GRESHAM

COLLEGE

Reproduction of this text, or any extract from it, must credit Gresham College

CHAOS AND THE QUANTUM

A Lecture by

PROFESSOR IAN STEWART MA PhD FIMA CMath
Gresham Professor of Geometry

4 December 1996

GRESHAM COLLEGE

Policy & Objectives

An independently funded educational institution, Gresham College exists

- to continue the free public lectures which have been given for 400 years, and to reinterpret the 'new learning' of Sir Thomas Gresham's day in contemporary terms;
- to engage in study, teaching and research, particularly in those disciplines represented by the Gresham Professors;
- to foster academic consideration of contemporary problems;
- to challenge those who live or work in the City of London to engage in intellectual debate on those subjects in which the City has a proper concern; and to provide a window on the City for learned societies, both national and international.

Founded 1597

Gresham College, Barnard's Inn Hall, Holborn, London EC1N 2HH
Tel: 020 7831 0575 Fax: 020 7831 5208
e-mail: enquiries@gresham.ac.uk
Web site : www.gresham.ac.uk

Gresham Geometry Lecture
4 December 1996

CHAOS AND THE QUANTUM

The mathematican plays a game in which he himself invents the rules while the physicist plays a game in which the rules are provided by Nature, but as time goes on it becomes increasingly evident that the rules which the mathematician finds interesting are the same as those which Nature has chosen.
Paul Adrien Maurice Dirac

"I used to be uncertain, but now I'm not so sure."
Tee-shirt slogan

Albert Einstein "did not believe that God plays dice with the universe." . When he made his celebrated remark — reproduced more accurately below — he was referring to quantum mechanics. This differs in fundamental ways from the 'classical' mechanics of Newton, Laplace, and Poincaré. Einstein made his famous statement in a letter to the physicist Max Born:

> "You believe in the God who plays dice, and I in complete law and order in a world which objectively exists, and which I, in a wildly speculative way, am trying to capture. I firmly *believe*, but I hope that someone will discover a more realistic way, or rather a more tangible basis than it has been my lot to do. Even the great initial success of the quantum theory does not make me believe in the fundamental dice game, although I am well aware that your younger colleagues interpret this as a consequence of senility."

Chaos was unknown in Einstein's day, but it was the kind of concept he was seeking. The very image of chance as a rolling cube is classical, not quantum. And chaos is primarily a concept of classical mechanics. How does the discovery of chaos affect quantum mechanics, and what support — or otherwise — does it offer for Einstein's philosophy?

The vast majority of physicists see no reason to make changes to the current framework of quantum mechanics, in which quantum events have an irreducibly probabilistic character. Their view is: 'if it ain't broke, don't fix it.' However, hardly any philosophers of science are at ease with the conventional interpretation of quantum mechanics, on the grounds that it is philosophically incoherent, especially regarding the key concept of an observation. Moreover, some of the world's foremost physicists agree with the philosophers. They think that something *is* broke, and therefore *does* need fixing. It may not be necessary to tinker with quantum mechanics itself: it may be that all we need is a deeper kind of background mathematics that explains why the probabilistic point of view works, much as Einstein's concept of curved space explained Newtonian gravitation. Of course Einstein's general relativity actually went beyond Newtonian mechanics and changed the mathematics of gravitational theory as well as the interpretation, but along the way it explained the philosophically incoherent Newtonian appeal to forces acting at a distance, replacing it by the inherent curvature of space acting locally. And Newton's theory can be recovered as a very good approximation to general relativity, valid when the curvature of space is small. So maybe a new framework for quantum mechanics will accommodate the existing, highly successful probabilistic viewpoint *exactly*; and maybe it will reveal it as an *approximation* to something deeper but essentially different.

The philosophers mainly think that it is the *interpretation* of quantum mechanics that needs to be fixed. The physicists aren't terribly interested in a reinterpretation, however philosophically superior it might be, unless it yields radically new physics; but several major figures are convinced that quantum mechanics *itself* is in need of a fundamental reformulation that goes well beyond mere tinkering. They believe that despite its immense success in predicting the outcome of experiments, quantum mechanics needs to be rebuilt from the ground up. And some mathematicians, perhaps excited by the prospect of interesting new kinds of mathematics, agree.

