

A question-led introduction to quantum physics

Quantum before symbols

States, measurement and matter

David Elliman

Neuro-Symbolic

Modern physics through questions

Volume 1

Preface

Modern physics is fascinating, often surprising, and sometimes hidden behind unfamiliar mathematics. Popular accounts can awaken curiosity without showing enough of the reasoning to satisfy it. University textbooks supply that reasoning, but often assume mathematical fluency that a reader has not yet acquired, or has not used for some time.

At Neuro-Symbolic, we needed material to help able graduates understand the physics relevant to their work. These four books are our attempt to meet that need. They may also interest undergraduates, researchers entering a new field, and readers who want to understand more than a memorable analogy.

The books use questions and answers. We found this approach useful in workshops: a concrete puzzle gives people a reason to examine an idea and discuss their predictions. Each group of questions begins with the physical setting. Explanations use experiments, diagrams and simple arithmetic before introducing symbols. Optional mathematical sections connect that reasoning to the notation used in more advanced texts.

Our aim is to explain what the theories say, why physicists take them seriously, and where their explanations remain incomplete. Established results, assumptions of an example and unsettled proposals are identified separately. A useful calculation need not explain every ingredient that went into it.

We have taken care to check the physics and the examples, but that is no guarantee against error. Mistakes are our responsibility, and we welcome corrections from researchers and readers. The most useful feedback often begins with the exact sentence where an explanation stopped making sense.

Acknowledgement

OpenAI's Codex assisted with organisation, prose, diagrams and teaching checks. Responsibility for the content rests with the author.

How to use this volume

Begin with the introduction to the first part, which sets up the experiment and the reason for its questions. If you start with a later topic, read that part's introduction before choosing a question. The first eight questions build one experiment slowly. Later questions use its lessons to discuss particles, spin and entanglement. No calculus is needed. You will add numbers, square them and occasionally use a right-angled triangle. Three companion volumes take up particles and interactions, space and gravity, and the history of the universe. The optional pages at the back show how physicists compress these operations into symbols. They are not a prerequisite for reading the rest.

At each opening, make a prediction. A guess is enough; write it down if you can. When an explanation changes your mind, the interesting part is what changed it. The numbered "Try" questions have answers at the back. They are there to locate a difficulty, not to assign a mark.

The experiments are idealised where that makes the reasoning clearer. Their assumptions are stated beside them. A probability is a prediction about repeated preparations, not a promise that a short run will match a percentage exactly. Established quantum physics supplies the rules used here. The final case study asks what it would take for a proposed information-based theory to explain those rules. It does not present that proposal as established physics.

Each part begins by describing the physical setting and the problem that motivates its questions. Read that introduction first. Short transitions then prepare the next question; make the prediction after that context. The optional mathematical pages explain each new operation before using its compact notation.

Contents

Part one: One experiment, one question at a time	5
1 What should happen when there are two routes?	5
2 Why use arrows to calculate a chance?	7
3 How can one exit be completely dark?	9
4 What changes when one route is delayed?	11
5 Does observation turn a wave into a particle?	13
6 Can a record be undone?	15
7 What does a quantum state tell us?	17
8 Why can't a qubit be copied like a bit?	18
Part two: Particles that spread, move and stay put	20
9 What is a particle?	20
10 If a particle is a pattern, what moves?	22
11 Why can't position and momentum both be sharp?	24
12 How can a particle get through a barrier?	26
13 Can a quantum state stay still?	28
Part three: Spin without a spinning ball	30
14 Is spin really angular momentum?	30
15 Why does an electron need two full turns?	32
16 Why can't two electrons occupy the same state?	34

Part four: Correlations and interactions	36
17 How do entangled particles know each other's results?	36
18 Do all possible interactions collapse into one?	39
19 Are virtual photons real photons?	41
20 Can other particles be virtual too?	42
Part five: What the probability rule does	44
21 Why do we square an amplitude to get a probability?	44
22 Does decoherence explain why one result is recorded?	46
A closing case study: Information and physical theory	48
23 When does a useful description become a physical theory?	48
Optional: A closer look at the notation	51
24 How can a number describe an arrow?	51
25 What are a ket and a bra?	53
26 What does a matrix actually do?	55
27 What does the uncertainty formula say?	57
For checking: Answers to the Try questions	59
For returning to: Words, questions and sources	62

Part one

One experiment, one question at a time

Light provides a good starting point because we can prepare it, guide it along different routes and count where it is detected. With a bright beam we see an intensity. With a suitable faint source and detector we can collect individual clicks. Repeating the same preparation lets us ask how often each outcome occurs. A probability of one half means half the trials on average, not a promise about the next click.

Our experiment uses a partly reflecting optical element called a beam splitter. It sends light towards two routes, which mirrors bring together again. Such an arrangement is an interferometer: it compares contributions that have travelled through different parts of an apparatus. We will first describe what its detectors do, then build a calculation that accounts for the results. No matrix notation is needed.

The puzzle comes from recombining the routes. Their separate detection probabilities are not enough to predict the final counts. The first four questions introduce the extra information, called phase. The next four ask what becomes of it when a route is recorded, and what this teaches us about a quantum state.

1 What should happen when there are two routes?

Before reading on. A single-photon pulse enters an apparatus with two routes to an exit. If we block one route, must that exit become darker? Keep your answer until we have worked through the apparatus.

There is a sensible everyday answer: less light gets through, so the exit should be darker. Quantum experiments make us examine the words “gets through” more carefully. Blocking a route can make a previously

dark exit light up. We will calculate a simple example before trying to explain it.

Imagine a source adjusted to send one photon at a time. A photon is a quantum of light; for now, its useful experimental feature is that one successful detection gives one local click. A balanced beam splitter has two outputs. With a photon entering through just one input, repeated trials put half the clicks in each output, on average. “Balanced” describes that measured split. It does not mean that every photon becomes two half-energy detections.

Two mirrors guide the outgoing routes to a second splitter. There are now two exit detectors. We call them 0 and 1 because the names need to be short, not because either has a special role in nature. A delay in one route lets us adjust the relative timing of the two contributions reaching the second splitter. The drawing shows where everything goes.




Figure 1: A two-route interferometer. Light travels along the arrows. Detector 0 is to the right of the second splitter; detector 1 is above it. The little diagonal strokes represent the splitters and mirrors. The delay changes the relative phase; it is not a detector.

Our ideal apparatus loses no photons and produces no false clicks. The routes remain coherent: their relative phase is controlled. Both carry the same fixed polarisation, perpendicular to the plane of the drawing, and the splitters preserve it. Polarisation describes the orientation of light’s electric-field oscillation; here we keep that property fixed so that it does not distinguish the routes. We are studying the route choice, not a second experiment on polarisation. Real apparatus needs alignment and calibration to approach these conditions.

If we remove the second splitter and put detectors in the two routes, the results are ordinary enough: half upper, half lower. Replacing the second splitter changes what is being measured. We can adjust the delay so that every ideal trial clicks at exit 0. Exit 1 is dark. The two-route measurement and the exit measurement ask different questions of the prepared light.

That is the observation we need to understand. Merely saying that the photon secretly chose one route, with no further account of what the other route does, will not predict it. We need a rule that remembers more than the two route probabilities.

Try 1. With the second splitter removed, would 100 trials have to give exactly 50 clicks in each route?

We have described an exit becoming dark even though either route on its own can supply clicks there. Ordinary addition of positive chances cannot produce that cancellation. We need a calculation that keeps track of how the two contributions meet, before it assigns a probability.

2 Why use arrows to calculate a chance?

Before reading on. Two contributions each have probability one quarter when tested separately. When both are available, must their combined probability be one half?

For mutually exclusive classical alternatives, adding probabilities is often the right operation. A card may be a heart or a diamond; the two possibilities do not cancel each other. Coherent quantum alternatives require an earlier step. We assign each contribution an *amplitude*. Add the amplitudes for the same complete outcome, then turn the resulting amplitude into a probability.

An arrow on a sheet of paper is a useful way to keep track of an amplitude. Its length says how large it is. Its direction keeps track of its phase. For the whole amplitude of an outcome, square the arrow's length to obtain the probability. A length of 0.3 gives $0.3^2 = 0.09$: a nine per cent chance. A length of 0.5 gives $0.5^2 = 0.25$: one quarter. These are probabilities, so they have no unit such as metres or seconds.




Figure 2: Amplitude arrows live on a calculation page. They do not show the direction in which the photon moves. Rotating a lone arrow changes its phase but leaves its length, and hence its probability, unchanged.

Two arrows of length 0.5 can point the same way. Placed head to tail, they make a length of 1. Squaring gives probability 1. They can also point in opposite directions. Then their sum has zero length and gives probability zero. The separate probabilities were one quarter in both cases. The difference lies in the relative direction of the arrows, information that the separate probabilities discarded.

There is no claim here that a photon contains a small rotating pointer. An amplitude arrow is a representation of a calculation. It earns its place because experiments agree with the probabilities calculated from it. Later, complex numbers will provide a more compact description of exactly the same two-dimensional arrow.

The instruction to square the length is the *Born rule*. We are using it as part of quantum theory. Drawing the arrows does not derive the rule, and knowing the name does not explain why nature follows it. For this book, the first task is to learn what the rule says and where it differs from ordinary probability arithmetic.

One warning will matter throughout: do not add every arrow you can think of. Add contributions to the *same complete outcome*. If another system keeps a distinguishable record of the alternatives, we must include that record in the description. We will return to this restriction shortly.

Try 2. An arrow of length 0.3 turns through half a turn. What is its probability now? What if it is added to an unchanged arrow of length 0.3?

The arrow rule has two stages: combine contributions, then square the length of their sum. We can now apply it to the actual two-splitter apparatus. First we must specify the arrow length for each complete route and the phase change made by the optics.

3 How can one exit be completely dark?

Before reading on. At one exit, the upper-route contribution is +0.5 and the lower-route contribution is -0.5. Should we square them separately or add them first?

The minus sign is just a short description of an arrow pointing backwards along our chosen horizontal axis. The arithmetic becomes simple enough to do without a diagram. To get it right, we must first specify the apparatus phases. The fact that a splitter gives equal probabilities at its outputs does not by itself specify those phases.

