system.

Why is it incorrect to send a single qubit successively through three projective analysers and use the resulting counts in this inversion?

▼ Work through the solution

A projective measurement changes the state entering the next analyser. The likelihood above assumes independent preparations of a common input state. Sequential data instead require the conditional measurement operations and their ordering.

Reconstruct how a partial trace produces a mixed state in the entanglement and measurement chapter.

Computational quantum experiments · 5. Read wave-packet overlap from a coincidence dip

Send one photon into each of two input ports of a balanced lossless beam splitter. Two alternatives lead to one click at each output. Derive the coincidence probability when their wave packets overlap only partially, then determine how a timing scan separates pulse duration from distinguishability in another degree of freedom.

Track the two indistinguishable alternatives

Choose input creation operators $a^\dagger, b^\dagger$ and outputs $c^\dagger, d^\dagger$. A possible symmetric unitary convention maps $a^\dagger$ to $(c^\dagger + id^\dagger)/\sqrt{2}$ and $b^\dagger$ to $(ic^\dagger + d^\dagger)/\sqrt{2}$. For input $a^\dagger b^\dagger|0\rangle$, expand the product. The two one-in-each-output terms have coefficients $1/2$ and $i^2/2 = -1/2$. They cancel for otherwise identical photons. The same-port outputs remain and their squared probabilities add to one after using the $\sqrt{2}$ normalisation of a two-photon mode. The photons do not interact with each other; bosonic amplitudes for indistinguishable output records interfere.

Allow the internal wave packets to differ

Let the photons carry normalised internal states f, g, which may include arrival time, polarisation and spectrum. In the coincidence sector the two possible output states are proportional to $c_f^\dagger d_g^\dagger|0\rangle$ and $c_g^\dagger d_f^\dagger|0\rangle$. Their overlap is $\langle f|g\rangle\langle g|f\rangle = |\langle f|g\rangle|^2$. Squaring the difference of their amplitudes gives

$$P_{\mathrm{coinc}} = \tfrac{1}{2}(1 - |\langle f|g\rangle|^2).$$

Orthogonal internal states give one half, the value obtained by adding alternative paths as independent probabilities. A common overall optical phase of either photon leaves the modulus unchanged. Delay matters because it changes the overlap in time, even without a locked optical phase between the two sources.

Background · Can a polariser restore a missing coincidence dip?

Prepare one horizontally polarised and one vertically polarised photon in otherwise identical temporal and spatial modes. Their internal overlap is zero, giving the distinguishable coincidence probability of one half at a balanced splitter. Put identical 45° polarisers in the two input arms. Each photon passes with probability one half. Conditional on both passing, their polarisation states now agree.

$$P(\text{two pass}) = \frac{1}{4},$$

$$P(\text{coincidence} \mid \text{two pass, zero delay}) = 0.$$

At large delay the conditional coincidence probability returns to one half, so the coincidence probability per incident pair after the polarisers is one eighth. The dip can regain ideal visibility while the available event rate falls. Reporting the transmission or raw count rate along with visibility distinguishes filtering from an improvement in the source brightness.

Compute a dip with a specified width convention

Take a transform-limited Gaussian amplitude proportional to $e^{-t^2/(4\sigma^2)}$. Its detection-time probability has standard deviation σ. Delay one photon by τ. Completing the square in the overlap integral gives amplitude overlap $e^{-\tau^2/(8\sigma^2)}$ and squared overlap $e^{-\tau^2/(4\sigma^2)}$. If both photons also have linear polarisations separated by α, multiply the squared overlap by $\cos^2\alpha$. Thus the dip has full width at half depth $4\sigma\sqrt{\ln 2}$, independent of α as long as a dip remains measurable.

For $\sigma = 80$ fs and $\alpha = 30^\circ$, the coincidence probability is 0.125 at zero delay and approaches 0.5 at long delay. The half-depth level is 0.3125 at $|\tau| \simeq 133$ fs, giving full width about 266 fs. A width measured at probability 0.25 instead would not in general represent half depth. Quote the baseline, minimum and width convention if comparing different measurements.