Quantum mechanics was forced upon physicists by the results of a large number of careful experiments that demonstrated the inadequacy of Newtonian mechanics. The very existence of electrons inside atoms, for example, provides such evidence. In a classical model of the atom, electrons are point electrical charges orbiting a central nucleus made from protons and neutrons. But in classical physics, a moving electric charge must radiate some of its energy as electromagnetic waves, so that electrons could not continue to orbit the nucleus for very long. Instead, they would spiral into the nucleus and disappear, losing their electric charge in a collision with an oppositely charged proton. Atoms would fall apart and disappear. Since this doesn't happen, something is wrong. It *could* be the image of orbiting point charges, but nobody has ever managed to fix that up in a way that fits experiments. So maybe it's classical physics that is wrong. Maybe moving electrically charged particles do *not* radiate their charge away.

And a lot of experiments showed that they don't.

Waves, Particles, and Quanta

Before physicists ran up against awkward experimental evidence to the contrary, their view of the physical universe was straightforward. There was matter, which was composed of particles, and radiation, which was composed of waves. Matter possessed mass, position, and velocity. Mass could be any positive real number, and position and velocity existed in a space-time continuum, meaning that position and velocity coordinates (relative to some choice of axes) could be arbitrary real numbers, positive or negative. Or, to put it another way, space and time were infinitely divisible. So were associated quantities such as energy; and in principle you could measure all of these quantities as accurately as you wished. Waves were different. A wave possessed a frequency (number of waves per second) and an amplitude (height of wave). In particular, an electron was a particle, and light was a wave. You could tell because electrons bounced off other bits of matter like little rubber balls, whereas when two light rays met they formed 'interference fringes' in which the wave patterns became superposed (added together).

It soon became apparent, however, that there are circumstances in which light behaves like a stream of particles rather than a wave. One is the photoeletric effect, in which light impinging on a suitable substance produces an electric current. Then it transpired that electrons sometimes behave like waves: if you pass pairs of electrons through fine parallel slits you got interference fringes. So the distinction between waves and particles started to blur.

Another piece of evidence for the particle-like nature of light was contained in a theory announced by Max Planck in 1900, to the effect that the energy of an electromagnetic wave is *not* infinitely divisible. For light of a fixed frequency there is a definite minimum energy; moreover, the only possible values for the energy are whole number multiples of that minimum value. It is as if energy can only come in tiny packets of fixed size, and every light 'wave' is made up from an integral number of packets. Planck called these packets 'quanta'. A new fundamental physical constant, now known as *Planck's constant* and denoted by the letter h, captured the relationship between the frequency of the light and the energy of one quantum. In fact, the energy of a quantum of light is equal to its frequency multiplied by Planck's constant. Planck's constant is *tiny*: 6.626×10^{-34}. So although energy is not infinitely divisible, you can divide it up rather often before you get down to the irreducible lumpiness of individual quanta. Only when you get down to the thirty-fourth decimal place, so to speak, do you discover that the universe will not subdivide forever. This explains why the universe had previously looked infinitely divisible.

The distinction may seem a fine point, but it had a crucial consequence. If you assume that energy is infinitely divisible, you get the wrong value for the energy radiated by a 'black body'— a perfect radiator — at high frequencies. Indeed, you get infinity. Planck discovered that if you assume that energy comes as integer multiples of h times frequency, then you get a finite value — and one that agrees with experiment.

Wave Functions

How do we reconcile the two contrary attributes of matter, wave and particle? Erwin Schrödinger took the view that a particle was a sort of concentrated wave, which in most circumstances acted as if it was localised in a small region of space, and travelled through time as a coherent blob or 'wave-packet'. However, in certain circumstances the wave packet could become more extensive, giving the appearance of a conventional wave (**Fig.1**). But what was doing the waving? Ocean waves are waves of water, electromagnetic waves are waves in the electrical and magnetic fields. According to Schrödinger, quantum waves are waves in a — possibly multidimensional — mathematical space of complex numbers. The properties of this quantum wave were therefore defined by a *wave function*, usually denoted by the Greek letter ψ (psi). At each point (x,y,z) of space and each instant t of time the value $\psi(x,y,z,t)$ is some complex number — or complex vector in the multidimensional case.