We choose a calibrated convention for the whole interferometer, including its fixed reflection phases. The first splitter gives equal, aligned route amplitudes. At the second splitter, the upper and lower contributions arrive aligned at exit 0. At exit 1 they arrive with opposite signs. This is a physically possible balanced, lossless recombiner. It is not a rule that every reflection in every optical device adds the same half-turn.

A half-probability split needs an amplitude length whose square is one half. Finding the positive square root means asking which positive number, multiplied by itself, gives the required number. Here that length is about 0.707; squaring the rounded value gives approximately one half. There is one such factor at each of the two splitters. Successive amplitude factors multiply: the exact factors give $\frac{1}{\sqrt{2}} \times \frac{1}{\sqrt{2}} = \frac{1}{2}$. Thus the amplitude for either complete route to a specified exit has length 0.5. The displayed decimals below use that exact one-half value. We multiply amplitude factors for successive stages of one route. We add the contributions from alternative routes when they reach the same complete outcome. These are two different operations with different physical jobs.

At exit 0 the arithmetic is

$$(0.5 + 0.5)^2 = 1.$$

At exit 1 it is

$$(0.5 - 0.5)^2 = 0.$$

The probabilities add to 1, as they must in our lossless experiment. There is one photon to detect, and we have included all its exit possibilities. We have not cancelled energy out of existence: changing the phase redirects the detection probabilities between the exits.

Now block the lower route with an absorber. Only the upper contribution to each exit survives. Each exit then has probability one quarter per emitted photon, and the absorber takes the remaining half. The previously dark exit clicks in a quarter of the trials. Conditional on a photon getting past the absorber, the two exits are equally likely. This distinction between “per emitted photon” and “per surviving photon” prevents a misleading claim that blocking a route creates light.

The first prediction can now be answered. Blocking one route need not make every exit darker. The blocked route had contributed to cancellation at one exit. Removing it removes that cancellation. What fails is the attempt to describe the coherent experiment using route probabilities alone.

Try 3. If we square 0.5 and -0.5 separately and then add, what wrong probability would we predict for the dark exit?

The bright and dark exits were obtained at one adjusted setting. A real interferometer lets us change the relative optical delay while keeping both routes open. This gives a controlled way to test whether the direction of an amplitude arrow has physical consequences.

4 What changes when one route is delayed?

Before reading on. The two route probabilities stay at one half each while we adjust the delay. Could the exit probabilities nevertheless change from all at 0 to all at 1?

A delay changes how the two route contributions line up when they meet. For sufficiently narrow-band light and a delay that preserves their overlap, we can describe the change by rotating one amplitude arrow. The other route supplies the reference. A relative phase is this angle between contributions. We will express it in turns or degrees here; a quarter-turn is 90 degrees and a half-turn is 180 degrees.

Begin at the setting of the previous question. Turn the lower-route contribution through half a turn. At exit 0 the two arrows now oppose each other; at exit 1 the extra reversal makes them agree. The answers are exchanged: zero chance at 0 and certainty at 1. Neither arrow's individual length changed.

A quarter-turn is more interesting. One contribution points right and the other points up. Put them head to tail. The resultant is the diagonal of a right-angled triangle with sides 0.5 and 0.5. Pythagoras tells us that its squared length is $0.5^2 + 0.5^2 = 0.5$. We need that squared length, not the length itself, because we want the probability.




Figure 3: At a quarter-turn setting, the two amplitude contributions make a right angle. The orange diagonal is their sum. The probability at this exit is one half. These axes describe phase, not the upper and lower routes in physical space.

The other exit also has probability one half. A three-quarter-turn gives the same probabilities; a full turn restores the original result. Here is the whole cycle for the ideal, calibrated apparatus.

Relative turn	Exit 0	Exit 1
None	1	0
Quarter	0.5	0.5
Half	0	1
Three quarters	0.5	0.5
Full	1	0

The phase is experimentally meaningful because changing it changes the click statistics. It is not an additional property inferred from the upper-versus-lower counts: those counts remain unchanged. This gives us a practical reason to distinguish a *state* from the probabilities for one chosen measurement. Different states can agree on that one set of probabilities and disagree on another.

A much larger delay could separate the arriving pulses so that arrival time reveals the route. Then the simple phase-only calculation would no longer apply. Our assumptions about overlap and indistinguishability are part of the experiment, not optional decorations on its description.

Try 4. At three quarters of a turn, what are the two exit probabilities? Would the route counts, measured before recombination, distinguish that setting from no delay?


So far, nothing in the apparatus has left a record that distinguishes the two routes. A measurement changes that situation by correlating the route with another physical system. The next question asks what that new correlation does to the exit counts.

5 Does observation turn a wave into a particle?

Before reading on. A device marks which route the photon takes, but nobody reads the device. Will the unmarked interference pattern survive?

The usual double-slit question contains a trap. It pictures a wave moving through the apparatus, then a person looking at it and making it become a particle. The experiment does not require that sequence. Individual detections are local clicks both with and without interference. What changes is the distribution of those clicks over many trials.

Our two-route apparatus asks the same essential question in a simpler form. Instead of a screen with many possible positions, it has two exits. Put a route marker into it. Ideally, the marker neither absorbs the photon nor changes the route probabilities. It becomes correlated with the route: one distinguishable marker state for upper, another for lower. In the quantum description, we now have a joint state of photon and marker.



Different marker states remain distinguishable
even when no person reads them.

Figure 4: A perfect route marker makes the two alternatives distinguishable in the joint system. Calling its states “upper” and “lower” describes their correlation with the routes; it does not require a human-readable display.

At exit 0 there are now two complete outcomes: photon at 0 with an upper record, and photon at 0 with a lower record. Their amplitudes have length 0.5 each. Because the records are perfectly distinguishable, the probabilities for these alternatives add: $0.25 + 0.25 = 0.5$. The same happens at exit 1. The sign difference that previously produced cancellation cannot cancel two different marker states.

This is why merely leaving the room makes no difference. The physical correlation exists whether somebody inspects it or not. Nor does the explanation require a badly designed device giving an uncontrolled mechanical kick. Route information can be stored in another degree of freedom, with the loss of interference accounted for by the joint state. Partial distinguishability gives an intermediate reduction of contrast; our all-or-nothing calculation uses an ideal perfect marker.

In the double-slit experiment, many positions on a screen replace our two exit labels. Contributions through the slits can reinforce at some positions and cancel at others. Reliable slit information changes the joint outcomes and removes the corresponding interference. The local clicks themselves do not suddenly begin at that moment.

There remains a deeper question about why we experience one particular outcome. Describing the joint state and the lost interference does not, by itself, choose between interpretations of quantum mechanics. It does settle a narrower issue: consciousness is not an ingredient needed in this prediction. The apparatus and its correlations are enough.

Try 5. A marker has recorded the route. Its display is covered with paper. Does covering it undo the record and restore the dark exit?

A route marker has removed the original interference from the unselected exit data. Calling that process a measurement does not tell us whether it can be reversed. We must ask what happened to the marker and whether other systems acquired copies of its record.

6 Can a record be undone?

Before reading on. Compare reversing the interaction that made a route record with deleting a file that contains the result. Are they the same physical operation?

In the ideal quantum description, a controlled marking interaction can be reversed. A second device can act jointly on photon and marker so that either route leaves the marker in its original ready state. If the route coherence and the relevant phases have been preserved, the alternatives are again indistinguishable and can interfere at the recombiner. This is a demanding physical operation, not an instruction to forget what happened.

A mundane deletion is quite different. The record may already have been copied into a detector, a computer, scattered light and heat. Changing the displayed file does not reverse all those correlations. Any remaining route information in an unobserved system can still prevent the original interference from returning. There is no need for a person to remember it.




Figure 5: A result can spread into several records. Reversing one local operation need not undo the other copies. The arrows show the spread of information, not a universal microscopic law for detectors.

Another experiment is called a quantum eraser. Instead of reading the marker with a test that distinguishes upper from lower, we use a test whose alternatives are coherent combinations of those routes. Such a

choice of distinguishable alternatives is called a measurement basis. We then sort the exit data according to the marker result. In a suitable ideal arrangement, one selected group has one interference pattern and the other has the complementary pattern. Add the groups back together and the unsorted output still has equal exit probabilities.

For example, suppose half the runs belong to a selected group that, at our chosen setting, all click at 0. The other half belong to the complementary group and all click at 1. The full data are still half at each exit. Sorting reveals correlations. It does not alter a previously stored detector result or send information backwards in time. To discover which group a distant result belongs to, the observers have to compare records by ordinary communication.

These distinctions make “information was erased” an insufficient description of an experiment. Was a joint interaction coherently reversed? Was a classical file reset? Were results sorted according to another measurement? Each operation has its own prediction. A good account says which was performed.

Real detectors are useful precisely because they make records that survive and can be copied. That practical stability makes a full reversal extraordinarily hard. The difficulty explains much about why familiar measurements seem irreversible, while leaving the interpretational question about individual outcomes open.

Try 6. If two equally numerous sorted groups give opposite perfect interference outcomes, what will someone see who receives only the unsorted exit data?

The same upper-and-lower counts can accompany different final exit patterns. We therefore need a description that predicts more than one chosen measurement. Physicists call that description a quantum state. Our experiments now give the term a concrete job.

7 What does a quantum state tell us?

Before reading on. Two preparations give upper and lower route results equally often. Is that enough to say that they are the same state?

We already have a counterexample. At the zero-phase setting, recombining the routes sends every photon to exit 0. At the half-turn setting, every photon goes to exit 1. Yet measuring the routes directly gives equal upper and lower probabilities for either preparation. A state must contain enough information to predict more than one chosen test.

In this two-route example, two amplitude arrows provide a description of a pure state. Here that means one coherent amplitude description, rather than a statistical mixture of different preparations. Their squared lengths add to 1, and their relative direction matters when we mix the routes. Rotating both arrows by the same angle changes no prediction. That common rotation is called a *global phase*. Rotating one relative to the other can change the predictions. It is a *relative phase*.

The state belongs with a specification of the physical alternatives. “Upper” and “lower” mean something because we have identified two routes and a way to distinguish them. The same pair of numbers could describe two spin results in a different apparatus. Numbers alone do not identify the physical system, the measurement or its dynamics.

