THE THIRD ASTRA VOLUME

A Story of Quantum Mechanics

What the World Allows

A project by David Jones

Written by GPT-6 (Astra)

GPT-6 (Astra) edition · 2026

A Story of Quantum Mechanics

What the World Allows

A project by David Jones

Written by GPT-6 (Astra)

GPT-6 (Astra) edition

2026

Before the story

Historical scenes follow documented evidence; hypothetical examples are identified in the text. Interpretive questions remain open where the evidence leaves them open. Sources and notes follow the story.

This edition was written by GPT-6 (Astra) in September 2026. It is a separate telling in David Jones's continuing *A Story of Quantum Mechanics* project.

The third Astra volume follows *A Story of Systems* and *A Story of Information*. The contents and source notes are linked for navigation.

Contents

The Detector That Goes Quiet 5

The Color of Heat 10

What Light Can Give 16

An Atom with No Way Down 22

A Rule for the Unseen 29

The Arithmetic of Possibility 35

The Width of a Place 42

The Shape of Matter 48

The Wall That Can Be Crossed 57

Two Questions Far Apart 63

Three Out of Four 70

Where the Interference Goes 78

The Unfinished Answer 86

A World We Can Build 93

Keeping the Possibilities Together 101

What the Darkness Holds 110

Sources and notes 115

OPENING

The Detector That Goes Quiet

Open another way through, and one counter stops.

To see how this can happen, imagine a small experiment assembled from a source of light, mirrors, two balanced beam splitters, and two detectors. Each splitter divides an incident beam equally between its outputs. We will use an ideal version first, with perfect components, so that imperfections do not conceal the question.

The source sends single photons into the apparatus. A beam splitter gives the light two possible routes. Mirrors bring those routes together at a second beam splitter, beyond which sit the detectors. Each detection registers as one event: a click, or a number added to a count.

Call the detectors A and B.

Block one route. Some photons are absorbed by the obstruction. Those that reach the second beam splitter are equally likely to register at A or B. Over many repetitions, both counters rise at approximately the same rate.

Now remove the obstruction. Both routes are open. Adjust their relative timing with the care this experiment requires.

Detector A receives the photons.

Detector B goes quiet.

B has not been switched off. Put the obstruction back, and it clicks again. Nor have the missing detections been lost from the total. With both routes open, they appear at A.

The oddity is that a route capable of contributing detections at B, when available together with the other route, can help prevent detections there.

If you imagine each photon as a little bead choosing a road, while the other road has no part in what happens, this is difficult to explain. Opening an unused road should not cancel the arrival of the bead.

Waves can cancel. A crest can meet a trough. But the light arrives in individual detection events, even when the source is arranged to send photons one at a time.

The experiment requires us to keep both facts.

Real experiments have losses, background counts, and imperfect alignment. Their dark outputs are judged through measurements rather than declarations of absolute darkness. In 1986, Philippe Grangier, Gérard Roger, and Alain Aspect reported experiments that brought the relevant behaviors together: light with a strong single-photon character at a beam splitter, and interference when the alternatives were recombined.

The clean experiment imagined here is a way into those measured facts. It is not a story about a photon deciding to surprise someone.

What makes it strange is more demanding: a rule that works extraordinarily well requires a different account of possibility from the one we use for roads and beads.

We will learn that rule. By the end, the silent detector should feel less like an inexplicable trick. You will be able to say what makes it quiet, what can make it click again, and which questions that explanation leaves unanswered.

Along the way, the same theory will give an atom somewhere to stop falling, give matter its patterns, let particles cross barriers, and place a sharp limit on what two distant systems could achieve by carrying separate local instructions.

The surprising darkness belongs to the theory that helps explain why there is a visible, durable world at all.

Quantum mechanics is often introduced as permission for anything to happen. Its actual power comes from being exacting about what can happen, with what probability, under which conditions. Sometimes it allows something a familiar picture forbids. Sometimes it forbids something that picture freely allows.

The empty counter is evidence of a prohibition.

Our story begins earlier, when people were trying to account for a much more familiar kind of light: the glow of something hot.

Sources and notes

01

The Color of Heat

A piece of metal grows hot enough to glow. First a dull red becomes visible. With further heating, the light grows brighter and spreads toward shorter wavelengths. The change is ordinary enough to be useful to someone working at a furnace.

For a physicist, it poses an exact question. At a given temperature, how much radiation should appear at each wavelength?

The question can be studied in an oven with a small opening. Radiation inside the cavity is repeatedly absorbed and emitted by its walls. The light escaping through the hole closely approaches a universal thermal spectrum, one determined by temperature rather than by every incidental feature of the surface.

The hole becomes a way to study an ideal emitter known as a blackbody.

The name can be misleading beside something luminous. A perfect absorber, brought to thermal equilibrium, is also an ideal emitter. The darkness and the glow belong to different aspects of the same physical arrangement.

By the end of the nineteenth century, physicists possessed powerful theories of heat and electromagnetic waves. They also had increasingly careful measurements of thermal radiation. The difficulty was to bring the theories and the measurements into agreement across the spectrum.

A formula could work in one range and fail in another. A new measurement at longer wavelengths could expose the failure of an account that had seemed convincing elsewhere.

This is how a revolution can approach without looking like one. There is an instrument, a curve, and a disagreement too definite to be dissolved by admiration for the theory.

In 1900, Max Planck found a formula that matched the observed spectrum. In December, he presented a statistical argument for it. In that argument, the energy associated with the material oscillators was counted in discrete elements, each proportional to the oscillator's frequency.

The constant of proportionality was what we now call Planck's constant.

A quantity that would become fundamental entered through the work of accounting for the color of heat.

Imagine trying to distribute an amount of money among many accounts. If the money can be divided without limit, the ways of distributing it differ from the ways available when it comes in fixed units.

The size of the energy elements affects the number of arrangements that can be counted, and therefore the statistical prediction.

Planck used a larger energy element when counting arrangements among higher-frequency oscillators. The resulting distribution did not pour unlimited energy into the shortest wavelengths.

In later classical treatments, assigning energy across endlessly many high-frequency modes produced a disastrous divergence, eventually called the ultraviolet catastrophe. That label belongs to the history after Planck's first result. It should not make us imagine him in December 1900 solving a crisis already understood in exactly the form later textbooks would give it.

He had a measured spectrum, a successful formula, and a derivation containing a step whose significance was still unfolding.

The unfamiliar unit did not announce the finished theory.

Planck had not provided the modern account of a photon traveling between a source and a detector. Nor had he shown that every measurable quantity in nature comes in separate allowed values. Quantum physics will permit continuous ranges in many circumstances and discrete ones in others.

What had appeared was a scale at which the familiar freedom to divide and distribute energy required reconsideration.

At everyday scales, a very small unit can be easy to miss. A beach appears continuous from far away; closer inspection reveals grains. Yet even that image must be handled carefully here. Quantum theory is not simply the discovery that the world is built from smaller hard lumps. The grains analogy captures discreteness and leaves out much of what is coming.

Planck's constant will govern something deeper than the size of a packet. It will connect energy with frequency, momentum with wavelength, and the spreads of measurements that cannot both be made arbitrarily narrow in the same preparation.

One small constant will keep returning in apparently different questions.

The thermal spectrum gives us a first glimpse of why restrictions can be explanatory.

If a theory allows too many ways for a system to take up energy, its prediction can become wrong precisely because it is too permissive. A better account must identify which changes are available and how strongly they contribute under the conditions being studied.

That is a different kind of surprise from finding a hidden object. The oven contained no new colored substance waiting to be extracted. The revision concerned the rules by which ordinary matter and radiation could exchange energy.

The measurements were asking for a change in the grammar of the account.

And the change reached beyond the oven. The same relation between energy and frequency would help explain why light of one color could eject electrons from a surface while light of another color could fail, even when made more intense.

The difference would lie in what a single act of absorption could give.

Sources and notes

02

What Light Can Give

Light falls on a metal surface, and electrons leave it.

Change the light, and the departing electrons change their behavior. The effect gives physicists a way to ask what radiation can deliver to matter in an individual interaction.

Albert Einstein addressed the question in 1905. He proposed that, in certain processes, light behaves as if its energy is carried in localized quanta. The energy of one quantum is proportional to the light's frequency.

Planck's constant supplies the proportion.

The proposal took the relation out of the thermal calculation and gave it a more radical role. An electron absorbing one light quantum could receive a definite amount of energy set by the frequency. Some of that energy was needed to free the electron from the material. What remained could appear as its kinetic energy.

Frequency could determine the energy available in one such encounter. Intensity could influence how many encounters occurred.

The two changes need not do the same thing.

Suppose each of many people is given a ticket worth less than the price of admission. Adding more people does not increase the value of one person's ticket. Changing the ticket's value does.

This analogy isolates the difference between the number of arrivals and what an individual arrival can provide. In the photoelectric regime Einstein described, increasing the intensity of light below a material's threshold frequency does not substitute for giving one absorbed quantum enough energy to release an electron.

Above the threshold, the maximum kinetic energy of emitted electrons increases with frequency. At a fixed frequency, more intense light can release more electrons without raising that maximum energy in the way a frequency change does.

The rules become more complicated at very high intensities, where several photons can participate in an event. The simple case remains a clean

entrance into the distinction: energy per quantum and number of quanta are different physical levers.

It is possible to make a light brighter without making its individual quanta more energetic.

Einstein's proposal faced resistance, including from scientists who respected his work. The wave theory of light had explained interference and diffraction with remarkable success. A localized energy quantum seemed difficult to reconcile with those achievements.

Robert Millikan undertook careful photoelectric measurements and eventually verified the frequency dependence in Einstein's equation. His 1916 paper reported that agreement while remaining critical of the light-quantum explanation behind it.

The experimental result and the proposed picture did not become acceptable in one motion.

This is a useful historical complication. Science does not advance only when a researcher sets out to confirm the explanation that later survives. A skeptical experiment can strengthen a theory its

experimenter does not yet accept. The apparatus can return an answer that travels farther than the intention with which the measurement began.

The relation was becoming harder to avoid.

In the early 1920s, Arthur Compton studied X-rays scattered by matter. The scattered radiation could have a longer wavelength than the incoming radiation, with a shift that depended on the scattering angle.

An account treating the interaction as an exchange of energy and momentum between a light quantum and an electron explained the shift. Light's quantum character was reaching beyond the energy threshold of a surface. It was participating in the accounting of a collision.

The word *photon* would later name the quantum of light.

Yet no return to tiny classical pellets could preserve everything the experiments showed. A pellet picture could help organize the energy and momentum transfer. It could not by itself account for the interference that makes a detector go quiet when both routes are open.

The new description would have to make room for both kinds of evidence without requiring light to change identities according to a storyteller's convenience.

That demand is the beginning of a discipline.

When a detector registers a photon, the event is localized. When alternatives leading to that event are brought together coherently, the probability of detection can depend on their relationship. The localization of the outcome and the interference in its probability belong to the same experimental account.

We will build that account through what can be prepared, changed, and measured.

The stakes extend into ordinary color.

A red photon carries less energy than a blue photon. A material's possible changes help determine which light it can absorb or emit. The color of a substance, the operation of a detector, and the frequencies available from an excited atom all depend on the relations between light's energy and matter's allowed responses.

A beam can pass through one material and be absorbed by another because the materials offer different ways to take up its energy. What looks to the eye like a simple color is also an encounter between a frequency and a structure.

This gives matter a new way to be investigated. Send light in, or let excited matter send light out. Separate the wavelengths. The resulting pattern can reveal distinctions inside something too small to resolve as an ordinary object.

Long before anyone could look directly at the interior of an atom, atoms were giving out sharply patterned light.

The light offered a clue to their construction.

The construction, according to a straightforward application of familiar physics, should not have lasted.

Sources and notes

03

An Atom with No Way Down

Hydrogen, made to emit light in a discharge tube, gives a pattern of narrow spectral lines. Pass the light through an instrument that separates wavelengths, and bright marks appear at particular positions.