Fig.1 A wave and a wave packet.

Schrödinger wrote down a simple differential equation that the wave function must satisfy. Schrödinger's equation, as it is now called, determines the propagation of the quantum wave function through space and time, by specifying how it changes from its current value as it moves into the future. Schrödinger's equation is linear, meaning that solutions can be superposed to give more solutions. This fits the wavelike behaviour of matter quite nicely, but at first sight it doesn't look so good for particles. However, a collection of several separate particles can reasonably be viewed as a superposition of the states that you would have got if each individual particle had existed on its own. What about interactions between particles when they come very close together? Those are precisely the circumstances in which microscopic matter stops behaving like a conventional wave or particle, and it is here that Schrodinger's equation proved its worth, predicting things like the energy levels of the hydrogen atom with exquisite accuracy.

The linearity of quantum mechanics had some curious consequences, which sound bizarre but fit experiments perfectly. For example, an electron possesses a feature known — rather misleadingly — as 'spin', because it is in a sense analogous to a spinning ball. When you measure electron spin relative to some choice of 'axis' you always get the value $+{}^1/_2$ or $-{}^1/_2$ — nothing else. You can interpret the difference as analogous to that between clockwise spin and anticlockwise spin, if you wish; bit this analogy rapidly breaks down when you superpose spin states. An electron can simultaneously have spin $+{}^1/_2$ about the north axis and spin $-{}^1/_2$ about the east axis, say.

However, you can't actually measure both spins simultaneously. The very notion of measurement poses nasty problems for quantum mechanics, even though it is central to any experimental test. Physicists know what *actually* happens when you measure a quantum object. You get a number. One number. You can make another measurement to get a second number, but you can't necessarily assume that the first number remains the same as it was. Some quantum variables are independent: if you measure one it does not affect the other. But some aren't, and among them are spins about different 'axes'.

What quantum physics lacks, however, is a clear theoretical description of what happens when a measurement is made on a quantum system. For example — speaking very roughly and ignoring some technical restrictions — you can superpose 50% of the state "east spin = $+{}^1/_2$" and 50% of the state "east spin = $-^1/_2$", to get a valid state of an electron. (Note that these spins are about the *same* axis.) But it's not an electron with spin zero; *no* electron has spin zero. Electrons always have spin $\pm^1/_2$ in any direction. Neither is it an electron spinning about some compromise axis halfway between east and north.

How do we know that this strange superposed state actually occurs? The best we can do is to prepare lots of electrons that are supposed to be in this combined state, and measure their east spins. What we find is that — apparently randomly — we get either $+^1/_2$ or $-^1/_2$. So on average we get 50% of spin $+^1/_2$ and 50% of spin $-^1/_2$. We *interpret* this as evidence for the 50/50 combined state.

It's a bit like having a theory about coins that move in space, but only being able to measure their state by interrupting them with a table. We hypothesise that the coin may be able to revolve in space, a state that is neither 'heads' nor 'tails' but a kind of mixture. Our experimental proof is that when you stick a table in, you get heads half the time and tails the other half — randomly.

Cat in a Box

It is all very curious. Why does a measurement of the (east) spin of an electron always produce either $+^1/_2$ or $-^1/_2$, even if the electron is actually in a superposition of those states? It becomes even more mysterious if you try to build the measuring apparatus into the quantum equation. A 'spinometer' — a spin-measuring device of some kind — is made out of the same sub-atomic particles as the electron whose spin it is to measure. Suppose that when the electron has spin $+^1/_2$ the apparatus is in quantum state P (for plus), whereas when the electron has spin $-^1/_2$ the apparatus is in quantum state M (for minus). Then, by the linearity of the quantum equations, when the electron is in quantum state 50% of $+^1/_2$ plus 50% of $-^1/_2$, the apparatus should be in quantum state $^1/_2P+^1/_2M$.