The two chosen alternatives are a *basis*: a set of possibilities used to describe the state. A different apparatus can test combinations of them. Our recombiner is the concrete example. Physicists use vectors and matrices to keep track of these combinations, but all the calculations so far used ordinary arrows and arithmetic. The notation can wait.

There is also a difference between preparing a coherent combination and choosing randomly between alternatives. Imagine a coin choosing upper

in half the runs and lower in the rest, with no phase relation between preparations. That mixture gives the same direct route counts as the coherent state. It will not produce our perfect interference when sent through the second splitter. A more general state description, called a density matrix, handles both such mixtures and pure states.

None of this makes the state a list of outcomes already waiting to be uncovered. It is the mathematical object from which the theory predicts the statistics of specified experiments. Whether it is also a complete description of reality is a further question, on which interpretations differ. We do not need to decide it to distinguish the two preparations in the laboratory.

Try 7. Turn both route arrows through a quarter-turn, keeping their angle relative to each other unchanged. Does our exit table change?

A detector's recorded answer can be copied into a notebook. A quantum state contains more than that answer: the phase comparisons also matter. The question of copying a state is therefore different from copying one distinguishable record of it.

8 Why can't a qubit be copied like a bit?

Before reading on. A machine can copy the labels upper and lower perfectly. Must it therefore copy every coherent combination of upper and lower?

A bit has two distinguishable values. Once a value is recorded in a suitable physical memory, it can be copied into another memory. A qubit is a quantum system described using two basis states and their coherent combinations. For a pure state, two complex amplitudes give the components in that chosen basis. This is what two-dimensional state space means; it is not a statement that the object is a flat sheet. Saying that it has "both values" misses the relative phase and suggests that we could simply read out two classical entries.

Our route marker already copies the chosen route label. Starting with a definite upper route, it produces an upper record. Starting with a definite lower route, it produces a lower record. Starting with a coherent

combination, however, it produces a correlated joint state. It does not leave two independent systems each carrying the original pure state. The change in interference shows why that distinction matters.

Try to specify what two genuine copies would require. Each copy of an equal coherent route state, tested separately in the route basis, would give half upper and half lower. Independent route tests on the two copies would then allow four pairs, each with probability one quarter: upper–upper, upper–lower, lower–upper and lower–lower. The route-marker correlation instead permits only the matching pairs, with probability one half each. The two physical states are different even though a single list of route probabilities looks the same.

Quantum evolution is linear: the same device acts on a combination of input amplitudes by combining its responses to those inputs. This rule makes the copying restriction general: there is no universal operation that perfectly copies an arbitrary unknown pure state into an independent second copy. This is the no-cloning theorem. It does not prohibit copying distinguishable basis labels, preparing many specimens of a known state, or using approximate copying machines with a specified loss of fidelity.

It also does not prevent quantum error correction. An error-correcting code distributes information across a larger joint system in a controlled way. The useful information is held in correlations; it is not a pile of independently readable copies of an unknown qubit. Repair works for a specified set of errors and a specified encoding. It is not a promise that all damage can be undone.

The distinction will matter in the closing case study. A register label, a propagating quantum state and a durable detector record play different roles. Calling all three “bits” would hide the very interference we have just learned to calculate.

Try 8. Does the no-cloning theorem forbid copying the result 0 from a detector’s classical display into a notebook? Explain why.

Part two

Particles that spread, move and stay put

The interferometer gave us two labelled routes. An electron moving through a laboratory has a much larger set of possible detection positions. We now extend the same idea: a preparation determines amplitudes, and amplitudes predict a distribution of results. Imagine a screen divided into many small collecting regions. Repeatedly preparing an electron and recording its arrival builds a pattern across those regions.

A concentrated pattern is called a localised packet. Its width, its motion and its internal phase pattern are different properties. To discuss them we also need momentum. In the familiar slow-motion limit, momentum is mass times velocity; it has a direction and is the quantity exchanged in a collision. Quantum theory relates a momentum component to a wave pattern.

These questions ask what is spread out, what moves, and why localisation limits the possible momentum distribution. We then use those ideas for tunnelling and for an atomic state whose probability pattern stays fixed. The drawings will represent distributions and energies, not photographs of little objects hidden inside them.

9 What is a particle?

Before reading on. Two electrons are prepared differently: one in a tightly confined state and one in a broad beam. Must they have the same spatial extent because they are the same kind of particle?

An electron has a reproducible charge and mass. Electrons also have characteristic spin properties, which we will meet in the next part. These identify the kind of particle. They do not tell us where a particular electron will be detected, or how widely its possible detection positions are spread.

Prepare an electron repeatedly in the same state and measure its position each time. The results form a distribution. A narrow distribution describes a well-localised preparation; a broad one describes a more spread-out preparation. The width is a property of that state and of the preparation, not a fixed diameter carried by every electron.

It helps to keep four questions separate. What species is it? What quantum state has been prepared? What observable are we measuring? What result has the apparatus recorded? For example, electron, a packet moving to the right, position and a click near the centre of the screen answer four different questions. Substituting one answer for another causes much of the confusion around the word particle.

In quantum field theory, a particle can be described as a quantum excitation of an appropriate field. A field is a physical system with quantities that can vary from place to place. Those independently variable quantities are its degrees of freedom; for an electromagnetic field they include the electric and magnetic field components. Its quantum states include particle excitations. This description allows creation, destruction and interactions, where a fixed number of little objects would be inadequate. The one-electron wavefunction is a useful reduced description in regimes where a one-particle treatment is appropriate. It is not the entire interacting field theory.

A particle is a wave is therefore not a sufficient replacement for a particle is a tiny ball. A water wave displaces water, and a string wave displaces string. A quantum amplitude predicts measurement probabilities. Its interpretation, its observables and its evolution have to be specified. Our amplitude arrows were already an example of a wave-like calculation without a little water surface waving inside the interferometer.

Nor is wavelength the same as packet width. A mode is one component wave pattern used to build a more general state. Its wavelength is the distance over which its phase completes a turn. The envelope describes how the overall amplitude grows and fades across a packet assembled from several such patterns. A nearly single-wavelength state can extend a very long way. A localised packet needs a range of wavelengths or momenta. The next two questions make that distinction useful.

There is no general size ranking in which every electron is more spread out than every quark. Quarks are confined within composite systems,

and assigning a free quark packet the role of an electron beam would be inappropriate. Comparisons must identify the physical state and the particular measure of extent.

Try 9. Which can change when we prepare a narrower electron packet: the electron's species charge, the position distribution, or both?

The size of an electron packet depends on its preparation. We can now ask how that distribution changes while the electron is left to move. A wavefunction assigns position-dependent amplitudes; its evolution changes their phases and hence the pattern of possible detections.

10 If a particle is a pattern, what moves?

Before reading on. A moving packet passes through a region. Must each little piece of the packet keep its identity and travel with it, as the parts of a rigid object would?

The stadium-wave analogy gets one thing right: a pattern can travel without the same material pieces travelling the whole distance. It also has a serious limitation. A quantum packet is not a crowd of people each making a definite classical movement. Its evolution concerns amplitudes and phases, and those must be combined before we calculate the probabilities of a measurement.

For a freely moving, nonrelativistic particle, we can construct a packet by combining momentum components. Each component acquires phase at a rate fixed by its energy. As their relative phases change, the position-dependent contributions reinforce in different places. A peak can move, and the packet can spread at the same time. Neither behaviour requires a permanent little ball to be hidden at the peak.




Figure 6: A schematic free packet at two times. The later distribution is broader and its centre has moved. The shapes illustrate those two effects; this is not a plotted solution for specified experimental parameters.

We can describe the experimental meaning without assigning a trajectory. Prepare many identical packets. In one set of runs, measure position after a short wait. In another set, wait longer. Compare the distributions. The later peak is farther along and may be less sharp. These are statistics from different repetitions, not repeated photographs that leave one isolated quantum state undisturbed.

Suppose, purely as an example of counting, the early detections fall into three adjacent collection regions with probabilities 0.2, 0.6 and 0.2. If a device translates the entire state one region to the right, the corresponding probabilities become 0, 0.2, 0.6 and 0.2 across four regions. They still total 1. That describes the effect of an ideal translation. It does not derive a law of motion or specify how a real device implements the translation. Many dynamics can move a peak while changing its shape as well.

The distinction matters when someone proposes a microscopic model. Labels for the internal states do not by themselves say how amplitudes propagate between places. A separate dynamical law must determine the evolution. A drawing of neighbouring registers is not yet such a law.

In the familiar theory, the energy and interaction rules determine that law. A potential can deflect a packet, bind it or scatter part of it. An environment can take away coherence or energy. Free spreading, however, is not automatically dissipation: a packet can broaden under reversible quantum evolution while its total probability and energy remain conserved.

Try 10. A free packet broadens without interacting with an environment. Does that fact alone show that it has lost energy as heat?

A narrow pulse needs a range of wave patterns to build it. Momentum measurements provide another way to describe those components. We are therefore led to ask whether a preparation can have both a very narrow position distribution and a very narrow momentum distribution.

11 Why can't position and momentum both be sharp?

Before reading on. If we remove every imperfection from our instruments, can we prepare a state with arbitrarily narrow position and momentum distributions at the same time?

The uncertainty principle is often described as the trouble caused by looking at a very small object. Measurement can certainly disturb a system, but that is not the whole principle. Even before any measurement, a state has a position distribution and a momentum distribution. Quantum theory places a lower bound on the product of their spreads.

Think of a long, regular train of waves. Its wavelength is well defined, but there is no single narrow region containing it. To make a short pulse, combine waves with different wavelengths. Arrange their phases so that they reinforce near the pulse and cancel elsewhere. The shorter the pulse, the broader the range needed in this construction. In quantum mechanics, wavelength and momentum are related, so this familiar mathematical trade-off becomes a physical restriction on state preparation.




Figure 7: Two schematic minimum-uncertainty packets. The narrow position distribution in the upper row goes with a broad momentum distribution. The lower row reverses the trade-off. Each horizontal axis is a different measured quantity; their scales are not being identified.