Decide what a coincidence counter can identify

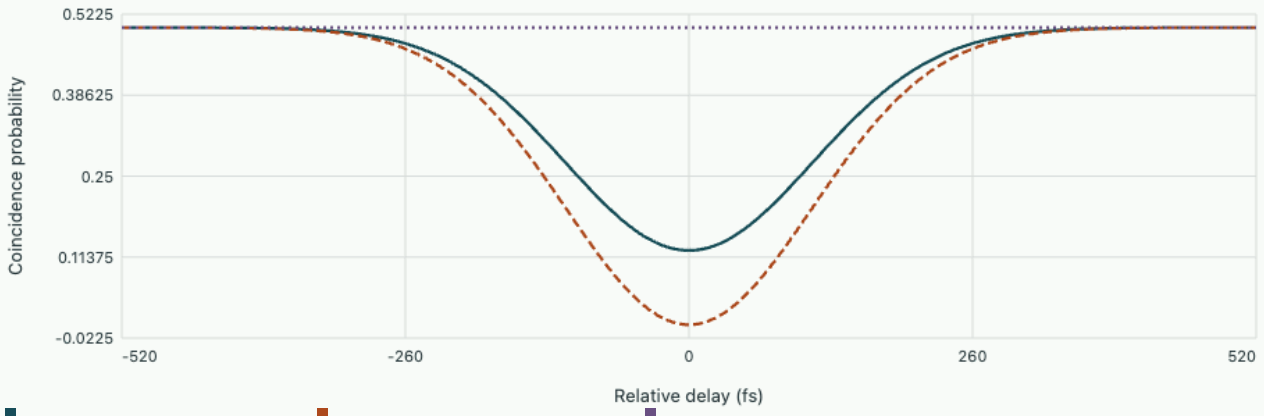
Adjust pulse duration and relative polarisation. Study peak-to-baseline depth before using the width to infer duration. A reduced visibility might equally arise from unobserved spectral distinguishability, multi-photon contamination or uncorrelated background counts. Blocking each input arm provides an independent accidental-count estimate. The ideal probability plotted assumes exactly one photon in each input and observation of both outputs irrespective of arrival time within the accepted window. A narrow coincidence window makes detection conditioning part of the model.

Open this experiment online

Try the calculation

Single-photon rms intensity width σ (fs) **80**

Relative polarisation angle (degrees) **30**



Zero-delay coincidence 0.125. Full width at half depth 266.42 fs.

Tutorial

Orthogonal linear polarisations yield no dip even at zero delay. Does a slower detector that cannot resolve polarisation automatically restore it?

▼ **Work through the solution**

No. The outgoing internal states remain orthogonal. Summing over an unobserved label is equivalent to tracing it out and removes the interference term. Deliberate projection of both polarisations onto a common state is another experiment with a different success probability.

For a lossless beam splitter of intensity reflectance R and transmittance $T = 1 - R$, predict the coincidence probabilities for perfectly distinguishable and identical photons.

▼ **Work through the solution**

The distinguishable alternatives give $R^2 + T^2$. For identical packets the amplitudes subtract, giving $(R - T)^2$. A perfectly indistinguishable pair need not give a zero minimum at an unbalanced splitter.

What delay has squared temporal overlap one half if the individual photon intensity has 100 fs rms width?

▼ **Work through the solution**

Set $\tau^2/(4\sigma^2) = \ln 2$, giving $|\tau| = 2(100)\sqrt{\ln 2} = 166.5$ fs. The measured full width spans both signs, about 333 fs.

Follow the single-photon and many-particle descriptions in the exchange chapter.

Computational quantum experiments · 6. Can an atom leave its initial state and still pass the final test?

A driven two-level atom undergoes repeated ideal measurements. Compare two laboratory records, one retaining only runs that passed every test and another retaining every run. The distinction is visible even when every measurement is made with the same apparatus.

Count histories or evolve a population

Write the resonant Hamiltonian in the rotating frame as $H = \hbar\Omega\sigma_x/2$. A pulse of duration T/N rotates the initial ground state into $\cos(\alpha/2)|0\rangle - i\sin(\alpha/2)|1\rangle$ with $\alpha = \Omega T/N$. The probability of passing one ground-state test is $c = \cos^2(\alpha/2)$. Each successful projection prepares the same ground state for the next interval. The probability that a complete record contains only successes is consequently $P_{\mathrm{all}} = c^N$.