Other elements give other patterns.

An atom is not offering every color with equal freedom. Its emission has a structure precise enough to help identify the element at a distance. The same kinds of pattern allow astronomers to recognize material in stars they cannot visit.

The lines were known before physicists understood the atom well enough to explain them. Mathematical regularities had been found in their spacing. A successful account would have to explain why those regularities existed.

It would also have to explain why the object producing them could remain an atom.

Experiments by Hans Geiger and Ernest Marsden, interpreted by Ernest Rutherford, helped establish a nuclear picture. Positively charged alpha particles were sent toward thin material. Most continued with small deflections. A small fraction were deflected through remarkably large angles.

Rutherford's 1911 account concentrated the atom's positive charge in a very small central region. The large deflections could then be understood as close encounters with a concentrated electric influence.

The nucleus would become part of the enduring picture. The obvious planetary arrangement of electrons around it would not.

An electron circling a nucleus is an accelerating electric charge. Under classical electrodynamics, it radiates. Radiation carries away energy. A straightforward classical orbit would therefore change as it lost energy, spiraling inward instead of providing a stable, persistent atom.

Nor did that picture naturally give the sharp set of emitted frequencies.

An atom could not be saved merely by drawing the orbit neatly. The difficulty belonged to the equations governing the drawing.

In 1913, Niels Bohr proposed a model that combined Rutherford's nucleus with new rules.

The atom could occupy certain stationary states. While in such a state, it did not radiate in the manner of the classical orbiting charge. Light was associated with transitions between states, its frequency fixed by the energy difference divided by Planck's constant.

The frequency belonged to the difference between two allowed energies.

That relation gave a new meaning to a spectral line. It could be read as evidence of a particular interval in the atom's possible energy states. A collection of lines became a set of clues to an internal structure.

Bohr's model reproduced important features of hydrogen's spectrum. It connected previously observed regularities to a physical account, however incomplete. The atom could have a lowest allowed state, with no lower state into which it could radiate its remaining energy.

There was somewhere for the descent to end.

The word *allowed* is doing enormous work.

In a classical orbit, an electron can in principle occupy continuously varying energies, adjusting its motion as it loses more. Bohr restricted the possibilities. A transition could occur between permitted states; an arbitrary halfway energy was not automatically another bound state available to the atom.

Imagine a staircase with no basement. You can move down through the levels that exist. Reaching the lowest landing does not leave you suspended above another landing that you have failed to find. Within that structure, there is no further step.

A later theory must explain the structure itself, rather than merely decree it.

Bohr's proposal was a productive intermediate achievement. It gave correct results in an important domain and made the inadequacy of the old picture impossible to ignore. It also left questions about why the stationary states should exist and how the same approach could handle more complicated atoms.

The rules worked beyond the reach of their accompanying image.

That can be an uncomfortable stage in understanding.

A picture feels like an explanation partly because it lets us rehearse what is happening. We can imagine a little system turning, pulling, colliding. We can run the scene in our heads.

But a picture earns its place by matching the evidence. If an appealing scene predicts the wrong light and an unstable atom, its familiarity cannot repair the result.

The new theory would retain the task of explaining measurable outcomes while changing what had to be imagined between them. Electrons in atoms would be described by quantum states, not by miniature planetary paths taken as the full account.

Some of those states would have definite energies. Their spatial structure would matter to bonding and chemistry. Their possible transitions would explain why atoms absorb and emit particular frequencies, with further rules governing which transitions can occur and how readily.

The spectral lines were leading toward a kind of structure no ordinary orbit could supply.

The magnitude of the problem is hidden by the reliability of the answer in everyday life.

Matter stays available. A glass left on a shelf does not lose its identity because its electrons continuously radiate their way into atomic ruin. The atoms can interact, combine, absorb energy, and change state. Their persistence nevertheless requires a deeper account than a collection of classical charges following planetary tracks.

The quantum ground state will help provide that account for an atom. The stability and scale of large amounts of matter will require additional ingredients, including the way identical electrons can occupy states together.

We will reach those ingredients. For now, one decisive possibility has appeared: explaining durability may require identifying the changes that a system cannot make.

What an atom is forbidden to do helps determine what can continue to exist.

The light from the hydrogen tube remains a few narrow lines. One could list their wavelengths and stop there. One could also read them as the public consequences of a structure that has not yet been pictured correctly.

A new generation of physicists would follow that second route. They would begin with the relations among the things experiments could reveal, and ask what kind of mathematics could hold those relations together.

The task was no longer simply to find a smaller version of a familiar machine.

Sources and notes

04

A Rule for the Unseen

A spectral line concerns a change between states. That fact invites a different way to organize the atom's behavior.

Instead of beginning with an electron's imagined orbit, one can begin with quantities related to possible transitions: their frequencies and strengths, and the rules connecting one transition with another.

Werner Heisenberg took this direction in 1925. His initial paper worked through restricted problems using a new calculation built around such observable relationships. Max Born recognized the mathematical structure of matrices; Born, Pascual Jordan, and Heisenberg developed the approach further.

The strange arithmetic had an important feature. The order of operations mattered.

Multiplying the mathematical object associated with one quantity by that associated with another did not necessarily give the same result when the order was reversed.

This was no clerical inconvenience to be removed once the notation improved. It became part of the structure of the theory.

Order matters in plenty of familiar activities. Put on a sock and then a shoe; reverse the order, and the resulting arrangement changes. That example does not explain quantum mechanics. It makes room for a basic possibility that ordinary arithmetic can obscure: operations need not commute.

In quantum theory, the noncommuting relations between quantities such as position and momentum are exact mathematical statements. They help determine the limits on how sharply their measurement results can be concentrated in the same prepared state.

The consequences would require a new account of uncertainty. Before we reach it, another route into the theory is arriving.

Louis de Broglie proposed that matter, like light, should have a wave-related description. He connected a particle's momentum with a wavelength using Planck's constant.

Erwin Schrödinger developed wave mechanics in 1926. His equation described the evolution of a wavefunction, a mathematical object from which predictions about a system could be obtained. Applied to the hydrogen atom in its appropriate nonrelativistic approximation, the method yielded the energy structure that the earlier orbit model had approached through special rules.

The relation to waves offered a way to understand why only certain bound energies occur. Conditions on an allowed state restrict which solutions fit. A vibrating string gives a useful analogy: its ends constrain its possible standing-wave patterns. The allowed patterns arise from the equation and its boundaries together.

An atomic wavefunction is not a tiny string. In a many-particle system, its mathematical structure is not even generally a wave spread through ordinary three-dimensional space in the way water is. Yet the idea of solutions constrained by the physic-

al situation gives us a better starting point than an electron forbidden by decree from radiating along an ordinary orbit.

The permitted states can emerge from a general rule.

Matrix mechanics and wave mechanics looked strikingly different. Their mathematical connection showed that they could be ways of expressing the same underlying theory, within the domains where the formulations corresponded.

This was a change in what counted as continuity. The familiar picture of a mechanism could fail while a precise mathematical structure remained coherent across different representations.

But what did the wavefunction mean?

Born's work on scattering in 1926 helped establish a probabilistic interpretation. The relevant amplitudes determine probabilities through their squared magnitudes. The wavefunction could provide a distribution of possible results rather than a miniature visible material spread that simply became the observed event.

The theory could predict how outcomes would be distributed across repeated experiments while generally declining to pick the result of one trial with certainty.

This was a major shift, developed through more than a single paper or correction. Probability became part of the account of quantum measurement, not merely a temporary admission that a classical calculation was too difficult.

The wave relation for matter could also be tested directly.

In 1927, Clinton Davisson and Lester Germer reported electron scattering from a nickel crystal with a pronounced dependence on angle. The observed structure agreed with diffraction using de Broglie's wavelength. Work by George Paget Thomson and Alexander Reid provided independent electron-diffraction evidence through a different arrangement.

Electrons could exchange energy and momentum in localized interactions. Their distributions could also show the structured reinforcement and suppression associated with waves.

The crystal acted as a carefully spaced arrangement through which this structure became measurable. It brought the matter-wave relation into contact with an angle on a detector rather than leaving it as a mathematical suggestion.

The atom's theory and the laboratory's pattern were beginning to speak the same language.

We have reached the point where the opening experiment can begin to make sense.

Two routes do not simply contribute two ordinary chances that must be added. Before a probability is obtained, the theory combines something else: amplitudes.

An amplitude carries a size and a phase. The phase allows contributions to reinforce or cancel.

The distinction sounds small until it is put to work. It separates a world in which another route can only add opportunities from a world in which another route can make an outcome impossible.

The next chapter will ask you to perform that new arithmetic yourself.

Sources and notes

05

The Arithmetic of Possibility

There are two routes to detector B.

If we were calculating the arrival of a bead that simply takes one route or the other, we could add the chances associated with the alternatives. A quarter plus a quarter would be a half. The alternatives would contribute more opportunity for arrival, never less.

But the photon experiment requires a calculation in two stages.

First combine the amplitudes for the alternatives that lead coherently to the same result. Then turn the combined amplitude into a probability.

The order is essential. Turning each amplitude into a probability too early throws away precisely the relationship that makes the dark detector possible.

A convenient picture for an amplitude is an arrow.

Its length represents a size. The direction in which it points represents phase. This is a mathematical direction, not the direction in which the photon is flying through the room.

To combine two arrows, place the tail of the second at the tip of the first. The arrow from the original starting point to the final tip is their sum. Two equal arrows pointing the same way make an arrow twice as long. Two equal arrows pointing exactly opposite ways return you to the starting point: their sum has zero length.

To obtain a probability, square the length of the resulting amplitude arrow.

The square is what makes an amplitude different from a chance written in another notation. A negative number cannot be an ordinary probability. An amplitude can carry an opposing phase, and that opposition can matter before the probability is calculated.

The cancellation happens in the amplitude. The resulting probability is zero.

Now return to the ideal apparatus.

For a suitable alignment, the two contributions to detector B can be represented by equal half-length arrows pointing opposite ways. One is conveniently written as plus one-half, the other as minus one-half.

Their sum is zero. Squaring its length still gives zero.

At detector A, the two corresponding contributions point the same way. One-half plus one-half gives one. Squaring gives one again.

The probability has gone to A. Nothing needs to disappear in order for B to remain dark.

These simple numbers incorporate both balanced splitters and the phases introduced by the complete optical arrangement. Another arrangement can exchange the bright and dark outputs. The meaningful claim is about the whole experiment, not a universal preference photons have for a detector named A.

Block one arm, and only one of the two contributions remains at each detector. Its length is one-half. Squaring gives one-quarter.

There is also a one-half probability of absorption at the obstruction. The complete accounting is therefore one-half absorbed, one-quarter at A, one-quarter at B.

Over many runs with a hundred incident photons per run, the average counts approach fifty absorbed, twenty-five at A, and twenty-five at B. An individual finite run need not produce those exact numbers.

Open both arms and restore the coherent alignment, and the ideal probabilities become one at A and zero at B. Opening the second route changes the calculation from the square of one contribution to the square of a sum.

You have now performed the operation responsible for the opening surprise.

The new route did not add a negative probability. It supplied an amplitude that could oppose another amplitude.

The phase can be changed continuously.

Alter the relative optical path, and the arrows need no longer point exactly together at one output and exactly against one another at the other.

The probabilities redistribute. Both detectors may click. As the phase varies, their rates rise and fall in a coordinated pattern.

This is interference made into a controllable experiment.

The arrows are not merely a metaphor pasted onto a mysterious result after it happens. They represent a calculation that predicts how the result changes when the apparatus changes. A phase shift becomes a change in the relative lengths of the final sum arrows, which becomes a change in observed frequencies of detection.

The theory gives the experimenter something to adjust and something to expect.

An ordinary mixture behaves differently.

Suppose the preparation randomly sends a photon into one arm or the other in such a way that there is no coherence between the alternatives. We can lack knowledge of the chosen route without possessing the phase relationship needed for cancellation. That lack of knowledge does not by itself produce the dark output.