Here is a numerical comparison that avoids the constants for the moment. Take a family of ideal Gaussian packets that reaches the minimum uncertainty product. Gaussian names a particular smooth bell-shaped distribution; the minimum family here has no position-momentum correlation that enlarges the product. If one member has half the position spread of another, its momentum spread is twice as large. If its position spread is one quarter as large, its momentum spread is four times as large. The product stays at the same minimum.

The qualification about this family matters. General states can have a larger product. Merely doubling the position spread of an arbitrary state does not force its momentum spread to halve. The principle is a bound, not an equality obeyed by every preparation. The optional pages state the precise bound and explain what spread means.

A momentum spread does not mean an instrument has failed to find the one momentum the theory had secretly assigned to each member of the preparation. The operational statement is about the range of results predicted for repeated, equally prepared systems. Different interpretations can discuss underlying properties differently; they must recover that measurable bound.

There is also no permission here to violate energy conservation for a short time. Energy–time relations require care because time does not play exactly the same operator role as position in ordinary quantum mechanics. The slogan about borrowing energy will cause trouble when we reach tunnelling and virtual particles, so it is best not to adopt it.

Try 11. Within the minimum-uncertainty Gaussian family, we reduce the position spread by a factor of three. What happens to the momentum spread? Would the same equality follow for an arbitrary state?

A ball below the top of a hill rolls back unless it gains energy. The corresponding quantum problem uses a potential barrier: a region that would require more potential energy than the incoming particle has. The test is whether any transmitted detections are nevertheless possible.

12 How can a particle get through a barrier?

Before reading on. An incoming particle has less energy than the top of a finite barrier. If some detections occur beyond it, must the particle have briefly borrowed the missing energy?

A classical particle in one dimension cannot cross a barrier higher than its total energy unless something supplies more energy or changes the barrier. Quantum mechanics poses a different problem: find a wavefunction that obeys the evolution law on both sides and inside the barrier, with the required matching at its edges.

For a simple static barrier, the wavefunction inside includes exponentially varying terms rather than the freely travelling oscillations found outside. For a finite width, matching the solution across both edges generally leaves a nonzero transmitted amplitude. That amplitude predicts some detections on the far side. This is tunnelling.




Figure 8: A static one-dimensional barrier higher than the incoming particle's energy. The dashed line represents energy, not a trajectory through the barrier. In the elastic problem with equal potentials far to either side, the transmitted particle has the same energy as the incoming one.

The exponential dependence explains why thickness matters so strongly. For a simple illustration in the thick-barrier regime, suppose adding a particular thickness approximately halves the transmission amplitude. It then approximately quarters the transmission probability. Adding that thickness again gives another factor of one quarter in probability. These ratios illustrate exponential attenuation; they are not universal numbers for all barriers. The barrier shape, mass and incident energy determine the actual factors, and interface matching can matter.

No energy loan is needed. In a time-independent, elastic scattering problem, the incoming and outgoing energies agree. Inside the barrier, the quantum state cannot be read as a classical trajectory with a temporarily negative kinetic energy. The classical trajectory picture is the part that has ceased to apply.

How long was the particle inside? That sounds like a single question, but it hides a choice of measurement. We might compare the peak arrival time of a transmitted pulse with a reference pulse. We might weakly couple a clock to occupation of the barrier region. We might ask about a selected subset of transmitted events. Those procedures need not measure the same quantity.

A reshaped transmitted pulse can have a surprisingly early peak. Peak timing alone is not a signal speed, and does not establish faster-than-light communication. For a credible time claim we need the preparation, the clock, the selection of events and a causal account of the signal. Quantum mechanics allows well-defined timing experiments without supplying one universal classical residence time for every tunnelled particle.

For some opaque-barrier delay definitions, the delay saturates as the barrier width increases: the Hartman effect. Such a delay is not automatically a clock reading for one transmitted particle's time inside. The preparation, measurement and reshaping of the transmitted packet matter.[12]

Try 12. In the stated attenuation example, two added thicknesses reduce amplitude by a factor of four. By what factor is the probability reduced?

Not every wave pattern changes its measured shape. A standing wave on a string has a fixed pattern of nodes even while the string moves. Quantum stationary states raise a related but different question, because their fixed pattern is a probability distribution.

13 Can a quantum state stay still?

Before reading on. A stationary atomic state has a time-independent position distribution. Must its electron therefore be sitting motionless at one unknown point?

Nothing changes can mean several things. A probability distribution can remain fixed even though the state vector acquires a phase. A system can have steady average properties while individual measurements still give a range of results. A gas in thermal equilibrium can have unchanging temperature and pressure while its molecules continue to move and collide. These are different kinds of stillness.

An energy eigenstate is a state with one definite energy: an ideal energy measurement on it has a certain result. If the energy law does not change with time, such a state evolves by a common phase factor. That energy law is represented by the Hamiltonian, the mathematical operation that specifies energy and drives the time evolution. Because a common phase changes no measurement probabilities, the statistics of time-independent observables remain fixed. That is why such a state is called stationary.

We have already seen the arithmetic idea in the interferometer. Rotate every amplitude arrow by the same angle. Relative angles and squared lengths stay unchanged. A stationary energy state undergoes that kind of common rotation as time passes. It need not have sharp values of position or momentum.

For a simple illustration, suppose three position regions have probabilities 0.2, 0.6 and 0.2 in a stationary state. Waiting changes none of those numbers. Repeating the position measurement on fresh preparations still gives a range of answers. Time independence of the distribution has not reduced it to one position.

An atomic orbital is therefore not an orbit in the planetary sense. Its stationary probability density does not describe a tiny planet making a definite circuit. Some stationary states carry probability current; others do not. A stationary state's kinetic energy can also be nonzero. These statements concern different observables and need not contradict one another.

Now combine two states with different energies. Their phases advance at different rates, so their relative phase changes. Measurements sensitive to the combination can then show oscillations. A superposition of energy states is not generally stationary. This is the same reason relative phase mattered when we varied one route's delay.

Equilibrium is broader than a pure energy eigenstate. A thermal state is a mixture with weights determined by the preparation and temperature. Its steady statistical description does not imply that every microscopic process has stopped. Nor does a ground state's nonzero spread provide a source of unlimited usable energy: a ground state is already the lowest-energy state of the specified system.

Try 13. Two components acquire different phase changes during a wait. Can we assume the wait is unobservable merely because each component's own probability stayed unchanged?

Part three

Spin without a spinning ball

A position measurement does not exhaust what we can learn about an electron. There is also an internal property called spin. Its name comes from angular momentum, the quantity associated with rotations. A turning wheel and an orbiting body carry angular momentum, but an electron's spin does not require a rotating material surface.

Experiments can prepare and measure spin along a chosen direction. For an electron there are two possible results along that direction. The same two-amplitude method used for our two routes can describe these two alternatives. The physical apparatus has changed; the calculation has a familiar structure.

This part separates three questions. What does spin measure? How does a spin state change when we rotate it? What restrictions appear when two identical electrons are described together? The last question leads to the exclusion principle, which helps explain the arrangement of electrons in atoms. A spatial orbital is the spatial part of an electron state; the spin description is needed as well.

14 Is spin really angular momentum?

Before reading on. If an electron has angular momentum, must it have a material surface rotating around an axis?

Angular momentum is associated with rotations and with the torques that change rotational motion. Familiar objects have it because their parts move around an axis. Quantum systems also have an intrinsic contribution called spin. It affects how states transform under rotations and how angular momentum is exchanged. Those are physical effects, not just a convenient name.

For an electron, measuring one component of spin has two possible results. They are often called up and down along the chosen axis. Up is not an absolute direction built into the electron. The axis belongs to the measurement. Change the apparatus axis and you change the question.




Figure 9: An ideal spin measurement with two possible records. This is a diagram of the measurement’s logical structure, not a design for a free-electron Stern–Gerlach apparatus.

Prepare an ideal spin-half state up along a vertical axis. Measure along that same axis immediately and the up result is certain. Measure along a perpendicular axis instead and the two results are equally likely. The state was not incomplete merely because the second measurement has a spread of outcomes.

The number one half in spin one half refers to the quantum rotation representation and its angular momentum scale. It does not mean that an electron turns half as fast as a photon, or that half of its surface has been left out of a drawing. A classically spinning charged ball is not an adequate model of these properties.

How do we know the angular momentum is real? It enters conservation laws, rotation experiments and couplings to magnetic fields. The relation between spin and magnetic moment depends on the particle. Spin itself is not defined as whatever makes a magnet work, and a neutral particle can have spin. Likewise, orbital angular momentum and spin are distinct contributions that can interact and be transferred between systems.

Why do the familiar elementary matter particles have half-integer spin? Established relativistic quantum field theory connects half-integer spin to fermionic exchange statistics under its usual locality and consistency assumptions. It does not derive the observed list of particle species merely from the word matter. Composite objects can have integer total

spin: pairing constituent spins can give a bosonic atom. Thus all matter has half-integer spin would be too broad.

The following two questions separate two effects that are often mixed together: changing a spin state's phase by a rotation, and exchanging identical particles. They are connected in relativistic theory, but they are different operations in an experiment.

Try 14. For a spin-half particle prepared up along the vertical axis, what are the ideal probabilities for a perpendicular-axis spin measurement?

A rotation changes the orientation of an apparatus relative to the prepared spin. Ordinary objects return to their original orientation after one full turn. Quantum amplitudes introduce an extra possibility: the orientation may return while a phase comparison changes.

15 Why does an electron need two full turns?

Before reading on. Suppose a full rotation multiplies every amplitude of a spin state by minus one. Would measuring that isolated state immediately reveal the minus sign?

The familiar statement that an electron needs 720 degrees to return to itself is easy to misread. It does not describe a tiny ball whose surface is still misplaced after one turn. It describes the mathematical response of a spin-half state to a rotation.

A full turn, 360 degrees, multiplies its spinor by minus one. A spinor is the mathematical object used to represent this rotation behaviour. A second full turn supplies another minus sign, and their product is plus one. We could write that as $(-1) \times (-1) = +1$, but the arithmetic alone does not tell us whether the first minus sign is observable.