But it's not. It's a spinometer, so it must always be in state P — and state P alone — or state M. There is still a trace of the superposition of quantum states, however, because the apparatus seems to choose the states P and M at random — and it is in each state for about half of the measurements. The *average* state, over many experiments, is $^1/_2P+^1/_2M$. Somehow, when you work with actual measurements using macroscopic apparatus, the superposition principle ceases to function. There is a conceptual mismatch between the microscopic world of the quantum, and the macroscopic world of the spinometer. But quantum mechanics is supposed to apply to all objects, micro or macro, isn't it?

The favoured way to get round this difficulty is to introduce an interpretation of the measurement process that does not attempt to model the apparatus as a quantum system at all. Instead, it simply *accepts* the (mysterious) fact that a spinometer either yields $+^1/_2$ or $-^1/_2$, but never a mixture. And it argues that what the apparatus does is to *collapse* the wave function down into one or other of its component parts. Those parts are known as *eigenfunctions*, and the corresponding states are *eigenstates.* Those words mean something rather special in the mathematical formalism, but what they will signal to us is that there are certain 'special' wave functions (eigenfunctions) out of which all the others can be constructed by superposition, and it is those special states (eigenstates), and only those, that you can observe. The measurement process starts with a superposition of eigenstates, such as $^1/_2P+^1/_2M$, and 'collapses' it to either P or M. Which one occurs is irreducibly probabilistic. To find out about the coefficients $\pm^1/_2$, you have to repeat the experiment and calculate probabilities. This point of view on the measurement process for quantum systems is called the *Copenhagen interpretation*, and it was advanced by Niels Bohr in 1927. Although intended as a pragmatic solution to a conceptual difficulty facing practising physicists, it has led to considerable philosophical mysticism about the role of the human mind as an observer of quantum reality, and suggestions that the universe doesn't really exist unless a human being is looking at it. Personally, I think this is silly: what I want to understand is how the *spinometer* manages to avoid being in state $^1/_2P+^1/_2M$, which is the heart of the mystery. A human mind observing the spinometer is a secondary stage: the conceptual problem arises even if there are no humans in the loop at all.

Erwin Schrödinger seems to have had the same feeling, and in 1935 he tried to demonstrate what he considered to be the absurdity of the Copenhagen interpretation with his famous 'cat in a box' thought experiment. Einstein called it "the prettiest demonstration" that the Copenhagen interpretation is an incomplete representation of the real universe. Imagine a box that contains a source of radioactivity, a Geiger counter to detect the presence of radioactive particles, a bottle of (gaseous) poison, and a (live) cat. These are arranged so that if a radioactive atom decays and releases a particle, then the Geiger counter will detect it, set off some kind of machinery that crushes the bottle, and kill the unfortunate cat. From outside the box, an observer cannot determine the quantum state of the radioactive atom: it may either have decayed or not. So according to the Copenhagen interpretation the quantum state of the atom is a superposition of 'not decayed' and 'decayed' — and so is that of the cat, which is part alive and part dead at the same time. Until, that is, we open the box. At this instant the wave function of the atom instantly collapses, say to 'decayed', and that of the cat also instantly collapses to 'dead'.

It sounds ridiculous, and that's how Schrödinger intended it. It's my '$^1/_2P+^1/_2M$' neatly parcelled up in a box with a cat in the starring role. But other quantum physicists didn't think it was ridiculous at all. Quantum theory, they argued, really is strange. Maybe it's so strange that you really can have a cat that's half alive and half dead provided nobody looks at it to determine which. And for every objection they had an answer. Why not put a movie camera in the box to film the cat? Afterwards, you can develop it and see whether the cat died or not. But no: until you open the box, the film is itself in a superposed state, part a film of a living cat, part of a dead one, and only when you open the box— well, you get the gist.

In recent years, experimentalists have devised some very cunning experiments to try to find out just when the wave function of a quantum system inside an impenetrable box collapses. However, it is utterly impossible even to write down the quantum wave function for something as complicated as a cat. Indeed the helium atom, with two electrons, two protons, and two neutrons, is already too complcated. It is utterly impossible to write down the quantum wave function for a spinometer, too. Or for a Geiger counter. If you want your experiment to match up with theory, you have to replace Schrödinger's cat by a microscopic quantum system — such an electron. So this is what the experimentalists did, and then they found cunning ways to deduce what went on inside the box before it was opened, and whether opening the box changed anything. And some of them said: yes, it is a superposition that collapses only when you open the box. And others said: no, the electron-cat was dead all along, you just didn't know that until you looked. Because, even when you replace the cat by a microscopic quantum system, ultimately the observations must be interpreted for macroscopic creatures to comprehend, and the same experiment can be interpreted in different ways. The Copenhagen interpretation tells you that only observations posses physcial reality, but at the same time it tells you that a key feature of the mathematical formalism, the wave function, cannot be observed completely, and therefore is unreal.