A coherent superposition and an ordinary mixture can give the same statistics for one kind of measurement and different statistics for another. Recombining the routes can reveal the difference.

This is one reason quantum states cannot generally be understood as lists of familiar objects whose hidden locations we simply have not checked. Any deeper account must reproduce the interference as well as the local detections. Some interpretations retain definite particle positions, but they do so with additional structure and dynamics, not by treating the whole apparatus as an ordinary lottery over independent roads.

The experiment constrains the explanation without forcing us to pretend that only one interpretation has ever been proposed.

There is a further change we can make. Instead of blocking a route, let the photon pass while correlating its route with another physical system. Give the two paths distinguishable records.

If those records are fully distinguishable and we look only at the output detector statistics, the interference between the path alternatives disappears. The probabilities add as they do for distinguishable alternatives.

No one has to read the record. Its physical relationship to the alternatives matters. An incomplete record can leave partial interference; an interaction that marks no difference between the paths need not spoil it.

This will become our entrance to measurement and to the everyday world. First, the amplitude rule has another consequence to show us.

It changes how a particle can be confined. Making its possible positions more concentrated imposes a cost elsewhere in the state. The mathematics that makes an output dark also belongs to the mathematics that gives an atom room to exist.

Sources and notes

06

The Width of a Place

Imagine preparing many electrons in the same quantum state.

On some, measure position along a chosen direction. On others, measure momentum along that direction: the quantity related, in the ordinary nonrelativistic case, to mass times velocity.

You obtain two distributions of results. Each has a width, measured here by a statistical quantity called its standard deviation. The widths describe variation among the outcomes of the identically prepared systems, even with ideal measurements.

Quantum mechanics places a lower bound on the product of those widths.

Make the position distribution sufficiently narrow, and the momentum distribution cannot also be arbitrarily narrow. This is the preparation uncertainty relation associated with Heisenberg and given precise mathematical forms by Kennard, Robertson, and others.

Notice what the imagined experiment did not require. No electron was measured for position and then carelessly disturbed before its momentum was checked. The two measurements were made on different members of an equally prepared collection.

The restriction concerns the state being prepared. Other properties can be sharply defined: an ideal stationary state can have a definite energy. The theory specifies which combinations are constrained.

The word *uncertainty* can make this sound like a problem with our confidence. Here it names a quantitative spread.

A precise instrument can measure a broad distribution precisely. Improving the instrument does not automatically turn that distribution into a narrow one. The apparatus's error and the state's spread are different things.

The relation is often written with Planck's constant, divided by two pi, setting the scale. We do not need to calculate its tiny numerical value to understand the consequence. There is no ordinary quantum state whose position and corresponding momentum distributions both have zero width.

The statement leaves room for extremely sharp everyday predictions because the scale is small relative to many ordinary actions. It does not make a thrown ball's path useless to calculate. It identifies a limit that becomes essential in the account of small, tightly confined systems.

A brief sound offers an analogy for the mathematics.

A sustained pure tone can have a sharply defined frequency. A very short pulse requires a spread of frequencies to describe it. The pulse's brevity and its range of frequencies are related through the way waves combine.

This is not because someone disturbed the sound by listening too roughly. It is a relation between two descriptions of the same signal.

Position and momentum wavefunctions are connected by a corresponding mathematical transformation. A narrow concentration in position requires contributions from a broader range of momentum components. The combination that gives the state a small spatial extent also prevents it from having an arbitrarily precise momentum.

Confinement therefore has consequences.

Suppose an electron is attracted toward a proton. Bringing the electron's distribution closer to the proton can lower the electric potential energy. But concentrating it into a smaller region changes the range of momenta needed to describe the state, and kinetic energy becomes costly.

There is a competition, not an unlimited advantage in squeezing forever.

This gives a useful intuition for an atom's finite extent. The full quantum calculation finds a lowest-energy hydrogen state with a definite spatial scale. The uncertainty relation alone is not a complete proof of stability or an exact calculation of the atom. It helps explain why the classical picture of inevitable inward collapse is missing an essential ingredient.

The electron's lowest state is not the state of a tiny charged planet finally sitting motionless on the proton. It is a quantum state with a spatial distribution and a nonzero kinetic-energy expectation.

The lowest energy need not mean that every familiar contribution has become zero.

This changes the meaning of cooling.

Cooling can remove thermal excitations, allowing a system to approach its ground state. It does not generally remove all quantum fluctuations or make every component into an object with perfectly sharp position and zero momentum.

A quantum harmonic oscillator retains a lowest energy above the bottom of its classical potential. Its ground state has spread in both position and momentum. The minimum is a property of the allowed quantum states, not a tiny reserve of heat that a more skillful refrigerator could drain away.

Nor is that ground-state energy a free fuel supply waiting to be extracted over and over. A ground state has no lower state of the same Hamiltonian into which the system can descend to keep paying out energy. Changing the physical arrangement changes the accounting and requires its own work or resources.

The restriction that sounds like restlessness is also a restriction on what can be taken.

The atom now has a different kind of room.

Its size is not maintained by a hard electron shell surrounding an otherwise empty miniature solar system. It emerges from a state shaped by attraction, kinetic energy, and the quantum rules.

The spatial distribution matters even though an individual position measurement gives a localized result.

Add another electron, another nucleus, then many more, and new constraints enter. The particles are not independent tenants who can be placed wherever a classical diagram has space. Their combined state must obey rules about the identities of its constituents.

Those rules help turn a collection of charges into the varied matter of a world: a gas that refuses to react easily, a crystal that carries current, a molecule that holds its atoms together.

The next restriction will help make that variety possible.

Sources and notes

07

The Shape of Matter

Press a fingertip gently against a table.

The skin changes shape. So does the table, by an amount you may not notice. Increasing the pressure requires work. Releasing it allows much of the deformation to reverse.

The resistance is an ordinary fact of your life. Its microscopic explanation must account for what happens when the atoms of two materials are pushed closer together.

Electric interactions matter. So does the structure of the allowed electronic states. Changing the arrangement changes the energy; the forces resisting further compression reflect that change.

The quantum rules are present in the shape your fingertip can take.

An electron is one of a class of particles called fermions. For identical fermions, the combined quantum state obeys an exchange rule: exchanging two particle labels reverses the sign of the state amplitude.

This is not a claim that two electrons physically circle one another whenever a physicist swaps their names on a page. It is a symmetry requirement on the description of identical particles.

The rule has a particularly clear consequence if we try to put two electrons into exactly the same complete one-electron state. Exchanging them would then leave that proposed arrangement unchanged. Yet the fermion rule requires the amplitude to reverse sign.

The only amplitude equal to its own negative is zero.

That attempted state is unavailable.

This is the exclusion principle associated with Wolfgang Pauli. The same mathematical capacity for a minus sign to matter, which helped us understand the quiet detector, now appears in the possibilities for matter itself.

The word *complete* matters.

An electron has an intrinsic property called spin. A measurement of the spin component along a chosen axis has two possible results. Spin is not adequately pictured as a tiny solid sphere rotating, but it is a real part of the quantum state and of the way electrons respond in experiments.

Two electrons can share the same spatial orbital when their spin states differ appropriately. The rule does not prohibit their spatial distributions from overlapping. It prohibits them from occupying the same complete one-particle state.

As electrons are added to an atom, they therefore cannot all pile into one lowest state. The occupation of different states helps determine the atom's energy structure and the ways it can interact with other atoms.

The periodic table's recurring patterns reflect these electronic arrangements together with nuclear charge and the interactions among the particles. The table is more than a list of substances discovered in separate places. It expresses recurring possibilities for arranging matter.

The restriction creates variety because adding more constituents requires new arrangements.

The distinction between one atom and a large piece of matter is important.

A quantum ground state can make a single hydrogen atom stable. Explaining why ordinary bulk matter has an energy and volume that behave sensibly as more particles are added requires more. The fermionic character of electrons is essential to that larger stability problem.

There is no simple picture of every electron as a hard bead with a private patch of forbidden space around it. Instead, the combined state, its symmetry, the electric interactions, and kinetic energy determine what is energetically possible.

As materials are compressed, their electrons cannot all reorganize into the same convenient low-energy state. The required reorganization can be energetically costly. That cost is part of why a solid resists being squeezed into a much smaller volume.

Exclusion is not an extra fundamental contact force added beside electromagnetism. It is a constraint on the states through which the energy and forces of electronic matter must be calculated.

The distinction is worth keeping because the actual account is more powerful than a new version of miniature billiard balls. It explains a pattern

across atoms, molecules, and bulk material through rules that do not depend on the particles having little hard surfaces.

Now bring two neutral hydrogen atoms together.

Far apart, each has its own proton and electron in an approximately separate atomic state. Closer together, the useful description must include both nuclei and both electrons as an interacting system.

A molecular state can have lower total energy than the separated atoms. If a process allows the excess energy to be carried away, the atoms can form a bound hydrogen molecule. Pulling the molecule apart again requires energy.

Walter Heitler and Fritz London gave a quantum account of this bond in 1927. Their calculation made the combined electronic state central. The electrons could no longer be treated as individually owned possessions that remained attached to their original atoms throughout the description.

The state extends across the pair of nuclei. Electric attraction, repulsion, kinetic energy, and the exchange properties of the electrons all enter. The lower-energy arrangement belongs to the molecule as a whole.

That last point is the opening into chemistry. The parts can have possibilities together that are not captured by listing what each isolated part was doing before the encounter.

The bond has a preferred range of separations.

Bring the nuclei too close, and the energy rises sharply. Pull them too far apart, and the binding advantage is lost. Around a stable separation, the molecule can vibrate; its rotational and vibrational motions acquire their own quantum energy structure.

The result is an object with ways to hold together and ways to respond.

A molecule can absorb particular radiation, exchange energy in a collision, or participate in a reaction that rearranges its bonds. The energy balance of the entire process matters. Breaking a bond by itself takes energy; a reaction can release energy when forming new arrangements more than compensates for the cost of disrupting old ones.

The language of bonds becomes a practical guide to a much larger calculation. It helps people reason about structures whose full quantum description grows difficult very quickly.

Quantum mechanics supplies principles. Chemistry develops ways to use them among the enormous number of possibilities those principles permit.

Add more atoms, and geometry becomes consequential.

An arrangement can make one direction easier to bend than another. It can permit charge to move readily or make that movement difficult. It can hold neighboring atoms in a rigid network or allow parts of a molecule to turn relative to one another.

The macroscopic differences among materials are not simply differences in how many identical bits of solidity were packed into a shape. They arise from which constituents are present, how they are arranged, which states are occupied, and what changes are available at the relevant temperature and pressure.

This is why the very small can explain the variety of the very familiar. The same general rules allow a world containing both the table and the softer fingertip pressing against it.

In a living body, chemical structures are also continually made, altered, and used. That further organization belongs to the story of systems. Here we can see part of the condition that makes it possible: matter has reliable ways to bind and specific ways to change, rather than either collapsing without structure or refusing all interaction.

Release the pressure.

The fingertip recovers its shape through the organized behavior of its materials. You do not need to calculate the electron states for them to matter to the movement. The theory describes conditions your body has always depended on.

At the beginning of the book, a forbidden detection looked like an exotic surprise. Here, unavailable states help explain ordinary persistence. A world with fewer arbitrary permissions can have more durable forms.

Yet the quantum rules can also allow a change that a classical energy barrier appears to forbid.

Some things remain because there is no lower available state. Others are only temporarily confined, and can escape by a route the old mechanics never supplied.

08

The Wall That Can Be Crossed

An unstable nucleus emits an alpha particle.

The particle leaves with a measurable energy. The nucleus has changed. Repeat the observation across a large population of such nuclei, and the rate of decay becomes a characteristic that can be measured too.

In the classical picture, there is a difficulty. The alpha particle's energy can be below the height of the potential barrier separating the interior from the outside. A classical particle with that energy cannot simply climb over the barrier.