Recall global phase. If every amplitude of an isolated state changes sign, all squared lengths and relative phases stay the same. The physical predictions are unchanged. In that sense the isolated physical state already returns after one full turn. The two-turn statement concerns the spinor representation, not a new probability for an isolated spin measurement.

There is a way to observe the sign relative to something else. Split a particle's coherent path into two routes and rotate its spin through a full turn in one route, while leaving a reference route alone. After allowing for other dynamical phases, the minus sign is now a relative phase between the route contributions. Recombine them, and the interference can change.




Figure 10: The rotation is applied in one arm only. Its phase can then be compared with the untouched arm. A real experiment must also calibrate phases from the interaction used to perform the rotation.

Our original exit calculation gives the arithmetic. Two aligned route contributions of $+0.5$ gave certainty at one exit. Change one to -0.5 and that exit becomes dark. Rotate through another full turn in the controlled arm and the sign returns, restoring the original ideal pattern. The effect is a comparison between coherent alternatives.

The belt trick is sometimes used to illustrate the special topology of rotations: a twisted belt can behave differently after one and two turns when its ends and allowed motions are specified. It can help the imagination, but it is not a derivation of the electron's spin, nor evidence that the electron is literally attached to an invisible belt.

What earns the spinor description its place is the collection of rotation and interference predictions it gets right. A theory proposing to explain the electron must recover that behaviour; drawing a twice-twisted object is not enough.

Try 15. Why can a minus sign be unobservable when applied to the whole state but observable when applied to only one of two coherent routes?

We have described one electron's spatial state and its spin. In an atom there are several identical electrons, so the joint state must also obey the rule for exchanging their labels. That rule determines which combinations of complete one-electron states are possible.

16 Why can't two electrons occupy the same state?

Before reading on. An atomic orbital holds two electrons with opposite spin states. Is that a violation of the rule that no two electrons can share a state?

The important word is *complete*. A one-electron state includes its spatial and spin description, and any other relevant labels. Two electrons can share the same spatial orbital if their spin states differ. They then occupy different complete one-electron states.

Electrons are identical fermions. Exchanging their particle labels changes the sign of the joint wavefunction. This antisymmetry is not an ordinary repulsive force that grows stronger when two electrons get close. It is a restriction on the allowed states of the combined system.

Imagine trying to put both electrons in the same complete state. One term assigns the first copy of that state to the first label and the second copy to the second label. The exchanged term is identical, but antisymmetry gives it the opposite sign. The two cancel. There is no nonzero two-electron state of this form to normalise. Normalising means scaling amplitudes so that the probabilities of all possible outcomes add to one. A set of zero amplitudes cannot be scaled into a possible preparation. That is the Pauli exclusion principle in the independent-orbital description.




Figure 11: Opposite spin states allow two electrons in one spatial orbital. Two copies of the same complete one-electron state do not produce an allowed fermionic state. The arrows label spin states, not rotating surfaces.

This changes the structure of atoms. Electrons cannot all pile into a single lowest complete state. Additional electrons occupy other available states, with their energies also affected by interactions. The periodic table depends on both the exclusion principle and the physical forces; exclusion alone is not a numerical theory of every atomic energy.

The principle also contributes to the pressure of a dense fermion gas. Confining many fermions forces occupation of states with larger momenta and energies. This does not require a new short-range force called an exclusion force. Effective forces used in approximate descriptions must not be confused with the underlying state restriction.

Bosons have different exchange statistics. Identical bosons can occupy the same single-particle state; their exchanged contributions add with the symmetric sign. That contrast does not mean every collection of bosons automatically condenses, since temperature, interactions and confinement also matter.

In a strongly interacting system, the simple picture of electrons assigned to named orbitals can become inadequate. Fermionic antisymmetry still applies. The orbital example is a way into the rule, not a claim that every many-electron state is a list of occupied classical slots.

Try 16. Do two electrons in the same spatial orbital with opposite spin states occupy the same complete one-electron state?

Part four

Correlations and interactions

So far we have mostly prepared one system and measured it. We have also met a joint system: a photon together with its route marker. A joint state predicts combinations of results. These can be correlated even when neither result, considered on its own, tells us much.

To see a correlation, two people can compare their lists of results after an experiment. The question is whether the pattern in the paired lists can be explained by ordinary instructions fixed at preparation. Entanglement lets quantum theory make predictions that go beyond that explanation, under conditions we will spell out.

We then turn to interactions. A scattering experiment prepares incoming particles and records what emerges after they interact. Several contributions may lead to the same final record, so the amplitude arithmetic returns. Physicists often organise those contributions with diagrams. Learning what a diagram calculates, and what it does not show, will help us discuss virtual particles without treating every line as an observed trajectory.

17 How do entangled particles know each other's results?

Before reading on. Two distant observers obtain strongly correlated results. Does that fact alone prove that a message travelled between them?

Ordinary correlations need no message sent at the moment of measurement. Put a left glove in one box and a right glove in another. Open one box far away and you immediately know what is in the other. The gloves had their handedness before the boxes were opened.

An entangled state cannot in general be described as independent states for its parts, even allowing for classical uncertainty about which independent states were prepared. Our photon and route marker provided one example. With two separated particles, the interesting test is whether their correlations for different choices of measurement can all be explained by local instructions agreed in advance.




Figure 12: Each observer chooses a measurement and records a result locally. Correlations become visible when the records are compared. The entangled pair does not let either observer choose the other’s local result.

A version of the Bell test can be expressed as a game. Each observer receives a randomly chosen question, numbered 0 or 1, and gives an answer, also 0 or 1. In the physical test, a question labels a measurement setting, such as one of two orientations of an analyser. An answer labels its two possible recorded results. These are instructions to the apparatus and detector outcomes, not questions the particles are asked to understand. They cannot communicate after receiving their questions. The four question pairs are equally likely. They win if their answers agree, except when both questions are 1, when their answers must differ.

Could a prearranged answer sheet win every time? To win on question pairs 00 and 01, the right observer must give the same answer to both questions, matching the left observer’s answer to 0. Winning on 10 then requires the left observer’s answer to 1 to match it too. But now 11 also produces matching answers, and loses. At least one of the four pairs must lose. Always answering 0 wins three of them, so the best classical success rate is 75 per cent. Randomly mixing answer sheets cannot raise that maximum.

Suitable measurements on shared entangled qubits achieve about 85.4 per cent in the ideal quantum prediction. This is the CHSH form of a Bell test, expressed as a game; the reading notes give the formal sources. Actual experiments estimate correlations from finite data and need careful controls. Bell violations have been observed in experiments designed to close the main locality and detection loopholes.[4, 5, 6]

The assumptions matter. The questions must be chosen independently of the hidden instructions, and the local response cannot depend on a message carrying the distant question during the trial. The conclusion excludes that class of local models. It is not a proof against every conceivable hidden-variable interpretation or a demonstration of controllable faster-than-light messaging.

In the ideal entangled strategy, either observer's own answer is still a fair random bit. Changing the distant measurement changes the joint correlations without changing that local half-and-half distribution. An observer who sees only one local list cannot read the other observer's choice from it.

The phrase the particles know therefore adds a mechanism that the prediction does not supply. Quantum theory gives a joint state and probabilities for joint measurements. Interpretations disagree about how best to describe that fact. The experimentally useful distinction is between a correlation revealed by comparing records and a signal encoded in one local record.

Try 17. Both observers always answer 0 in the game. Which question pair loses, and what fraction of uniformly chosen rounds do they win?

The Bell experiment compared measurements on two separated systems. A collision asks a different joint question: what results after incoming particles interact? We again need amplitudes, because different contributions may lead to the same complete outgoing record.

18 Do all possible interactions collapse into one?

Before reading on. Two contributions lead to the same scattering record. Is their combined probability necessarily the sum of their separate probabilities?

A scattering experiment prepares incoming particles, lets them interact, and records outgoing detections. The calculation predicts a distribution of final outcomes: directions, energies, particle species and whatever else the apparatus measures.

This is a larger version of our interferometer problem. Contributions to an indistinguishable complete outcome are combined as amplitudes. Only after that do we obtain the probability. If extra records distinguish alternatives, the joint description must include them, just as it included the route marker.




Figure 13: Two contributions to one amplitude. The branches in this diagram are terms in a calculation, not observed paths through a hidden intermediate apparatus.

Take a deliberately small arithmetic example. Suppose two real contributions to a selected outcome are $+0.3$ and -0.1 . Their sum is 0.2 , giving probability $0.2^2 = 0.04$. Squaring first and adding would give $0.09 + 0.01 = 0.10$, a different prediction. Other final outcomes must supply the remaining probability in a complete experiment. These

invented amplitudes illustrate interference; they are not a calculation of an electron cross-section.

In quantum field theory, a common approximation expands an interaction amplitude in powers of the coupling strength. A Feynman diagram labels a contribution in that expansion. A simple diagram may give the leading term, with more complicated ones providing corrections. Whether the expansion is useful depends on the physical regime and the approximation. The diagrams are not snapshots of little events that a detector watched between the incoming and outgoing particles.

There is also a path-integral formulation that sums contributions associated with histories. That is a mathematical organisation of the prediction. We should not silently turn it into the claim that a detector first saw every history occur and then compressed them into one recorded history.

What the apparatus supplies is a particular final record in each run. Repeating the preparation lets us compare its distribution with the calculated probabilities. The question of how to interpret a single outcome remains, but it does not change the amplitude arithmetic.

This is why a convincing interaction theory needs more than an attractive picture of particles meeting. It needs the amplitudes, their phases, the allowed final states and a rule connecting them to observed counts. Without those, all the possibilities interact is too vague to test.

Try 18. For contributions $+0.3$ and -0.1 to the same unmarked outcome, calculate the probability. What step would incorrectly give 0.10 ?

Interaction calculations are often arranged as diagrams with incoming and outgoing lines. The detector is associated with the specified external states. Lines drawn between interaction points have a different role, which the word virtual is meant to identify.

19 Are virtual photons real photons?

Before reading on. In a diagram for two charged particles interacting, an internal photon line appears. Could we place a detector on that line and read off a little photon's flight?