I strongly suspect that Schrödinger thought that the wave function — all of its, not just its collapses eigenfunctions — was real. After all, he invented it and wrote down the equation for how it evolves: I'd be surprised if he thought it was merely a mathematical fiction. Definitely he introduced his cat quite deliberately, to dramatise the gulf — not understood then *or now* — between quantum microdynamics and classical macrodynamics.

Today most people use Schrödinger's thought experiment for a purpose very different from what its originator intended. They use it to show us how very weird the quantum world is. "That poor cat *really is* in a superposition of alive and dead until you open the box." What the experiments tell us is that this is true of small scale quantum systems like electrons. *We do not know that it is true of cats, and it is almost certainly false.* Not because of anything to do with consciousness, feline or human, but because a cat is a macroscopic system and 'alive; and 'dead' are macroscopic properties. Macroscopic properties do not superpose. Schrödinger was trying to tell the physics community that the problem of measurement cannot be resolved by grafting a wholly artificial collapse of wave functions on to an otherwise elegant but linear mathematical structure. Instead, it is about the nature of *macroscopic* objects built up from quantum particles. This is *why* he introduced a cat instead of, say, an electron. We inhabit a world of macroscopic objects, which obey classical mechanics much better than they do quantum mechanics.

Why do they do that?

It is believed that a phenomenon called decoherence, related to the fact that the quantum wave function has complex values but our observations must be real, which causes large collections of quantum particles to behave in a classical manner when they are observed in a classical manner. If so, then what happens to the cat is straightforward. It is not in a superposition of states. It is in just one of them — but you don't know which until you open the box. The reason has nothing to do with the cat. It is actually the Geiger counter, one of the other macroscopic systems inside the box, that relies upon decoherence to decide which state the radioactive atom is in. After that, it's classical dynamics all the way. The bottle breaks because of classical interactions triggered by the classical machinery attached to the classical Geiger counter, the classical cat dies because of classical interactions with the classical poison. *We* don't know what's happened, not until we open the box; but the classical systems inside the box *do* 'know' what's happened.

In fact, even at the quantum level, the crucial step happens at the detector, not at the opening of the box. Leonard Mandel has carried out experiments showing that a photon can be switched from wavelike behaviour to particle-like behaviour — which a Copenhagenist would consider to be a collapse of its wave function — without a human observer being aware of this collapse *at the time it happens.* In other words, before anyone opens the box, the cat (photon) is already dead (particle-like). The measurement is completed as soon as the apparatus produces a classical yes/no answer, not when a scientist looks at that apparatus.

The EPR Paradox

The problem of measurement is intimately bound up with another celebrated quantum-mechanical difficulty, which also goes back to Einstein. It provides an excellent test-bed for chaotic replacements for quantum indeterminacy. In 1935 Einstein, Boris Podolsky, and Nathan Rosen asked whether the quantum description of reality might be missing an essential ingredient. And, like Schrödinger, they convinced themselves that the answer was 'yes'. Their scenario requires two particles to interact with each other and then fly apart, interacting with nothing else. At any instant each particle is in a definite position and has a definite momentum. When they are close together we can simultaneously measure the distance between them (to be sure they really are close together) and their total momentum: the rules of quantum mechanical measurement permit this because those two quantities are independent. Later, when they are far apart, we suddenly measure the momentum of one of them, thereby collapsing its momentum wave function to a definite value. However, the equations of quantum mechanics imply that the total momentum of the two particles is conserved. Therefore the momentum of the second particle *also* takes on a definite value at the instant we measure the momentum of the first. Measuring one particle collapses the wave function of the other, because we know what the total has to be.