Yet alpha particles emerge.

In 1928, George Gamow, and independently Ronald Gurney and Edward Condon, developed explanations using the new wave mechanics. The quantum state need not stop abruptly at the classical boundary. It can extend into the barrier and connect with an escaping component beyond it.

The resulting escape probability can be small without being zero.

The barrier is an energy landscape, not a wall made of tiny bricks.

For a classical comparison, imagine a ball in a valley separated from the outside by a hill. If it has too little energy to reach the top, it remains in the valley. The landscape tells us where the ball can go with the energy it possesses.

The quantum calculation works with a state that can extend across regions of that landscape. Within a barrier region, its amplitude can decay rapidly with distance. If the barrier has finite width, an amplitude can remain on the far side.

That is the basis of tunneling.

The word can tempt us to imagine a concealed passage bored through the hill. No such passage is required. Nor does the particle borrow energy for a moment and repay it before anyone can catch the violation. The complete process respects energy conservation.

What changes is the rule connecting the energy landscape with the possible outcomes.

A small change in the barrier can have an enormous effect on escape.

An amplitude that decays through a barrier can become much smaller when the barrier is made a little wider. Changing the particle's energy also changes the region it must penetrate and how strongly its amplitude falls there.

This sensitivity helps explain a striking feature of alpha decay: nuclei whose emitted alpha particles have somewhat different energies can have vastly different lifetimes. The relationship had been observed before the quantum explanation made it intelligible.

The barrier is acting as an exceptionally sensitive control on a probability.

A half-life does not give each nucleus an appointment. It describes the behavior of a population. For ordinary radioactive decay under the relevant conditions, the theory supplies a rate, and repeated observations yield the corresponding statistical pattern. It does not generally tell us which member of the population will be the next to emit.

The regularity resides in the distribution of events.

This does not mean every ground state must eventually leak away.

The unstable nuclear configuration has a lower-energy decay channel available to it. Its confinement is temporary because a transition to a different state of the complete system is possible. An isolated hydrogen atom already in its electronic ground state is a different situation: there is no lower electronic state of that atom waiting beyond a barrier.

Quantum theory distinguishes stability from long residence. A system can be at its lowest available energy, or it can be caught in a configuration whose escape is merely improbable.

Those are different explanations for why something has not changed yet.

Inside the Sun, barrier penetration helps particles approach one another rather than escape from one another.

Protons repel electrically. At solar-core temperatures, a classical account of overcoming that repulsion gives far too little reaction to explain the observed power. The thermal distribution does

contain unusually energetic particles; the difficulty is quantitative, not a literal prohibition on every imaginable collision.

Quantum tunneling makes close encounters possible at energies below the classical barrier. But proximity is only part of the first crucial reaction.

Two protons must form a deuteron, a nucleus containing a proton and a neutron, while emitting a positron and an electron neutrino. A weak interaction changes one proton into a neutron in this process. That conversion is rare.

Crossing the electric obstacle does not guarantee the nuclear success.

The rarity is consequential on a stellar scale.

The first step of the dominant solar fusion chain is slow. Later reactions can continue toward helium, with the overall process releasing energy associated with the difference in mass between the initial and final products.

Tunneling helps make the chain accessible. The weak interaction helps regulate how infrequently its difficult beginning succeeds. Together with gravity, pressure, temperature, and the star's structure, these rates help determine how the Sun shines over long periods.

The light is supported by a process that is possible but reluctant.

Make the relevant reactions much easier, and the history of a star with a given structure and fuel supply would change. The immense duration available to life on Earth is bound to the particular rates and balances of the star it orbits.

A detector can register the alpha particle. A telescope can measure the Sun's light. The particular events become evidence for the theory's probabilities and energy balances.

So far, much of our attention has been on what one prepared system can do. The next question concerns what happens when a system is prepared jointly and its parts are carried far apart.

Can everything important about the pair be packed into a separate local account for each part?

The answer will put pressure on an assumption more deeply embedded in everyday reasoning than the idea that a wall cannot be crossed.

Sources and notes

09

Two Questions Far Apart

Two systems are prepared together. Then they are taken to distant measuring stations.

For a particularly simple example, consider two spin-one-half particles in a state called a singlet. At each station, an apparatus measures the particle's spin component along a chosen axis. There are two possible results.

Use the same axis at both stations, and the singlet predicts opposite results. Either station by itself receives an evenly divided distribution over many trials. Compare their records, and the opposition becomes visible.

The pair has a precise relationship that neither station's isolated list reveals.

At first, that can seem unremarkable.

Put a red card in one envelope and a blue card in another. Shuffle the envelopes and send them apart. Whoever opens one can infer the color in the other without any message crossing between them at that instant.

The colors were settled before the envelopes separated. The later discovery changes the observer's knowledge. It does not have to change the distant card.

A correlation, even a perfect one, is therefore not enough by itself to establish the surprising character of quantum entanglement.

The difficulty appears when the stations can choose among different measurements.

The singlet is not merely a pair whose results are opposite along one preselected axis. The quantum prediction gives opposition for any common axis the stations choose, together with specific statistical relations when they choose different axes.

Can those several possible questions be answered by a complete set of local instructions carried out from the source?

Albert Einstein, Boris Podolsky, and Nathan Rosen brought a related problem into sharp form in 1935. Their original argument used position and momentum. The spin version, developed later by David Bohm and used by John Bell, offers a simpler route through the issue.

EPR asked whether the quantum description of reality was complete.

Their reasoning made locality consequential. If a measurement here permits a prediction about a distant system without physically disturbing that system, perhaps the predicted property should belong to the distant system already. Yet the quantum state does not assign simultaneous sharp values to all the quantities they considered.

For Einstein and his collaborators, that was a reason to seek a more complete account. They were not merely complaining that a theory used probabilities. Statistical descriptions were familiar and useful. The question was whether the quantum state exhausted what could be said about the physical situation.

A local account of an independently existing world seemed worth defending.

The defense deserves to be understood before it is tested.

Imagine giving each particle a very elaborate instruction manual before the pair separates. The manual can tell the local apparatus how to respond to every measurement setting it might encounter. The two manuals can be coordinated in any way at their common preparation.

Once separated, each side uses its own setting and its own instructions. It does not need to learn the other side's current choice in order to answer.

Perhaps the apparent randomness comes from our ignorance of which coordinated manuals were issued on a particular trial. Perhaps a deeper theory could describe that issuing process and recover the observed quantum statistics.

This possibility cannot be dismissed just by noticing that the outcomes are unpredictable to us. The hidden instructions were introduced precisely to explain why our incomplete description would look random.

The challenge is to find a prediction that separates the two accounts.

In 1964, Bell found a way to make the distinction mathematical.

Under a specified set of assumptions, local descriptions impose limits on the correlations between separated measurements. Quantum mechanics predicts that some prepared states and choices of measurements can exceed those limits.

The assumptions matter. The local account cannot allow the response at one station to depend on the other station's current setting through an unacknowledged influence. The choices of settings must also be suitably independent of the hidden variables supplied with the pair.

Given those conditions, the issue is no longer only whether one picture feels more reasonable than another. There is a quantitative difference to seek.

The measurements can ask a question that the philosophical argument alone could not settle.

The word *entanglement* names a feature of the joint quantum state. In an entangled state, the full description cannot be replaced by a suitable ordinary mixture of separate states for the parts while preserving all its predictions.

That is not a thread stretched between the particles. It is a property of how the state describes the whole and its parts. Its consequences appear in the relations among measurement results.

Those consequences do not give either station a controllable message from the other. In the singlet example, each local stream remains evenly divided, whatever axis is chosen at the distant station. The correlations are established by comparing records, which requires communication through an ordinary channel.

Strong joint structure and local inability to signal can coexist.

Bell's argument can be developed into a game simple enough to understand without advanced mathematics.

Two players agree on a plan. They separate. Each receives one of two questions and gives one of two answers. A referee compares their answers according to a short rule.

With local plans and independent questions, there is a ceiling on how often they can win.

You can derive that ceiling yourself. Then the quantum prediction will cross it.

Sources and notes

10

Three Out of Four

The players have as much time as they like to prepare together.

They may discuss strategies, share notebooks, arrange complicated plans, and agree how to use shared random numbers. Then they separate. During each round they cannot exchange the questions they have received or communicate before answering.

Each player receives a question labeled either 0 or 1. The questions are chosen independently and with equal probability, so the four possible pairs—00, 01, 10, and 11—are equally likely.

Each player must return an answer, also labeled 0 or 1.

The winning rule is short: their answers must agree unless both questions were 1. In that case, their answers must differ.

The labels could as easily be two colors. The difficulty is in the pattern of required agreement.

Try to make a plan that always wins.

For the question pair 00, the left player's answer to 0 must match the right player's answer to 0.

For 01, that same left answer to 0 must also match the right player's answer to 1. So the right player's two planned answers must match each other.

For 10, the left player's answer to 1 must match the right player's answer to 0. Now all four planned answers must be the same.

But for 11, the rule requires the left answer to 1 and the right answer to 1 to differ.

The plan cannot satisfy all four requirements.

You can win three of them easily: always answer 0. The answers agree, so the plan wins 00, 01, and 10, then loses 11. Since the question pairs are equally likely, this wins three-quarters of the time on average.

A different plan can move the failure to a different question pair. It cannot eliminate every failure.

Could randomness help?

The players can choose a different coordinated plan each round using shared random information. But every plan in the mixture still fails on at least one of the four equally likely question pairs.

Averaging plans that each win at most three-quarters of the time cannot lift the average above three-quarters.

More elaborate local randomness can be included in the same reasoning. The limit is not a punishment for insufficiently clever players or inadequate notebooks.

It follows from the structure of a local strategy, the independent question choices, and the winning rule.

If the questions were secretly correlated with the plans, the argument would change. If one player could learn the other's current question before answering, it would change again. The conditions are part of the theorem.

Under those conditions, seventy-five percent is the ceiling.

Now give the players an appropriate entangled pair and allow each question to select a local quantum measurement.

With suitable measurement orientations and answer conventions, quantum mechanics predicts a success probability of approximately 85.36 percent.

The state and conventions can be chosen for this agreement game; they need not use the unchanged same-axis labeling of the opposite-result singlet example in the previous chapter.

The prediction still does not promise a win on every round. It gives a higher long-run success rate than any allowed local plan can achieve. The gain is about ten percentage points, which can sound modest until its meaning is clear.

No quantity of extra preparation, hidden local detail, or shared randomness of the permitted kind can purchase those points.

The difference is between something difficult and something unavailable within the assumed class of explanations.

This is the CHSH game, a formulation of the correlation test associated with John Clauser, Michael Horne, Abner Shimony, and Richard Holt's 1969 work, following Bell's theorem.

The name is less important than the argument you have just followed. Four equally likely combinations, three demands for agreement, one demand for disagreement. Local answers cannot satisfy the whole pattern. Quantum measurements can outperform the best compromise.

In a laboratory, tests are expressed through measured correlations or related statistical quantities. Practical experiments have detector losses, imperfect states, timing constraints, and finite samples. One must show that these features have not made an ordinary local account appear to exceed its limit.

If a detector misses many events and the retained events form a biased sample, a misleading violation can result. If the stations can exchange relevant information before their outcomes are registered, the local-strategy condition has not been enforced.

The difficulty of the experiment lies in closing these opportunities while collecting enough reliable data.

In 2015, a team at Delft reported a Bell test using electron spins associated with defects in diamonds separated by about 1.3 kilometers.

The distance gave a physical role to the speed of light. Rapid choices and measurements could be arranged so that a light-speed signal carrying the current setting from one station could not reach the other in time to determine its recorded out-

come. Efficient readout and the method of heralding the entangled pairs addressed the problem of missing detections.

The reported correlations violated the relevant local bound. Independent experiments that year used entangled photons and different apparatus, also addressing the major detection and locality vulnerabilities together.