The electromagnetic field mediates interactions between charged systems. Quantum electrodynamics allows us to calculate those interactions. In a perturbative calculation, an internal photon line is a compact representation of one mathematical ingredient: a propagator connecting interaction vertices. A vertex is a point in the diagram where fields interact. A propagator describes how the field enters the amplitude between such insertions.

The line is called a virtual photon. It is not an extra species of photon, and it is not a detected free photon that temporarily evades the usual energy–momentum relation. Some internal variables are fixed by the measured external momenta; others, such as the momentum circulating around a loop, must be summed or integrated over. Integration is the continuous counterpart of adding contributions from many alternatives. They are not individually measured trajectories.

An outgoing photon that reaches a detector is different. It belongs to the specified physical final state. If we add a detector inside a scattering experiment and require it to absorb radiation, we have changed the physical process and the outcomes being calculated. We cannot infer that we have simply photographed a previously hidden internal line.

There is a further reason for caution. Intermediate descriptions can depend on a gauge choice, a mathematical redundancy in the description of the electromagnetic field. Observable predictions must be independent of that choice. Assigning a literal picture to each intermediate contribution can attach physical meaning to features that disappear when the same calculation is written another way.

None of this makes the measured force imaginary. Charged objects really accelerate; atoms really have energy levels; scattering rates really change with angle and energy. The theory's success is assessed through those observables. Virtual marks the role of an internal contribution, not a judgement that the interaction has no physical effect.

The popular account in which a photon borrows energy, survives for a time allowed by uncertainty, and then pays it back should be avoided. The calculation obeys the appropriate conservation laws. Its internal lines need not satisfy the free-particle relation imposed on real incoming or outgoing states. These are different statements; neither requires a short-lived exception to conservation of energy.

Try 19. Would adding a detector for emitted radiation leave the original scattering experiment unchanged?

The distinction between an internal contribution and a detected outgoing particle is not peculiar to light. The same method uses contributions from other fields. The physical test remains the measured reaction, including its energy and momentum balance.

20 Can other particles be virtual too?

Before reading on. If an electron line occurs inside a diagram, does that imply an additional electron could have been detected there without changing the process?

The same distinction applies to fields other than the electromagnetic field. Perturbative calculations can contain internal electron, quark or other field lines, where the theory's interactions allow them. Virtual describes their position within the calculation. It does not define a second collection of particles with their own relaxed physical laws.

Consider a correction to electromagnetic propagation involving an electron loop. The corresponding contribution affects measurable predictions. Calling it an electron-positron fluctuation can be a useful mnemonic, but a cartoon of a pair appearing, following two ordinary paths and vanishing is not a literal reconstruction of events observed by a detector.

Producing a real electron and positron in the final state is a different process. It needs enough available energy and a way to satisfy momentum, charge and the other relevant conservation laws. A photon travelling alone in empty space cannot turn into an isolated real electron–positron pair while conserving four-momentum, the combined energy and three momentum components. Another particle or an external field can change the kinematics.

Internal contributions can affect an observable even when the corresponding real-particle production is not energetically available. That fact is useful: precision measurements can be sensitive to effects of particles that are not produced directly in the experiment. Extracting such information requires a controlled calculation and experimental uncertainties. A suggestive loop drawing does not by itself identify a new particle.

Nor does every internal line represent an independently observable particle species. Some formulations introduce auxiliary fields to organise the calculation. Physical interpretation belongs to the consistent theory and its observables, not simply to counting the lines in a chosen diagram.

This returns us to the distinction with which the volume began. A calculation may contain more structure than one measurement records. Some of that structure is essential for predicting interference and correlations; some reflects a chosen representation. To decide which is which, ask what preparation and measurement would distinguish two proposals. That question is more useful than asking whether every mark in a diagram is really there.

Try 20. Does a virtual-pair contribution permit a lone photon in empty space to produce a real electron–positron pair without any other participant?

Part five

What the probability rule does

We have used amplitude arrows to predict detector counts. Those predictions work only after we supply a rule connecting the arrows to probabilities. We have also described how a detector becomes correlated with a particle. These are two different pieces of the explanation.

The questions below ask what has been assumed. A successful rule for frequencies does not automatically explain why an individual run has the particular result written in a notebook. Likewise, explaining how a record spreads into an apparatus does not by itself derive the probability rule. Keeping these tasks separate makes the measurement problem less mysterious to state, even though it does not settle it.

21 Why do we square an amplitude to get a probability?

Before reading on. If two contributions cancel as arrows, could we square their lengths separately and still predict the dark exit?

No. Two equal and opposite contributions have a zero sum, but each has a positive squared length. The interference experiment tells us to combine the amplitudes for an indistinguishable outcome first. We then need a rule that turns the resulting amplitude into a nonnegative probability.

The Born rule uses its squared length. An amplitude of length one half gives probability one quarter. Doubling that length to one gives certainty in a normalised example. Squaring is therefore a specific physical rule, not another word for counting clicks.

It has useful consistency properties. Rotating every amplitude by a common angle leaves all probabilities unchanged. For distinguishable alternatives, the squared lengths add to one in a normalised state.

Unitary evolution preserves that total. These features explain how the rule fits the rest of quantum theory; they do not derive it merely from the existence of arrows.

There are mathematical theorems that make the rule more constrained than an arbitrary choice. Gleason's theorem starts with a Hilbert space: a space of state vectors with an inner product, the rule that compares two states and assigns their overlap. A projection test asks whether the state lies in a specified set of alternatives; testing for the upper route is a simple example. The theorem assumes probabilities assigned consistently to these ideal tests, adding to one for a complete set of exclusive outcomes. In its usual form it applies in dimension at least three. Under those assumptions the allowed probability assignments have the quantum form.[11]

The consistency requirement includes assigning the same probability to a projection whenever it appears in different complete tests. This is a substantive assumption about the probability assignments, not something proved merely by giving outcomes record labels.

The qualifications matter. An isolated two-state qubit is outside that original theorem's scope. Extensions involving more general measurements can cover it, but require their own assumptions. Nor does the theorem produce the Hilbert space, the physical measurements or a detector from an unspecified collection of bits.

For this book, the honest starting point is that the Born rule is part of the quantum theory tested by the experiments. Deeper reconstructions may explain why it follows from other physically motivated requirements. The requirements must be stated so that we can see what has been gained. A ledger of outcomes alone supplies frequencies, not the missing rule that predicts them before the experiment.

Try 21. Two opposite contributions have lengths one half each. What probability follows after adding them? What incorrect result follows from adding their separate squared lengths?

A probability rule predicts the distribution of results over many runs. An individual run leaves one result on the display. To connect those statements, we need to examine what the quantum description says about the measuring apparatus itself.

22 Does decoherence explain why one result is recorded?

Before reading on. If a detector becomes correlated with two possible outcomes, does merely ignoring its environment select one of them?

Consider a detector with two possible pointer readings. If the incoming state takes one route, an ideal interaction correlates it with the first reading. If it takes the other route, the interaction produces the second. For a coherent combination of routes, linear evolution gives a joint combination of the two particle–detector alternatives.

Now let the apparatus interact with air, scattered light and its other degrees of freedom. These can retain distinguishable records of the pointer. Interference between the alternatives then becomes inaccessible to ordinary measurements on the pointer alone. This is decoherence. It helps explain why stable apparatus records behave so differently from the isolated amplitude arrows in the interferometer.

If the two alternatives had equal weights, the reduced description of the pointer predicts half the records of each kind. It may give the same local measurement statistics as an ordinary uncertain choice between two already definite readings. But the full quantum description still includes correlations with the environment. Discarding those correlations from a calculation is not, by itself, a physical operation choosing one result.

This distinction is part of the measurement problem. The laboratory contains individual records, while an unmodified unitary description can retain the joint alternatives. What connects that description to the outcome experienced and reported in one run?

Different accounts answer differently. A collapse postulate supplies an outcome rule; dynamical-collapse proposals change the evolution and must face experimental tests. Everett-type interpretations keep unitary

evolution and analyse records within branches, meaning the effectively separate correlated alternatives left by decoherence. Hidden-variable theories add further structure specifying outcomes. Mentioning these possibilities does not establish that they are interchangeable or that one has been derived here.

Decoherence remains a real physical process whichever interpretative account is used. What it does not do alone is turn the mathematical removal of unobserved variables from a description into a unique global outcome. That removal is called taking a partial trace. It also does not make consciousness necessary for a detector to leave a record.

For the calculations in this volume, the preparation, the measurement and the probability rule are enough to predict counts. For a proposed foundation, the account of records must be stated as well. Calling an apparatus a ledger is useful bookkeeping, but renaming it does not settle this additional question.

Try 22. What does decoherence explain about the two pointer readings, and what extra question about a single run does it leave open?

A closing case study

Information and physical theory

Could an information-based description explain the physical rules we have used? This short case study considers a proposal by the author. It asks what a collection of quantum states and records would have to supply before it could explain a physical experiment.

A register is a system with specified internal states. An algebra gives rules for combining operations on those states. A ledger records selected outcomes or conserved quantities. These are useful kinds of description, but they answer different questions. A detector ledger can list every click without containing the relative phases needed to predict a changed experiment.

The question below applies that distinction to the proposal. It is an assessment of what the existing work establishes, not a new research calculation. The quantum results in the preceding parts do not depend on accepting this framework.

23 When does a useful description become a physical theory?

Before reading on. A register has enough labels to name familiar particles, and a ledger keeps count of its records. Have we already explained how those particles move and interact?

A catalogue of labels is a useful start, but it does not predict an experiment. The same set of labels can be used by theories with different motions, interactions and measurement probabilities.

The proposal explored finite registers, error-correcting structures, algebras of operations and ledgers of records. An algebra specifies how operations combine. A ledger specifies what is recorded and how records

are updated. Both can make an account precise. Neither automatically selects the physical interactions or their strengths.

Suppose a register label is assigned the name electron. That is a proposed dictionary entry. To make it a physical identification, the model must recover the electron's charge, spin, statistics and dynamics in the relevant regime. A local label also does not fix the size of an electron packet. Internal identity and spatial state have to be related by an explicit field map and evolution law.

The ledger raises a different question. Our interferometer can have identical route counts at two phase settings and opposite exit predictions. A ledger containing only those counts has omitted information needed to predict the later experiment. A richer record scheme may preserve more, but that preservation must be demonstrated for specified operations. Naming a conservation rule does not establish which microscopic couplings nature uses.