For an unread measurement, keep the atoms detected in either state. The projection removes the off-diagonal density-matrix elements, leaving a classical mixture just before the next pulse. If p_j denotes the ground fraction after measurement j, the next pulse maps it to $p_{j+1} = cp_j + (1-c)(1-p_j)$. The excess population $z_j = 2p_j - 1$ then obeys $z_{j+1} = (2c-1)z_j$. Starting in the ground state gives

$$P_{\mathrm{final}} = \frac{1 + \cos^N(\Omega T/N)}{2}.$$

Leaving the ground state in one interval and returning in a later one contributes to the final fraction. It removes that run from the all-success record. The two probabilities therefore refer to different experimental questions, even though they use exactly the same sequence of pulses and measurements.

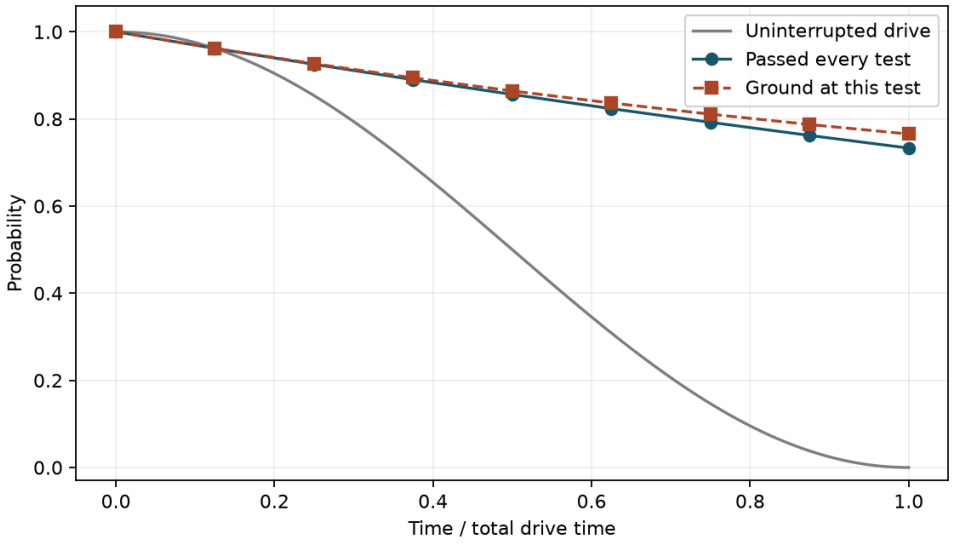
Follow a half-turn in two steps

Choose $\Omega T = \pi$ and make two evenly spaced measurements, including the one at the end. Each pulse is a quarter-turn of the Bloch vector and gives either result with equal probability following either previous outcome. The four two-result strings have probability one quarter each. Only 00 passes every test, whereas 00 and 10 both finish in the ground state. The fractions are therefore 0.25 and 0.50. With a single final measurement the undisturbed half-turn ends entirely in the excited state.

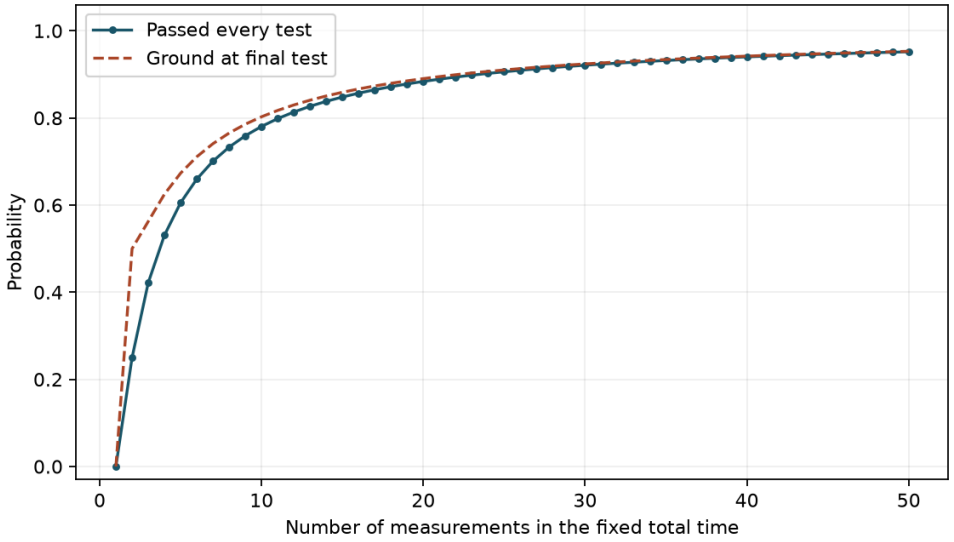
At fixed total duration and very large N, Taylor expansion gives