These were not literal measurements of the exact ideal game score above. They were carefully designed tests of closely related inequalities, with their own data and statistical analyses.

Their force comes from the disagreement between measured results and the allowed range of the tested local models, after controlling the identified experimental weaknesses.

The familiar phrase *loophole-free* refers to that experimental achievement. It does not mean that an experiment can operate without assumptions of any kind.

The relation between measurement choices and any proposed hidden variables still matters. So do the statistical procedures used to interpret finite

data. A serious alternative account can challenge an assumption, but it must say what changes and how it reproduces the results.

At one station, changing the distant measurement setting does not change the local outcome distribution. Each player sees a list of results whose special correlation with the other list becomes apparent only when the records are compared. That comparison requires ordinary communication. The violation supplies no way to send a controllable message faster than light.

The experiments reveal the importance of the whole quantum state. When parts become correlated, their behavior cannot always be understood by discarding the relationships and keeping only separate accounts of the parts.

That fact is about to return in a setting much closer than two stations a kilometer apart.

A particle inside an apparatus is rarely alone. Light, air, internal motion, and the apparatus itself can become correlated with it. To understand what

looks ordinary around us, we must ask what happens when the quantum relationships spread into those surroundings.

Sources and notes

11

Where the Interference Goes

Put a marker on the route without blocking it.

In our ideal interferometer, arrange for one path to become associated with one marker state, and the other path with a distinguishable state. The marker might be another degree of freedom of the light itself, or a separate physical system. What matters is that the two alternatives now leave different possible records.

The unconditioned detector counts lose the cancellation that kept B dark when the paths were coherent and unmarked.

The marker need not absorb the photon. No person has to inspect it. The physical arrangement has changed from two alternatives ending in the same complete kind of event to alternatives accompanied by distinguishable records.

The path system can no longer be treated as though those records were absent.

This has been tested in forms that go beyond the image of a measuring device knocking a particle off course.

In an atom-interferometer experiment published in 1998, Stefan Dürr, Thomas Nonn, and Gerhard Rempe used internal atomic states to hold information about the paths. The mechanical effect on the atoms' motion was too small to account for the loss of the interference pattern through the simple picture of random momentum kicks.

The internal state supplied a distinguishable marker.

The explanation remained physical. A microwave interaction changed relationships among degrees of freedom of the apparatus and atoms. Nothing immaterial called knowledge reached into the experiment. Yet describing only a rough collision missed the feature that mattered.

The question was what had become correlated with what.

In the quantum account, the joint state can retain both alternatives and their correlations with the marker. But measurements confined to the original path system no longer display the interference they would have displayed without the marking.

The distinction between the joint state and the part is essential.

It is possible for interference to become inaccessible in observations of a subsystem without the complete state having undergone a fundamental collapse to one alternative. The calculation of the part leaves out the degrees of freedom we are not controlling. Their correlations still matter to what the part can show.

If the marker states are only partly distinguishable, some interference can remain. If an interaction affects both alternatives in the same way, it may mark no difference at all. Interferometers depend on coherent interactions; they would not work if every physical encounter destroyed the relevant phase relations.

Under sufficiently controlled conditions, correlations can sometimes be manipulated to restore interference. The practical difficulty grows when the records spread into many uncontrolled surroundings.

A large object offers its surroundings innumerable opportunities to acquire such records.

Light scatters from its surface. Gas molecules collide with it. It exchanges thermal radiation. Its internal motions can become correlated with the variable being studied. If different alternatives lead to distinguishable states of these surroundings, interference between those alternatives becomes harder to reveal by manipulating the object alone.

This is decoherence.

The word names a process with rates and conditions that can be calculated and tested. There is no universal object size at which quantum mechanics turns off. The behavior depends on which alternatives are involved and how strongly their difference reaches into the environment.

A system may preserve one useful quantum relationship while other degrees of freedom interact strongly with the world. A large apparatus can therefore contain a carefully protected variable whose behavior remains distinctly quantum.

The border is made of physical relationships, not a line painted between the microscopic and macroscopic.

In 2023, researchers at ETH Zurich reported preparing a mechanical mode of a crystal in a superposition of opposite-phase oscillations. The mode had an effective mass of roughly sixteen micrograms. A superconducting circuit helped create and examine the state.

The evidence came through controlled preparation and measurements of the quantum state. It did not consist of watching a grain-sized object visibly appear in two places with the unaided eye. The separation between the relevant oscillatory alternatives was extraordinarily small on everyday scales.

Even so, the result brought a great many atoms into a collective quantum behavior that could be distinguished from an ordinary statistical mixture.

The experiment changed the size and kind of physical system over which that distinction could be maintained and tested.

It also made the care of the apparatus part of the result. Isolation, cooling, preparation, and readout were not inconveniences wrapped around a pure idea. They were how the chosen quantum relationship became accessible.

For ordinary objects in ordinary surroundings, the corresponding relationships are usually not available for us to recombine.

The environment has become involved too extensively. Recovering interference between macroscopically different records would require control far beyond what we possess in everyday life. A record on a meter can remain stable precisely because it is robustly connected with many physical consequences around it.

Decoherence theory helps explain which kinds of state remain useful as stable records under particular interactions. A record must survive the process of being encountered again. The coupling that suppresses some superpositions can favor descriptions in terms of other, more robust alternatives.

Together with dynamics and appropriate scales of description, this helps explain why classical trajectories and familiar objects work so well in their domains.

The quantum theory can contribute to the explanation of its own ordinary-looking limits.

There is an enormous achievement here, and a remaining question.

The achievement is explaining why interference between certain alternatives becomes unavailable in the measurements we can ordinarily perform, and why some records are stable. The remaining question concerns the particular outcome: how should we understand the fact that we encounter this record rather than another when the complete quantum description still contains correlated alternatives?

Simply discarding the inaccessible parts from a calculation does not, by itself, select one term in the full state.

Different interpretations or proposed modifications of the theory answer that question differently. Decoherence supplies essential structure to several of those accounts. It does not make their remaining differences disappear.

We can now see why the word *observer* need not summon a person.

An apparatus can register outcomes and store them while no one is watching. Reading its records tomorrow is another physical encounter. Ordinary quantum predictions do not require adding a special consciousness-triggered operation to the detector's mechanism.

Yet understanding the mechanism does not force us to declare every foundational question answered.

We know how to calculate the output probabilities. We know how distinguishable path records change the interference. We can model how the apparatus becomes correlated with its surroundings. The question of what the complete description means has become more precise through that knowledge.

In 1935, Schrödinger gave the unresolved difficulty an image that proved almost impossible for popular culture to leave alone.

The image was a cat.

Sources and notes

12

The Unfinished Answer

No cat had to enter a laboratory box for the argument to work.

Erwin Schrödinger proposed the situation in 1935 as a thought experiment. A radioactive event would trigger a mechanism with a large, unmistakable consequence: a cat would be alive or dead. If the complete arrangement were described by ordinary quantum evolution, the atomic alternatives would become correlated with those very different conditions.

Schrödinger intended the example to expose a difficulty in applying the description across scales. He was not reporting the discovery of an animal observed halfway between life and death.

The cat made an abstract question hard to evade: when a microscopic alternative becomes a macroscopic record, what exactly has happened to the alternatives in the quantum state?

A detector and a lamp could carry the argument with less emotional force. One lamp is on, another off; the microscopic event controls which record is made.

If the apparatus is treated as quantum, its state becomes correlated with the system it measures. Include a memory device, and the record becomes part of the joint account. Include the environment, and decoherence makes interference among the large-scale alternatives extraordinarily inaccessible under ordinary conditions.

The mathematics can describe those correlations. The interpretive question is what they imply about the outcomes we experience.

There are several ways to try to complete the account. Their differences concern which elements of the description are fundamental and whether the dynamics must change.

These are questions with consequences, even when particular approaches agree on the familiar laboratory statistics.

David Bohm's 1952 proposal adds definite particle configurations to the wavefunction, together with a law for their motion.

In this account, a particle has a position, but its behavior depends on the quantum state. In an interferometer, keeping a definite particle path does not make the unused arm irrelevant to the complete dynamics. The accompanying state still carries the structure needed to reproduce interference.

For entangled systems, the guidance involves the joint configuration in a nonlocal way. The familiar image of separate objects carrying independent local plans has not been restored.

Under the appropriate equilibrium conditions, Bohm's account reproduces the standard quantum statistics. It offers a definite configuration at the cost of accepting additional structure and a kind of dependence that ordinary local particle mechanics does not contain.

Its existence matters to the way the experiments are described. Interference does not entitle us to announce, as a neutral experimental fact, that no particle could have had a position between measurements. It requires any position-based account to do considerably more than the classical picture did.

Hugh Everett's 1957 proposal takes a different direction. Keep the quantum evolution for the whole arrangement, including whatever registers the result, and do not add a fundamental collapse.

The joint state contains different correlated records. A description of the measured system can be given relative to a particular record within that structure. Later many-worlds interpretations develop this family of ideas in terms of branching histories.

In that account, a conscious choice is not a special machine for creating universes. The relevant branching is associated with quantum interactions and the emergence of effectively distinct records. Everett's own observer models could be automatic devices with memories.

The attraction is the attempt to use one kind of quantum dynamics everywhere. The demands include explaining how the branch structure and quantum probabilities should be understood when the complete state retains the alternatives.

Decoherence makes the structure of robust, effectively separate records much clearer. Arguments about how to understand probability and the status of the branches continue within this approach.

The equations' success does not settle those arguments by itself.

Another possibility is that the usual quantum evolution needs an additional physical process.

In 1986, Gian Carlo Ghirardi, Alberto Rimini, and Tullio Weber proposed a model with spontaneous localization events. Their design made the effect rare for an individual microscopic constituent while allowing the involvement of many constituents to make large-scale superpositions rapidly localize.

This is a proposed change in dynamics. It is not simply a new way to describe learning the result of a measurement.

Such a change can have observable consequences. Depending on the model and its parameters, it can alter interference or produce other departures from standard quantum predictions. That gives experiments ways to constrain the proposal.

The price of a physical collapse is therefore an additional law that must survive testing. The advantage sought is an account in which the emergence of a particular result is itself a physical event in the fundamental theory.

These are not the only approaches.

Physicists also use operational formulations that focus on preparation, measurement, and prediction without committing to one detailed account of an underlying reality. Other interpretations put different weight on relations, information, or the status of the quantum state.

A catalogue of names would not improve our understanding unless each name helped specify what it explains and what it asks us to accept.

The useful distinction is already visible. One approach adds a configuration. Another retains the full evolving state. Another changes the evolution. Each must meet the experimental evidence, and each faces questions about the adequacy of its explanation.

Agreement on a prediction can coexist with disagreement about what the prediction describes.

The open question sits inside a field of detailed knowledge.

We can leave the cat where it belongs: inside an argument about the theory.

Outside that argument, laboratories learned to make quantum rules do practical work. A semiconductor could amplify a signal. A crystal could produce laser light. Atoms could supply a frequency reference steady enough to reveal a difference in the rate of time across a millimeter.

These achievements did not require the interpretive debate to finish first.

They required people to understand the allowed states well enough to build an apparatus around them.

Sources and notes

13

A World We Can Build

Two closely spaced gold contacts touch a piece of germanium. A plastic wedge holds them in place.

The object does not look like a new era. In December 1947, John Bardeen and Walter Brattain used such an arrangement to obtain amplification: a small change at one contact could control a larger electrical response through the powered circuit.

The point-contact transistor turned a semiconductor's behavior into a usable way to govern current.

The amplification drew energy from the circuit's supply. The incoming signal controlled how that energy was used. Making the control work depended on understanding and managing the material, including effects at its surface.

The theory of electronic states was becoming a way to build.

Begin with atoms far apart. Their electrons have allowed states associated with the separate atoms. Bring a great many atoms into a crystal, and the interactions create a much larger set of states extending through the material.

In a useful simple account, the allowed electron energies collect into bands. Within a band, many levels lie closely together. Between some bands are gaps: intervals of energy without the corresponding available states.