Some work in the framework produced conditional results worth keeping. The published P1 preprint studies specified gravitational and fermionic actions on a bond-centred lattice. It establishes properties of those chosen constructions within its stated assumptions. Its gravitational ingredients are supplied from existing gravitational theory. It does not derive gravity from information alone or establish one unified interacting theory of matter, light and geometry. The preprint and its evidence archive are separately available.[9, 10]

The conditional lattice results do not establish the proposed broader unification. That limitation concerns this construction; it is not a theorem ruling out all information-based approaches to physics.

A useful test of any future proposal would distinguish three levels. First, what has been defined or assumed? Second, what follows mathematically from those inputs? Third, which prediction discriminates the proposal from alternatives in an experiment? A good notation may help with the first two while leaving the third unanswered.

The questions in this volume are a starting set of demands. A theory needs the relative phase that changes an exit, the joint state that gives Bell correlations, the exchange rule that gives exclusion, and the

dynamics that moves an extended state. A successful record of a click is part of that account. It is not a replacement for the account.

Force-carrying fields are already quantum in the established theory. Putting matter variables on sites and gauge variables on links is one possible discretisation, not a proof that bosons are merely classical connections. A proposed alternative encoding must still supply their quantum states and interactions.

Try 23. Two models keep the same route-count ledger but predict different recombined exits. Which experiment would distinguish them?

Optional

A closer look at the notation

We have already added amplitude arrows, turned them, and used their squared lengths to predict counts. The optional pages translate those operations into the notation used in physics books. The notation is worth learning because it lets us keep track of a more complicated experiment without drawing every arrow again.

There are four steps. A complex number records the two components of one amplitude arrow. A column collects the amplitudes for different alternatives. A matrix tells us how a device changes that column. Finally, an inequality states a limit on the spreads of two measurement distributions. These are different jobs; recognising the job matters more than remembering a symbol.

Read one question at a time with paper beside you. Use the numerical example to check what each operation means. You can return to the main text whenever the notation gets in the way: none of these pages is required to follow the earlier argument.

24 How can a number describe an arrow?

Before reading on. An amplitude arrow has horizontal component 0.5 and vertical component 0.5. Do we need to calculate its length before finding its probability?

The horizontal component is how far the arrow travels across the calculation page; the vertical component is how far it travels up. An arrow with components 0.5 and 0.5 is the diagonal we used for the quarter-turn interferometer. Pythagoras says its squared length is the sum of the squares of those components: $0.25 + 0.25 = 0.5$. That squared length is already the probability. There is no need to take a square root and then square again.

A complex number is a compact way to record the two components. We write this one as $0.5 + 0.5i$. The first number is the horizontal component. The coefficient of i is the vertical component. The letter does not stand for an unknown quantity to be found later; it identifies a new unit with the multiplication rule $i^2 = -1$.

Here is why that rule fits the arrows. Multiplication by i turns an arrow through a quarter-turn anticlockwise. Start with an arrow pointing right. One such turn makes it point up; another makes it point left. Two quarter-turns therefore have the same effect as multiplication by minus one. Four bring it back to its initial direction. The algebra records these rotations.




Figure 14: One amplitude and its complex conjugate. Reflection reverses the vertical component, while the horizontal component and the arrow's length stay the same. These are directions on the amplitude calculation page, not in the laboratory.

The word imaginary is the conventional name for the marked component. It does not mean that this component can be ignored. Both components contribute to a real detector probability.

The conjugate reflects an arrow in the horizontal axis. Our example becomes $0.5 - 0.5i$. To multiply two complex numbers, multiply each term in the first bracket by each term in the second and then collect the results. For this conjugate pair, the two cross-terms are equal and opposite. The product of the two vertical terms is positive because it includes minus one times the square of the imaginary unit. The result is

$$(0.5 - 0.5i)(0.5 + 0.5i) = 0.25 + 0.25 = 0.5.$$

Thus multiplying by the conjugate gives the squared length. Multiplying a complex number by itself generally does not. When physicists say square the amplitude, the intended probability rule is the squared magnitude, or squared arrow length.

Adding complex numbers is simpler: add the horizontal components and add the vertical components. The interferometer's horizontal contribution 0.5 and vertical contribution of length 0.5 give the arrow we started with. The notation has introduced no new physical assumption; it makes the previous arrow calculation easier to repeat.

Try 24. An amplitude has horizontal component 0.6 and vertical component 0.8. What is its squared length? Does conjugating it change that value?

One complex number describes one amplitude arrow. Our two-route state requires two such numbers, one for each route. We need a way to collect them and to calculate the amplitude for a specified state test. Columns and rows will organise those two tasks.

25 What are a ket and a bra?

Before reading on. Testing a prepared state against itself should succeed with certainty. Can simply turning its amplitude column sideways always give the row needed to calculate that probability?

First separate two uses of the word component. One amplitude arrow has horizontal and vertical components on our calculation page. A two-route state, however, needs two complete amplitudes: one for the upper route and one for the lower. Each of these may itself be complex. We must not confuse a route label with an axis in the amplitude plane.

Put the upper amplitude first and the lower amplitude second in a column. This ordered list is a vector. Upper and lower are orthogonal alternatives: a perfect route measurement can distinguish them with certainty. Their squared amplitude lengths add to the total probability. A ket is a notation for the whole state vector, which saves us from displaying every entry whenever we refer to it.

How does a row act on a column? Multiply corresponding entries and add the products. As a real-number example, a row containing 0.6 and 0.8, applied to a column containing 0.6 and 0.8, gives 0.36 plus 0.64, which is 1. The row containing 1 and 0 selects just the first entry. The row containing 0 and 1 selects the second. This is an arithmetic rule, not a new apparatus.

Complex entries require one further step. Take an upper amplitude 0.6 and a lower amplitude $0.8i$. Their probabilities are 0.36 and 0.64, so the state is normalised: the probabilities add to one. To compare the state with itself, each amplitude must be multiplied by its conjugate. The column and the corresponding row are therefore

$$\begin{pmatrix} 0.6 \\ 0.8i \end{pmatrix}, \quad \begin{pmatrix} 0.6 & -0.8i \end{pmatrix}.$$

The row is called the bra. Obtaining it involves transposing, which means turning column into row, and conjugating each entry. Its product with the original column is 0.36 plus 0.64, or 1. Merely transposing without conjugating would give 0.36 minus 0.64, or -0.28. That cannot be the probability check for a state tested against itself.

The same operation compares different states. The row (1,0) selects the upper component of our example, giving amplitude 0.6 and hence probability 0.36. The row (0,1) selects the lower component, giving amplitude $0.8i$ and probability 0.64. A bra associated with another normalised state tests a coherent combination. That test requires an appropriate measurement arrangement, such as a recombiner followed by a detector.

If we call the input state ψ , its ket is $|\psi\rangle$ and its bra is $\langle\psi|$. The Greek letter is a name, not a new physical property. Calling a tested state ϕ gives the overlap notation $\langle\phi|\psi\rangle$. It is read as the amplitude for that specified state test. Its squared magnitude is the probability of the corresponding ideal yes-or-no projection measurement on a normalised pure input. Mixed states and more general measurements require their extended rules.

A device placed before the test changes the input state. The column must first be updated to include what that device does. The next question makes those instructions into a table.

Try 25. What is wrong with using the row $(0.6, 0.8i)$ as the bra of the displayed column? What total does that incorrect choice give when multiplied by the column?

We can now put a state's amplitudes into a column and calculate an overlap with a row. A device changes the state before it is tested. We need a repeatable set of arithmetic instructions for finding the new column from the old one.

26 What does a matrix actually do?

Before reading on. A coherent device mixes two routes without losing photons or leaving a route record elsewhere. Must each amplitude remain in its original route, or can probability move between the exits?

A matrix is a table of coefficients that tells us how to form new amplitudes from old ones. For our recombiner it has two columns, one for each input route, and two rows, one for each output. A row contains the instructions for one output amplitude: multiply each input by its coefficient, then add the products.

This is a linear operation. If we multiply every input amplitude by the same number, every output is multiplied by that number. If we add two input columns, the output is the sum of the outputs for those columns separately. That property is why a fixed table can describe coherent combinations as well as individual input routes.

Begin with the arithmetic instructions. For exit 0, add the two incoming amplitudes and multiply by approximately 0.707. For exit 1, subtract the lower amplitude from the upper one and multiply by the same factor. The exact factor is one divided by the square root of two. Writing the common factor outside the table gives

$$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

The factor outside multiplies every entry. The top row is the add instruction; the bottom row is the subtract instruction. The columns remain ordered upper route, then lower route. This is the recombiner used earlier, with the same calibrated phase convention.

Feed in amplitudes 1 and 0: a definite upper route. The two row sums are both 1, so both output amplitudes are approximately 0.707. Squaring their exact magnitudes gives one half each. Feed in amplitudes 0 and 1 instead. The top row gives a positive amplitude and the bottom row a negative one, both with that same magnitude. The route probabilities remain half and half, but the phase relation has changed.

Now feed in equal, aligned route amplitudes, each one divided by the square root of two. The top row adds them before the common factor is applied, giving output 1. The bottom row cancels them, giving output 0. This is the bright-and-dark calculation again. The matrix saves writing out the instructions for every preparation.

For a closed coherent evolution on this fixed state space, the overlaps of states are preserved. In particular, the total squared amplitude length stays fixed for every input, not merely one convenient example. Such an evolution is unitary. A proposed linear matrix must satisfy the rule for arbitrary input combinations.

Absence of absorption alone does not guarantee that the two-route description is unitary. A route marker can retain information while every photon survives. We must then include the marker in the joint evolution. Ignoring it changes the description of the photon alone.

Order also matters. Apply a half-turn phase change to an empty lower route and nothing changes. A subsequent splitter gives two aligned amplitudes. Reverse the order: split first, then change the lower phase. The amplitudes now have opposite signs, which a later recombiner can detect. Composing operations therefore requires a record of their order. This is one reason physicists use an algebra of operations; the actual devices still determine which operations belong in the account.