$$\log P_{\mathrm{all}} \simeq -(\Omega T)^2/(4N).$$

The survival probability tends to one because the transition probability per interval is quadratic in its short duration. This prediction assumes ideal instantaneous projective measurements and coherent rotations between them. A measurement with a fixed finite duration cannot be packed arbitrarily densely into a fixed time.



Markers represent measurements. The continuous reference curve is the ground population of an atom subjected to the same drive with every intermediate test omitted.



Keep the pulse area and total time fixed whilst shortening each free interval. An increasingly dense train of ideal observations preserves the initial state.

Exercise · recover a mixed population

If the atom initially occupies the ground state with probability 0.7 and has no coherences, find the final ground probability for four measurements and a total pulse area of π.

▼ **Work through the solution**

The initial contrast is 0.4. Each interval multiplies it by $\cos(\pi/4) = 1/\sqrt{2}$. Four steps give contrast 0.1, so the final ground fraction is $(1 + 0.1)/2 = 0.55$. The all-pass probability would instead contain a first measurement conditioned on the initially mixed preparation.

Exercise · budget the number of tests

Estimate the number of ideal tests needed for 99 percent all-pass survival during a coherent half-turn. Check the estimate with the editable input.

▼ **Work through the solution**

Equating the logarithm of 0.99 to $-\pi^2/(4N)$ gives $N \simeq 245.5$. Rounding up to 246 and evaluating the exact cosine power gives about 0.99002. This estimate concerns survival over an entire record. Filtering out only the atoms ending in the excited state would estimate a different probability.

Try it · Zeno

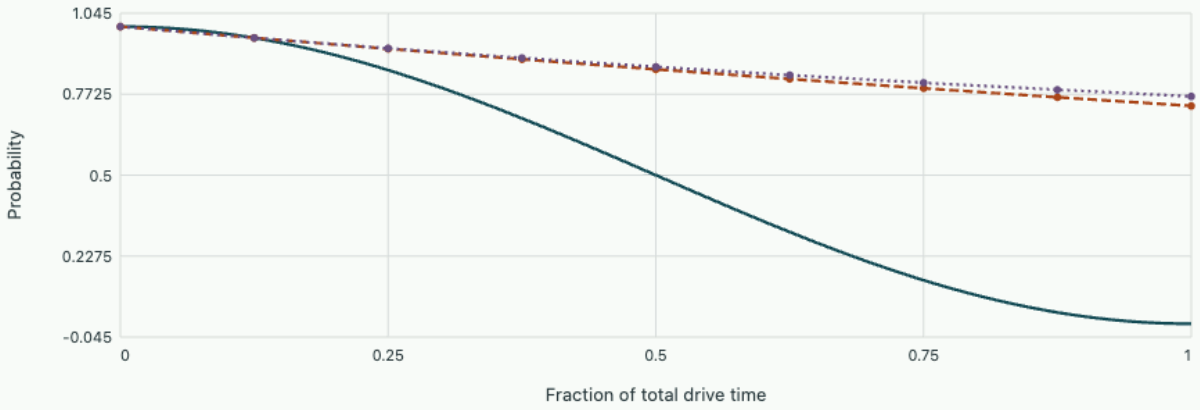
Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Total pulse area / π 1

Intermediate measurements 8



Selective survival 0.73313; unread final ground population 0.7654. Ideal instantaneous projective measurements are assumed.

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

Open editable notebook Download .ipynb