A band gap is not a crack running through the crystal. It is a feature of what electronic energies the material permits.

The occupation of the bands matters as much as their existence. A completely filled band does not provide ordinary electrical conduction under a small applied field in the way a partly filled set of accessible states can. Electrons need available states into which their response can develop.

In a semiconductor, heating or illumination can promote electrons across a gap, leaving mobile vacancies in the lower band. These vacancies are described as holes. A hole is a way to account for missing electron occupancy and its motion through the material.

The crystal's ability to carry current depends on these possibilities.

Carefully chosen impurities can change them.

A small concentration of suitable donor atoms can supply electrons that enter mobile states readily. Other dopants can create holes. Joining and shaping differently prepared regions gives engineers ways to arrange where charge moves, where it is depleted, and how applied voltages alter the situation.

The impurity is useful because its difference from the surrounding atoms is understood.

That is a remarkable change in the meaning of purity. A material does not have to be made as undifferentiated as possible to become useful. It can require an extremely controlled departure from uniformity: this kind of atom, in this concentration, in this region.

Later transistor designs used different structures and mechanisms from the original gold-contact device. The enduring opening was the ability to make a material's electronic response controllable enough to process signals.

A physical pattern could determine what another physical pattern would do next.

In another apparatus, a powerful flash lamp surrounds a synthetic ruby rod.

The lamp pumps energy into the material. Chromium ions in the ruby can occupy excited states and emit light associated with transitions between energy levels. With suitable populations and optical feedback, the apparatus can amplify radiation in selected modes.

Theodore Maiman and coworkers obtained the first working laser in May 1960. The output was pulsed, and the achievement had to be distinguished from the ordinary fluorescence also produced by the excited material.

The ruby was doing more than glowing red.

Einstein had identified stimulated emission in his earlier work on radiation. An existing radiation field can increase the likelihood that excited matter emits into the corresponding mode. Pumping supplies the energy; the field and resonator help organize how it is released.

An optical mode is a particular pattern of the electromagnetic field. The resonator returns light through the amplifying material and favors patterns compatible with the apparatus. When gain overcomes losses, laser action can occur.

The resulting light can have a narrow frequency range and strong coherence, with the actual bandwidth and modes determined by the device.

The practical advance is control over a relationship that our opening experiment required: a dependable phase structure.

Coherent light can be focused, interfered, modulated, and used to probe very small changes. The quantum account of light–matter interactions has helped make a source that lets people perform more demanding experiments on light and matter.

The explanatory tool has become an experimental tool.

An atomic clock uses a different consequence of allowed energy differences.

An atom responds strongly to radiation near a particular transition frequency. A clock can interrogate the atoms, use their response to correct an oscillator, and count the oscillator's cycles.

The atoms supply the reference. The clock's oscillator supplies a controlled repetition that can be counted.

The familiar cesium definition involves exactly 9,192,631,770 cycles per second of the specified radiation. It does not mean a single atom visibly hops between levels more than nine billion times each second. The count belongs to the radiation's oscillation, calibrated through its relation to the atomic transition.

A number once associated with an inaccessible internal difference becomes a way to organize public time.

Optical clocks use much higher-frequency transitions and can compare them with extraordinary precision. Their usefulness depends on keeping shifts and disturbances understood, rather than treating every atom in every setting as a flawless clock without correction.

In 2022, researchers at JILA compared different heights within one small cloud of strontium atoms held in an optical lattice.

The cloud extended over roughly a millimeter. The upper region's frequency differed from the lower region's in the way expected from gravita-

tional redshift. The upper part accumulated time faster relative to the lower part in the experiment's comparison.

A millimeter is a distance you can represent by holding finger and thumb almost together. It is a tiny division on a ruler. Here, it was enough height for the difference in the rate of time to become measurable.

General relativity supplies the gravitational prediction. Quantum transitions supply a reference precise enough to examine it across that small distance.

The experiment joins two theories through the work of a real instrument. Its precision makes a familiar room less uniform than it first appears.

There is another step beyond using materials whose behavior depends on quantum mechanics.

We can try to preserve and manipulate selected superpositions and entanglements themselves as parts of a computation. Instead of allowing the relevant phase relations to disappear into the surroundings, build an apparatus that keeps them accessible long enough to do work.

This requires turning the problem of decoherence into a problem of engineering.

The apparatus must protect what ordinary apparatus so readily records.

Sources and notes

14

Keeping the Possibilities Together

The spectrum belonged to a circuit.

In experiments published in 1985, John Martinis, Michel Devoret, and John Clarke applied microwaves to a superconducting Josephson junction. Particular frequencies changed the probability that the circuit would escape from its zero-voltage state.

The changes revealed discrete energy levels.

These were not the electronic levels of one isolated atom. The relevant quantum variable described a collective electrical degree of freedom of a fabricated device. Circuit parameters that could be measured independently entered a quantum calculation that predicted the spectrum.

The atom's selective response to radiation had acquired a surprising relative on a laboratory circuit.

A complementary experiment examined escape from the junction's zero-voltage state as the apparatus was cooled.

At higher temperatures, thermal fluctuations helped the system escape. Lowering the temperature reduced that contribution. At sufficiently low temperatures, the escape rate stopped following the thermal trend and agreed with quantum-tunneling predictions.

The barrier belonged to an effective energy landscape for the phase difference across the junction. The result was the appearance of voltage, not the whole metal device passing through a physical wall.

The distinction preserves the size of the achievement. A collective variable of a manufactured circuit could have quantized levels and a tunneling behavior that the researchers could prepare, excite, and measure.

Quantum behavior was not confined to the naturally occurring inventory of isolated atoms.

A quantum bit, or qubit, uses a pair of distinguishable basis states as the starting point for storing and manipulating quantum information.

In an ordinary binary memory, a stored bit can be read as 0 or 1. A qubit can be prepared in superpositions of its two basis states. The relative amplitude and phase affect what later operations and measurements can produce.

The extra structure is not an extra set of ordinary answers waiting to be read out all at once. A measurement in the basic readout basis still yields one of its two results. To make the phase structure useful, an algorithm must transform the state before the final measurement.

The opening interferometer already showed a small example of such control. Two coherent alternatives were brought together so that one output was suppressed and another favored. Quantum algorithms use much more elaborate sequences of operations to organize amplitudes in ways relevant to a problem.

The hard part is arranging that organization so the measured output contains useful information.

With several qubits, the joint state has possibilities that cannot generally be described by a separate state for each one. Entanglement becomes part of the computational resource.

The mathematical description can grow rapidly as qubits are added. That growth is a reason quantum systems can be difficult to simulate on ordinary computers. It is not, by itself, proof that a quantum device will solve every task faster.

Consider multiplying two large prime numbers. Finding their product is straightforward; given only that product, finding the original factors can be far harder. Peter Shor showed how a quantum algorithm could use interference to extract a related mathematical periodicity and recover the factors efficiently. For sufficiently large inputs, the algorithm offers an enormous advantage over the best known classical factoring methods, provided a quantum computer can carry out the required operations reliably.

That is a specific achievement of an algorithm. It gives no promise of faster answers to every problem. The input, the readout, the accuracy required, and the cost of maintaining the state all matter.

A large space of amplitudes is an opportunity for an algorithm, not a substitute for one.

The surroundings continually threaten the computation's useful relationships.

Unwanted interactions can entangle the qubits with degrees of freedom outside the intended calculation. Controls can be imperfect. A stored state can relax. A phase can drift. More components create more possibilities for something to go wrong.

Simply making backup copies of an arbitrary unknown quantum state is not available as a universal remedy. The no-cloning theorem rules out a device that perfectly copies every unknown state while preserving the original in the required way.

Known states can be prepared repeatedly. Restricted sets of distinguishable states can be copied. The obstruction concerns an arbitrary quantum state, including its unknown amplitude relationships.

Protecting that structure requires a different form of redundancy.

Quantum error correction spreads a logical state across several physical systems in a carefully designed joint encoding.

Additional measurements reveal information about errors without simply reading out the encoded logical value. Think of learning whether cer-

tain relationships have changed, rather than asking each component to disclose the information the computation is trying to preserve.

The results of those checks guide a correction or a revised interpretation of subsequent measurements. Classical processors participate in the work by decoding the patterns of error information. The quantum and classical apparatus must function together.

The protection has conditions. The operations used to diagnose errors can themselves introduce errors. If the underlying error rates and correlations are unfavorable, adding more hardware can make the situation worse. Below a suitable threshold, however, larger encodings can reduce the rate of logical failure.

The possibility is striking: a system made of unreliable parts can preserve a quantum state more reliably than any one unprotected part.

In work published online in December 2024, Google Quantum AI and collaborators demonstrated surface-code memories whose logical error rate fell as the tested code sizes increased. The largest memory in the reported comparison used

101 physical qubits. Its encoded information lasted longer than in the best individual physical qubit used for comparison.

These were quantum-memory experiments. They established a piece of the work required for reliable larger computations, rather than a machine already performing every hoped-for application.

Further scaling requires reliable operations, sustained decoding, and control of errors that can strike several components together. Those are concrete problems, with measured performance to improve.

The progress is in the work of preserving a state through time.

There is a reversal here.

For a detector, the desired result is often a durable record that can spread into the world. For a quantum computation, some of the relevant distinctions must be kept from becoming uncontrolled records until the algorithm has used their coherence.

The device must interact precisely where the computation requires it and remain protected where interaction would reveal the wrong information too soon.

This does not mean isolating an entire machine from everything. The control signals, cooling systems, error checks, and readout are essential interactions. The task is to engineer which relationships are maintained and which information is allowed to escape.

The same theory that explains loss of interference helps specify the conditions under which interference can remain useful.

The experiment at the beginning required only a few optical components in an ideal arrangement. The work described here can require immense technical effort to maintain and coordinate chosen quantum relationships.

Across that difference in complexity, one practical demand remains: preserve the amplitude relations until the experiment has brought them to the point where they can change an observable result.

We are ready to return to the quiet detector.

Sources and notes

CLOSING

What the Darkness Holds

Nothing arrives at detector B.

This time, you can account for the absence.

In the ideal apparatus we have been imagining, two coherent contributions lead toward that output. Their amplitudes have equal sizes and opposing phases. Add them before turning the result into a probability, and the sum is zero.

At A, the contributions reinforce. The photons are registered there.

Open paths, localized detections, a dark output: the facts that seemed to pull in different directions now belong to one calculation.

We can change the experiment again.

Block one route, and the surviving contribution gives B a nonzero chance of registering a photon. Mark the alternatives with fully distinguishable records, and the unconditioned output loses the in-

terference. Shift the relative phase while keeping the paths coherent, and the probabilities move between the detectors.

These are different physical operations. The theory tells us how to keep their consequences distinct.

Choose a phase that makes the two outputs equally likely. Now a single detection can occur at either. The quantum state supplies the probabilities; it does not generally identify the next result with certainty. Repeat the preparation and measurement, and the pattern of outcomes can test the prediction.

The question of how to understand the particular recorded result remains where we left it, among competing accounts. The ability to calculate and test the changing distribution is already ours.

The open question has acquired an exact setting.

The clues were available in objects that did not look as though they could overturn a worldview: a hot cavity's light, a metal releasing electrons, narrow lines from a gas, radiation scattered at a meas-

ured angle. People learned to let those details outweigh pictures that had made sense of almost everything they could easily imagine.

They did not all change their minds together. They made proposals that succeeded incompletely, found different mathematics for related results, and argued about what their success meant. The theory became powerful through that work without making its makers agree on every question it raised.

You do not need to remain in a laboratory to encounter the consequences.

A tabletop offers resistance. A molecule holds together. The Sun's difficult first fusion reaction occurs often enough across its vast interior to support a steady outward flow of energy, and slowly enough to belong to a long stellar life. A circuit governs a current. An atomic frequency helps keep time.

These events differ enormously in scale and purpose. Their explanations meet in the quantum rules governing the available states and transitions.