Try 26. Use the two rows of the displayed matrix on input amplitudes 0 and 1. What are the output amplitudes and their probabilities?

The position-momentum trade-off concerned the spread of results from repeated preparations. Before writing its formula we need a precise measure of spread and a way to distinguish an equality from a lower limit. The arithmetic below gives those symbols something to refer to.

27 What does the uncertainty formula say?

Before reading on. If an inequality gives a minimum product of two spreads, must every preparation attain that minimum?

We first need a numerical meaning for spread. Suppose a position measurement gives either 1 or 3 millimetres, equally often. The mean is 2 millimetres. The deviations from that mean are -1 and +1 millimetres. Averaging the deviations would give zero, hiding the spread we want to describe.

Instead, square each deviation, average the squares and take the square root. Here both squares are 1 square millimetre, so their average is 1 square millimetre. Its square root is 1 millimetre. That is the standard deviation, a particular measure of spread with the same units as the measured quantity. It is not the full width of the range of possible results.

If the outcomes are not equally probable, weight each squared deviation by its probability before adding. This is a probability-weighted average. Our two-position example explains the operation; it is not a model of a minimum-uncertainty wavefunction.

In Question 11, halving position spread doubled momentum spread within a chosen family of Gaussian packets. The general rule is an inequality, so it also allows states with a larger product. For one position coordinate and its corresponding momentum, it is

$$\Delta x \Delta p \geq \frac{\hbar}{2}.$$

The first factor, Δx , is the standard deviation of position. The second, Δp , is the standard deviation of momentum. Their juxtaposition means multiply them. The sign $\geq$ means greater than or equal to. The constant $\hbar$, pronounced h-bar, is the reduced Planck constant and sets

the quantum scale. Its units match position times momentum, so the two sides can be compared.

Read the formula as a restriction: if one spread is specified, the other cannot be arbitrarily small. Equality is possible for suitable Gaussian states. A general state need not sit at that minimum, and correlations can give a stronger uncertainty statement. Thus doubling one spread does not force the other to halve in an arbitrary preparation.

The distributions refer to repeated preparations in the same state, using one set of runs for position and another for momentum. A sharp position measurement followed by a momentum measurement on the same electron does not sample both original distributions without disturbing the preparation.

The bound assumes that the spreads are well defined and that position and momentum obey their usual quantum relation. It neither measures detector imperfection nor licenses a temporary energy loan. Its purpose is to limit which pairs of distributions can belong to one prepared state.

Try 27. For the equally likely positions 1 and 3 millimetres, what are the mean and standard deviation? Why would averaging the unsquared deviations fail to measure spread?

For checking

Answers to the Try questions

Try the questions before reading these answers. Compare the physical assumptions and the first step of the reasoning, not just the final number. A matching number can hide a different interpretation.

1. No. Half and half are the probabilities. One hundred independent trials can fluctuate around 50 in each route. A probability is not a quota imposed on each batch of trials.
2. The rotated arrow still has length 0.3, so its probability is 0.09. Adding it to an unchanged, oppositely directed 0.3 arrow gives zero amplitude and zero probability, provided they contribute coherently to the same complete outcome.
3. Squaring separately gives $0.25 + 0.25 = 0.5$. The coherent dark-exit calculation instead adds amplitudes first: the result is zero before squaring.
4. Each exit has probability 0.5. Direct route measurements remain half upper and half lower, so those route counts alone do not distinguish the two phase settings.
5. No. The marker's distinguishable states still exist. Covering a display does not reverse the interaction or remove the physical route correlation.
6. Half at each exit, on average. The complementary patterns appear only after sorting by the marker result; the unsorted data contain both groups.
7. No. This changes only global phase. Every relative phase and every output probability remains unchanged.
8. No. A detector result is a distinguishable record value. Copying that value is compatible with the no-cloning theorem, which forbids universal perfect copying of arbitrary unknown quantum states.

9. The position distribution can change. The species charge does not change because an electron packet is prepared more narrowly.
10. No. Reversible free evolution can broaden a packet without energy loss to an environment. Spreading and dissipation are different.
11. The momentum spread increases by a factor of three within that minimum-uncertainty family. For a general state, the uncertainty principle supplies a lower bound, not that exact equality.
12. By a factor of sixteen: the probability factor is the square of the amplitude factor, so one quarter in amplitude gives one sixteenth in probability. The assumed attenuation model is approximate and was stated for illustration.
13. No. Relative phase may change interference-sensitive measurements even when separate component probabilities stay fixed. Only a common phase is automatically irrelevant to all such predictions.
14. One half for each of the two results. Along the original preparation axis, by contrast, the up result is certain in the ideal immediate measurement.
15. Applied to the whole state, the sign is a global phase. Applied to one route only, it changes the relative phase between alternatives that can later interfere.
16. No. Their spin labels differ, so their complete one-electron states differ. The spatial orbital alone is not the complete state.
17. The question pair 11 loses because it requires different answers. The other three pairs win, giving three quarters, or 75 per cent, for uniform questions.
18. Add first: $0.3 - 0.1 = 0.2$; then square to obtain 0.04. Squaring each contribution before adding discards the interference term and gives 0.10 instead.
19. No. It defines additional physical interactions and different possible records. The calculation must describe that changed experiment.
20. No. Real pair production must obey the appropriate energy-momentum constraints. A virtual contribution does not remove them for a real final state.

- 21.** The sum is zero, so the probability is zero. Squaring separately gives one half and loses the interference term.
- 22.** It explains the suppression of locally accessible interference and the stability of distinguishable pointer records. A unique outcome in one run requires a further interpretative or dynamical account.
- 23.** Recombine the routes with a calibrated apparatus and measure the exits. The shared route ledger alone cannot decide the different exit predictions.
- 24.** $0.6^2 + 0.8^2 = 1$. Conjugation reverses the vertical component's sign but leaves its square, and hence the squared length, unchanged.
- 25.** It omits complex conjugation. The incorrect product gives $0.36 - 0.64 = -0.28$. The correct conjugate row gives $0.36 + 0.64 = 1$.
- 26.** The amplitudes are $1/\sqrt{2}$ and $-1/\sqrt{2}$. Their squared magnitudes are both one half. The minus sign matters in subsequent interference.
- 27.** The mean is 2 millimetres and the standard deviation is 1 millimetre. The unsquared deviations average to zero because they cancel, even though the results are spread apart.

For returning to

Words, questions and sources

A short working dictionary

Amplitude

A quantity combined with other contributions to a complete outcome before its squared magnitude gives a probability. Our arrows represented amplitudes.

Basis

A chosen set of alternatives used to express a state, such as upper and lower route.

Coherence

The phase relationships that allow appropriate alternatives to interfere.

Decoherence

Loss of accessible interference in a subsystem through correlations with degrees of freedom not retained in its description.

Entanglement

Joint-state structure that cannot be represented as a mixture of independent states of the component systems.

Field

A system with degrees of freedom assigned through space or spacetime. A quantum field has quantum states and specified dynamics.

Hamiltonian

The energy operator that generates time evolution in the usual quantum description.

Observable

A specified physical quantity or measurement question, represented mathematically within the theory.

Phase

The angle of an amplitude arrow. Relative phase can affect interference; a common global phase cannot.

Record

A physical state correlated with a result. Its stability, readability and possible copying depend on the apparatus.

State

The object used by the theory to predict the statistics of specified measurements. A state is not just the results already written in a ledger.

Unitary

A linear operation preserving inner products, and therefore state normalisation, for a closed system.

A route into further reading

The following references serve different purposes. Feynman's lectures give an established introduction to amplitudes, spin and identical particles. Tong's notes provide the more formal wave mechanics behind packets, barriers and stationary states. Preskill's course develops states, measurements and quantum information. These sources use more mathematics than this volume.[1, 2, 3]

For Bell tests, the Watrous notes give the game and its quantum bound; Hensen and colleagues report an experiment. Zurek reviews decoherence, while Englert treats the relation between which-way information and interference. These support the specific topics named, not a derivation of a unique interpretation of quantum mechanics.[4, 5, 6, 7, 8]

The P1 preprint and its evidence archive belong only to the closing case study. They do not serve as sources for established quantum mechanics. The preprint has not undergone external peer review.[9, 10]

References

- [1] R. P. Feynman, R. B. Leighton and M. Sands, *The Feynman Lectures on Physics*, vol. III. See especially chapter 1, chapter 4 and chapter 6. Caltech’s authorised online edition.
- [2] D. Tong, *Quantum Mechanics*, lecture notes. [Author’s course page](#).
- [3] J. Preskill, *Quantum Computation*, lecture notes for Ph/CS 219. [Caltech course and notes](#).
- [4] J. Watrous, *Quantum Computing*, course notes, example 25 (the CHSH game). [University of Waterloo notes](#).
- [5] J. Watrous, *Advanced Topics in Quantum Information Theory*, section 7.1. [University of Waterloo notes](#).
- [6] B. Hensen et al., Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres, *Nature* **526**, 682–686 (2015). [doi:10.1038/nature15759](#). [Open manuscript](#).
- [7] W. H. Zurek, Decoherence, einselection, and the quantum origins of the classical, *Reviews of Modern Physics* **75**, 715–775 (2003). [doi:10.1103/RevModPhys.75.715](#). [Open manuscript](#).
- [8] B.-G. Englert, Fringe Visibility and Which-Way Information: An Inequality, *Physical Review Letters* **77**, 2154–2157 (1996). [doi:10.1103/PhysRevLett.77.2154](#).
- [9] D. Elliman, *Gravity and fermions on a bond-centred lattice* (2026). Preprint, without external peer review. [doi:10.5281/zenodo.23458842](#).
- [10] D. Elliman, P1 reproducibility archive (2026). [doi:10.5281/zenodo.23458463](#).
- [11] A. M. Gleason, Measures on the Closed Subspaces of a Hilbert Space, *Journal of Mathematics and Mechanics* **6**, 885–893 (1957). [doi:10.1512/iumj.1957.6.56050](#). The original theorem assumes dimension at least three.
- [12] E. H. Hauge and J. A. Støvneng, Tunneling times: a critical review, *Reviews of Modern Physics* **61**, 917–936 (1989). [doi:10.1103/RevModPhys.61.917](#).