The ordinary world has gained depth. The solidity does not disappear when we understand more of it. The color remains visible. The clock still

gives a number we can use. What changes is our sense of how much structure an ordinary event contains.

That structure was there before anyone named it.

The three volumes have approached the same world by different routes.

Systems followed the growth of organized possibilities. Information followed the ways an absent event can continue to matter. Here we have reached rules that help determine which physical arrangements can exist, how they change, and what their records can reveal.

The structures that remember and the bodies that read are made of matter whose reliable possibilities include these quantum ones. Our familiar scale is one of the places where those rules have consequences.

Return the phase to the setting at which B is dark.

Imagine the source continuing its carefully prepared sequence. The apparatus has not become grander because we understand it better. Its components remain small. Its records remain simple. One count increases; another does not.

Yet the pattern of counts can now tell a precise story. Vary the phase, mark or unmark a route, and compare the results: the changing pattern lets us test whether the alternatives remain coherent. At this setting, the silence at B is the measured consequence of their relationship.

There is a way to know more that does not require the world to become more familiar.

Sometimes it requires learning to recognize an answer in the place where nothing arrives.

Sources and notes

Sources and notes

The story moves between documented experiments, historical proposals, and imagined apparatus whose ideal conditions are identified in the text. These notes provide the evidence and explain where an analogy or a simplified calculation ends. Competing interpretations remain distinct from measured results.

The order follows the narrative rather than strict chronology. Original papers are linked where available, alongside research reviews and explanations from the institutions involved. References and the cited recent paper's correction status were checked on 8 September 2026.

00 — The Detector That Goes Quiet

Return to The Detector That Goes Quiet

The opening uses an ideal Mach–Zehnder interferometer, with two balanced beam splitters and a single incident photon. Its clean probabilities make the problem visible before introducing imperfect sources, losses, background events, and incomplete alignment. Detector names are arbitrary: adjusting the relative phase can exchange which output is bright and which is dark. Darkness at one output is not a loss of probability from the complete experiment.

The historical anchor is Philippe Grangier, Gérard Roger, and Alain Aspect's *Experimental Evidence for a Photon Anticorrelation Effect on a Beam Splitter: A New Light on Single-Photon Interferences*, *Europhysics Letters* 1, 173–179, published 15 February 1986. An

atomic cascade supplied paired emissions; detection of one photon triggered a window for observing its partner. The researchers measured strong anticorrelation at a beam splitter and, in a second arrangement, interference using the same source. Those complementary tests establish more than the observation of isolated detector clicks alone.

The source does not supply the imagined run's perfect components or exact counts. The opening condenses an experimentally grounded arrangement into a worked question. Its reference to what an unused road would do concerns the ordinary independent-path picture, rather than every conceivable theory with particle trajectories.

01 — The Color of Heat

Return to The Color of Heat

Planck's paper read on 14 December 1900, "On the Theory of the Energy Distribution Law of the Normal Spectrum," begins by referring to the successful formula he had presented several weeks earlier. Its central statistical calculation distributes energy among material resonators in a finite number of equal elements, with the element's size proportional to frequency. The chapter describes this counting prescription without attributing the complete modern account of photons or oscillator energy levels to Planck at that moment. The English translation in Dirk ter Haar's collection preserves the sequence of the argument and its numerical estimates of the constants.

The later ultraviolet divergence is distinct from the immediate experimental and theoretical setting of that December paper. A primary example of the subsequent argument appears at the beginning of Einstein's 1905 light-quantum paper: applying the

classical assumptions across all frequencies gives an unbounded total radiation energy. The money analogy concerns the counting of arrangements; it supplies no mechanical model of light.

02 — What Light Can Give

Return to What Light Can Give

Einstein's 1905 paper, "On a Heuristic Point of View Concerning the Production and Transformation of Light," develops a thermodynamic argument for light quanta before examining photoelectric emission and other applications. In the photoelectric section he subtracts the work of escape from the absorbed quantum's energy to obtain the greatest kinetic energy of the departing electrons. He also allows for energy losses before an electron reaches the surface. This is why the chapter specifies the *maximum* energy rather than assigning identical energies to all emitted electrons. The original paper with an English translation provides the argument in context.

Robert Millikan's "A Direct Photoelectric Determination of Planck's 'h'," *Physical Review* 7, 355–388 (1916), records precise support for Einstein's relation alongside continuing criticism of the light-quantum account. Caltech's repository provides the article and its opening discussion. The chapter refers to the position expressed in that 1916 publication; it does not reconstruct a moment of conversion or infer Millikan's private feelings from later recollections.

The simple intensity–frequency distinction concerns single-photon photoemission. More complicated processes appear under intense illumination: for example, Fujimoto and colleagues' 1984 experiment observed multiphoton and thermally

enhanced photoemission from tungsten using very short optical pulses. That later work supports the chapter's qualification, rather than changing what Einstein proposed in 1905.

Arthur Compton's "A Quantum Theory of the Scattering of X-rays by Light Elements", published in May 1923, relates the wavelength shift to electron recoil and the scattering angle through energy and momentum conservation. This concerns the shift of the scattered radiation, not a claim that every feature of scattering had been explained. The later terminology is documented by Gilbert N. Lewis's "The Conservation of Photons", published in December 1926; the chapter does not adopt Lewis's proposed theory.

03 — An Atom with No Way Down

Return to An Atom with No Way Down

The nuclear argument appears in Ernest Rutherford's "The Scattering of Alpha and Beta Particles by Matter and the Structure of the Atom", *Philosophical Magazine* 21, 669–688 (1911). Large deflections in the experiments of Hans Geiger and Ernest Marsden called for an intense, concentrated central electric field. The scattering calculation itself did not determine the central charge's sign. Nor did it establish that electrons actually travelled in the familiar planetary paths later used to illustrate the atom.

Niels Bohr's *On the Constitution of Atoms and Molecules* (1913), especially Part I, gives both the problem of radiative orbital contraction and his proposed stationary states. He connected emitted frequencies to differences between allowed energies and recovered the principal hydrogen spectral regularities. The nonradiating states were a new assumption, not a consequence

of classical electrodynamics. The model's success did not make it a complete account of complicated atoms or a derivation of modern quantum states.

For the later mathematical account, Elliott Lieb's "The Stability of Matter", *Reviews of Modern Physics* 48, 553–569 (1976), separates atomic stability from the stability of matter in bulk. The staircase in this chapter illustrates a lowest energy; it does not explain how the energy spectrum is calculated. The full quantum theory supplies that structure.

04 — A Rule for the Unseen

Return to A Rule for the Unseen

Heisenberg's 1925 reinterpretation paper, available in D. H. Delphenich's English translation, starts from transition-related quantities and explicitly recognizes that reversing a product can change its value. He presents restricted examples rather than claiming an immediately complete solution of atomic physics. Born and Jordan's "On Quantum Mechanics", received in September 1925, identifies matrix calculus as the tool for developing this work systematically. These papers support the distinction between a new multiplication rule and its subsequent formal development.

De Broglie's 1924 thesis is identified in the Fondation Louis de Broglie's publication catalogue. His own *Waves and Motions* (1926) provides a near-contemporary exposition of the relation between matter, momentum and phase waves. It is cited as his later explanation, not as a substitute carrying the thesis's original publication date.

The University of Vienna's archival display documents receipt of Schrödinger's first wave-mechanics paper on 27 January

1926. His collected papers in English open with the explicitly nonrelativistic, unperturbed hydrogen problem and include his comparison of wave mechanics with the Heisenberg–Born–Jordan formulation. The chapter's standing-wave analogy follows the idea that conditions on admissible solutions produce discrete energies; it does not identify an atomic wavefunction with a material string or claim that the initial calculation included every observed spectral correction.

Born's short collision paper of June 1926, translated with the original pagination, connects the scattered amplitudes to probabilities for particular outgoing directions and atomic transitions. Its correction added in proof specifies the squared magnitude. Born also distinguishes what the formalism predicts from the further philosophical question of determinism. The chapter therefore treats this paper as an important part of the development of the probability interpretation, rather than presenting the footnote as the instantaneous origin of all quantum probability.

The experimental comparison rests on Davisson and Germer's 1927 nickel-crystal study and George Paget Thomson and Alexander Reid's thin-film experiment, published that June. The former measures angular structure in electron scattering; the latter records rings around the central spot on a photographic plate. They provide different experimental routes to electron diffraction. Neither is the later demonstration of an interference pattern accumulating from separately detected, individually arriving electrons.

05 — The Arithmetic of Possibility

Return to The Arithmetic of Possibility

The amplitude rule is set out in Richard Feynman, Robert Leighton, and Matthew Sands's *The Feynman Lectures on Physics*, volume III, chapter 1, especially §1–7. Contributions to the same coherent outcome add before their squared magnitude is taken. The arrows in this chapter represent complex numbers: their directions encode phase, rather than directions of flight. Equal opposite amplitudes can cancel, although ordinary probabilities cannot be negative.

The half-length contributions in the worked example include both balanced beam splitters and the phases of the complete arrangement. Blocking one arm leaves a probability of one-quarter at each output and one-half at the absorber. With both arms coherent and correctly phased, the output probabilities become one and zero. The 50/25/25 counts describe averages over repetitions, not mandatory totals for each hundred-photon run.

John Preskill's Caltech notes on states and ensembles, especially §§2.3–2.5, develop the difference between a coherent superposition and a statistical mixture. Two preparations can agree under a path measurement and disagree when paths are recombined. Merely being ignorant of a route does not supply the coherence required for a dark output. Distinguishable records introduce a further physical correlation, developed in chapter 11.

06 — The Width of a Place

Return to The Width of a Place

H. P. Robertson's *The Uncertainty Principle*, *Physical Review* 34, 163–164 (1929), defines uncertainty statistically and derives the general commutator bound. For position and corresponding momentum, the standard deviations obey $\sigma_x\sigma_p \geq \hbar/2$. This is a restriction on distributions associated with a prepared state. It is distinct from the errors of an instrument or the disturbance produced by a particular sequence of measurements. The argument does not require measuring both quantities successively on one electron.

Elliott Lieb's *The Stability of Matter*, *Reviews of Modern Physics* 48, 553–569 (1976), explains both the hydrogen problem and the limits of an uncertainty-based shortcut. Increasing confinement eventually brings a kinetic-energy cost, but the elementary uncertainty inequality is not by itself a complete proof of atomic stability. The full quantum calculation supplies the lowest energy and spatial scale; bulk matter introduces additional questions.

Richard Fitzpatrick's University of Texas treatment of the harmonic oscillator derives its lowest energy, $\hbar\omega/2$ above the classical potential minimum, and its spatially spread ground state. Ground-state fluctuations are not residual heat. For a fixed Hamiltonian, a ground state has no lower energy state available to provide an indefinitely repeated supply of extracted energy.

07 — The Shape of Matter

Return to The Shape of Matter

Lieb's "The Stability of Matter" explains why the existence of stable atoms does not by itself settle the stability and scale of a large collection. The exclusion principle is essential to the latter problem. Its restriction concerns complete electronic states, including spin, rather than a prohibition against electrons occupying overlapping regions. Two electrons with opposite spins can share a spatial orbital.

Walter Heitler and Fritz London's 1927 paper on neutral atoms and covalent binding, *Zeitschrift für Physik* 44, 455–472, brought hydrogen bonding within quantum mechanics. What matters is the combined state of the two-electron system, not a permanently identifiable electron belonging to each nucleus. Its bonding arrangement has lower total energy than two separated ground-state hydrogen atoms. Attraction, repulsion, kinetic energy, and the symmetry of the electronic state all contribute to the calculation.

Touch is a physical interaction. Pressing ordinary materials together changes their electronic arrangements and raises the system's energy. Electromagnetic interactions and the restrictions on possible electronic states both enter the explanation. Exclusion is not a separate fundamental force, and overlapping orbitals are not forbidden. The chapter's ordinary contact example brings these facts into everyday scale; it does not assign a universal fraction of a contact force to any one mathematical contribution.

08 — The Wall That Can Be Crossed

Return to The Wall That Can Be Crossed

Ronald Gurney and Edward Condon's "Wave Mechanics and Radioactive Disintegration", *Nature* 122, 439 (1928), proposed barrier penetration as an explanation of alpha decay; George Gamow developed the explanation independently. The wave-function can extend through a barrier that the emitted particle's energy would not permit it to cross classically. The resulting escape probability is highly sensitive to the barrier and energy, helping explain the enormous range of alpha-decay half-lives. This account conserves energy. It does not require a particle to borrow energy briefly, nor does it supply a measured trajectory inside the barrier.

Hans Bethe and Charles Critchfield's "The Formation of Deuterons by Proton Combination", *Physical Review* 54, 248–254 (1938), is an early quantitative treatment of the reaction that begins the Sun's principal fusion chain. Its historical numerical solar model should not be mistaken for today's measured solar parameters.

The Department of Energy's "Proton-Proton Fusion: Powering the Sun" (2017) states the two important obstacles: electrical repulsion and the small rate of the weak interaction. Forming deuterium requires a proton to become a neutron, with a positron and neutrino also produced. Barrier penetration does not make this conversion automatic. Tunneling helps enable fusion; the chain's released energy comes from the mass-energy difference between its initial and final products. The weak reaction's rarity helps account for slow hydrogen burning.

09 — Two Questions Far Apart

Return to Two Questions Far Apart

Albert Einstein, Boris Podolsky, and Nathan Rosen asked whether the wavefunction provides a complete account in *Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?*, *Physical Review* 47, 777–780, published 15 May 1935. Their argument concerned predictions about one system obtained by measuring another, previously interacting system. They used position and momentum; they were questioning completeness, not reporting a failed experimental prediction.

The spin example is a later, clearer version of that argument. John Bell explicitly uses the Bohm–Aharonov formulation in *On the Einstein Podolsky Rosen Paradox*, *Physics* 1, 195–200 (1964). For the ideal singlet state, measurements along identically oriented axes give opposite outcomes. Different relative orientations yield a characteristic pattern of correlations. Bell shows why a locally determined account cannot reproduce the complete pattern under the assumptions of the argument. Agreement or opposition at one shared setting alone would not establish that conclusion.

Preskill's discussion of faster-than-light communication, §2.5.3, explains the accompanying limit. A distant measurement choice cannot be read from the local outcome distribution alone. Records compared afterward can reveal correlations, but that comparison needs communication. The relevant fact is that the local distribution remains unchanged; the special fair-coin statistics of a singlet are not a requirement for every entangled state.

10 — Three Out of Four

Return to Three Out of Four

John Clauser, Michael Horne, Abner Shimony, and Richard Holt developed an experimentally useful inequality in *Proposed Experiment to Test Local Hidden-Variable Theories*, *Physical Review Letters* 23, 880–884 (1969). Their original derivation explicitly assumes that the hidden-variable distribution is independent of the measurement settings. Each local response also depends on its own setting rather than the distant one. Both assumptions matter to the conclusion.

The game formulation and optimal quantum strategy are explained in IBM Quantum Learning's CHSH lesson. Independent, uniformly random questions create four equally likely pairs. Three require matching answers and the fourth different answers. A deterministic local table can satisfy at most three; averaging such tables using shared randomness cannot improve that bound. Suitable measurements on an entangled pair achieve $\cos^2(\pi/8)$, approximately 85.36%, rather than the classical maximum of 75%. These are success probabilities. A finite sample can fluctuate, so an experimental excess needs statistical analysis.

Bas Hensen and colleagues' diamond-spin experiment, published 21 October 2015, combined efficient readout with rapid settings and 1.3-kilometer separation. Marissa Giustina and colleagues' photon experiment and Lynden Shalm and colleagues' independent photon experiment were published on 16 December 2015. These closed major detection and locality loopholes together. They used differing states, inequalities, and statistical procedures; they were not three literal performances of the chapter's ideal game. "Loophole-free" describes experimental controls, not the removal of every premise, including inde-

pendent settings. None supplied a controllable faster-than-light message.

11 — Where the Interference Goes

Return to Where the Interference Goes

Stefan Dürr, Thomas Nonn, and Gerhard Rempe's 1998 atom-interferometer experiment, *Nature* 395, 33–37, used internal atomic states to mark paths. Ordinary momentum kicks were too small to explain the observed loss of interference. A physical record can distinguish alternatives without being read by a person. Partial distinguishability permits partial interference; an interaction that marks no relevant difference need not destroy coherence.

Wojciech Zurek's *Decoherence, Einselection, and the Quantum Origins of the Classical*, *Reviews of Modern Physics* 75, 715 (2003), explains how environmental interactions favor stable records. Maximilian Schlosshauer's review of decoherence and measurement distinguishes the locally accessible description from the full joint state. Suppression of interference in a subsystem's reduced state does not by itself show that one branch of the complete state has physically survived while the others have vanished. This is why decoherence is a substantial explanation without automatically settling the question of definite outcomes.

Marius Bild and colleagues reported *Schrödinger Cat States of a 16-Microgram Mechanical Oscillator*, *Science* 380, 274–278 (2023). The resonator's effective oscillating mass was 16.2 micrograms; its constituent atoms participated in a superposition of two opposite-phase oscillations. The researchers controlled the superposition and studied its decoherence. This was a carefully pre-

pared vibrational degree of freedom, rather than an ordinary animal in two biological conditions. The preprint was posted in November 2022; the journal report and ETH's explanation followed in April 2023.

12 — The Unfinished Answer

Return to The Unfinished Answer

Schrödinger's cat appears in his 1935 *Die gegenwärtige Situation in der Quantenmechanik*. John D. Trimmer's English translation was published in 1980. The original proposes a radioactive sample and apparatus connecting a possible decay to the animal's fate. It is a critical thought experiment about extending the quantum description to an entire system, not a laboratory report or evidence that an observer's glance creates a living or dead cat. The argument amplifies the difficulty of relating a joint wavefunction to familiar definite records.

David Bohm's 1952 account with hidden variables adds definite particle configurations and a law governing their motion. With the standard wave equation and the appropriate equilibrium distribution, it reproduces ordinary quantum statistics. Its nonlocal dynamics prevents it from being reduced to the independent bead-and-road picture. Having trajectories does not allow it to discard the wavefunction's contribution to interference.

Hugh Everett's *"Relative State" Formulation of Quantum Mechanics*, published in 1957, instead retains the continuous evolution of the whole state, including correlated apparatus and observer records. His models of observers can be automatic recording machines. The proposal does not require conscious decisions to produce the mathematical branching structure. Its account of

relative records is the historical starting point for the family of views later commonly called many worlds.

Ghirardi, Rimini, and Weber's *Unified Dynamics for Microscopic and Macroscopic Systems*, published in 1986, changes the dynamics through spontaneous localization. Its collective effect can be large while microscopic effects are rare. Such collapse models predict departures from ordinary unitary theory; they are not simply alternative descriptions guaranteed to agree in every experiment. The chapter presents these as distinct proposals, without treating their questions of interpretation, probability, or empirical adequacy as settled by a single familiar demonstration.

13 — A World We Can Build

Return to A World We Can Build

The Computer History Museum's account of the point-contact transistor dates Bardeen and Brattain's successful amplification experiment to 16 December 1947 and their demonstration to laboratory officials to 23 December. Two closely spaced gold contacts, supported by a plastic wedge, touched high-purity germanium. The device controlled current in a powered circuit; it did not create the amplified signal's energy. Its construction differs from that of later silicon field-effect transistors.

For the underlying electronic structure, MIT's semiconductor lecture connects molecular states to crystal energy bands, then explains band occupancy, gaps, and doping. A band gap is an interval of unavailable electron energies, not a spatial opening inside a material.

Theodore Maiman reported stimulated emission from ruby in "Stimulated Optical Radiation in Ruby", *Nature* 187, 493–494

(1960). The American Physical Society's LaserFest history supplies the earlier operating date, 16 May 1960, and the coiled flash lamp surrounding the synthetic ruby rod. Supplied light excited chromium ions; the apparatus produced pulsed laser output. The paper appeared on 6 August.

Einstein's "The Quantum Theory of Radiation" (1917) distinguishes spontaneous emission from transitions induced by incident radiation. The induced downward transition supplies the historical basis for stimulated emission. A laser also requires a maintained supply of excitation and an arrangement in which amplification can overcome losses.

NIST's "How Do You Measure a Second?" explains how atomic transitions provide a frequency reference for an oscillator whose cycles can be counted. The cesium number describes cycles of resonant radiation, not billions of individual atomic jumps each second. In Bothwell and colleagues' 2022 experiment, a frequency gradient within a millimetre-scale strontium sample agreed with gravitational redshift. Quantum precision enabled a test of general relativity; it did not establish a quantum theory of gravity.

14 — Keeping the Possibilities Together

Return to Keeping the Possibilities Together

Two 1985 papers establish the circuit experiments. Martinis, Devoret, and Clarke measured discrete energy levels through microwave excitation, *Physical Review Letters 55*, 1543–1546. Devoret, Martinis, and Clarke measured escape from a junction's zero-voltage state, *55*, 1908–1911. At sufficiently low temperatures the escape rate stopped following temperature and agreed with quantum tunneling predictions using independ-

ently measured circuit parameters. The quantum variable was the collective phase difference across the junction. The experiment did not move the whole apparatus through a physical barrier.

Peter Shor's "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer" develops efficient quantum algorithms for these specific problems. The arXiv manuscript, posted in 1995 and revised in 1996, expands his 1994 conference paper. The factoring algorithm uses quantum period finding as part of a larger procedure. Its advantage is relative to the best known classical methods; the paper does not prove that no efficient classical factoring algorithm can exist. The algorithmic result assumes a quantum computer capable of executing the required operations and is distinct from a demonstration at a practically useful scale.

Wootters and Zurek's "A Single Quantum Cannot Be Cloned", *Nature* 299, 802–803 (1982), rules out a universal device that perfectly copies an arbitrary unknown quantum state. It does not forbid preparing multiple systems in a known state. Quantum error correction therefore protects information through a suitable joint encoding and measurements of error information, rather than by making ordinary independent backup copies of the unknown state.

Google Quantum AI and collaborators' "Quantum Error Correction below the Surface Code Threshold" appeared online on 9 December 2024 and in *Nature* 638, 920–926 (2025). Larger surface-code memories had lower logical error rates over the tested sizes. The largest used 101 physical qubits and preserved information longer than its best constituent physical qubit. These were quantum-memory experiments, not a demonstration of useful general-purpose computation. Scaling, correlated

errors, decoding, and reliable operations remain consequential requirements.

The 28 April 2026 author correction repairs Figure 3a's axis and legend labels: those data concern repetition-code distance, with symbols distinguishing earlier work from this experiment. It reports no changed data or conclusions. Repetition-code results should remain distinct from the paper's surface-code memory results.

15 — What the Darkness Holds

Return to What the Darkness Holds

The final detector is the ideal apparatus introduced in the prologue and calculated in chapter 05. Returning to it adds no new experimental claim. The distinction between interference becoming inaccessible and a unique outcome being explained remains the one developed in the measurement references above.

The reader's changed understanding is the literary frame. The suggestion that understanding can deepen rather than exhaust wonder belongs to the telling, not to a theorem of quantum mechanics. What the experiment establishes is specific: preparations, coherent transformations, and distinguishable records determine different outcome statistics. The broader meaning attached to learning those restrictions is an interpretation offered to the reader. It does not depend on asserting that uncertainty grants unlimited possibility or that an unexplained feature of measurement supplies evidence for a preferred account of mind or reality.

A Story of Quantum Mechanics
What the World Allows

GPT-6 (Astra) edition
Written by GPT-6 (Astra)
A project by David Jones

Set in Libre Baskerville.
A typeface by Impallari Type,
after the work of John Baskerville.

2026