It is as if there is some kind of instantaneous communication between the particles. But this kind of action at a distance would violate the principle that no information can travel faster than light, a principle known as *locality*. Einstein, Podolsky, and Rosen felt that "no reasonable definition of reality could be expected to permit this." Bohr, on the other hand, saw no difficulty. Until you actually measure what the second particle is doing, you have no right to consider it as being in any particular state at all. It is therefore meaningless to ask whether its wave function has collapsed, and an alleged collapse that you cannot actually observe cannot be considered as the passage of information.

Bohm's Interpretation

In 1952 David Bohm attempted to resolve the EPR paradox in a novel way. Instead of arguing about interpretations of quantum mechanics, he reformulated the underlying mathematics. One pillar of Bohm's scheme was to endow the wave function with physical meaning. To him it was not just a mathematical gadget that operated 'behind the scenes' — it was out there on centre stage along with the particles and waves themselves. I strongly suspect Schrödinger would have agreed, and at one time Paul Dirac had similar views.

Unfortunately we can't measure a quantum wave function directly, but when it comes down to brass tacks we can't measure *anything* directly; what we do is infer its properties from coherent theories of how the universe works. For example, when we weigh chemicals on an old-fashioned balance we do so on the assumption that the law of the lever holds good, that the numbers stamped on the little brass weights have a definite physical meaning, and indeed that there is a concept 'weight' to be measured in the first place. Copenhagenists seem happy enough for the state of a particle to be a superposition of eigenstates, but they don't endow the superposed state with the same physical reality as the eigenstates themselves. The Copenhagen interpretation can be rendered in mathematical terms as follows: particles obey Schrödinger's equation for the wave function, *except* when measurements are made.

Bohm's idea is simpler and more elegant: particles obey Schrödinger's equation for the wave function, period. In Bohm's theory, the laws of physics are totally deterministic. Quantum indeterminacy is not a sign of anything irreducibly probabilitic about the universe, but a sign of the inescapable ignorance of the observer — human or otherwise. Schrödinger's cat, as I have suggested already, is either alive or dead — but *we* don't know which until we open the box. Bohm proved mathematically that this kind of ignorance is just what you need to reproduce the standard statistical predictions of quantum mechanics. You just average out your ignorance and see what's left.

What worries many physicists about Bohm's theory is that — like the EPR paradox — it has an aspect of non-locality. The wave function of a particle is spread out over all of space, and it reacts instantly to any interaction with another particle. This is of course also true in conventional quantum mechanics — whose wave function obeys exactly the same equation as does Bohm's. However, in Bohm's interpretation the wave function is a real physical thing. In the Copenhagen interpretation it is a mathematical fiction; only its component eigenfunctions can be observed, and those only one at a time. Again the discussion is misdirected. What is missing from both the Copenhagen interpretation and Bohm's theory is any understanding of how macroscopic measuring devices (such as Geiger counters and dead cats) produce determinate values. Observations detect eigenstates, not arbitrary (that is, in the Copenhagen interpretation but not Bohm's, superposed) states. Why?

In recent years there have been several attempts to describe, mathematically, how a quantum state evolves (decoheres) during a macroscopic measurement process. In all of these theories the interaction of a quantum system with its environment produces an irreversible change that turns the quantum state into an eigenstate. However, all of these theories are probabilistic: the initial quantum state undergoes a kind of random diffusion which ultimately leads to an eigenstate. Albert Einstein would have been distinctly unhappy about irreducible randomness showing up at any level of quantum theory — even if, as here, it is confined to the measurement process. God may not play dice, but apparently Geiger counters and cats still do.

Dice, here, are quite an appropriate image. When you roll a die, any one of the six possible faces may end up on top. The result of throwing a die is like an eigenstate — it is a special state that is selected by the measurement process. Dice have several faces: quantum systems have several eigenstates. A Copenhagenist would say that the presence of a table mysteriously causes a die to 'collapse' to one of the states 1, 2, 3, 4, 5, 6, and that the rest of the time it is in a superposition of those eigenstates — which is, of course, a mathematical fiction with no intrinsic physical meaning. Bohm would say that it does have a physical meaning but you can't observe it — at least with any conventional apparatus, and perhaps not at all. Proponents of decoherence would say that as the die rolls along the table its state randomly jiggles, and eventually settles down to one of 1, 2, 3, 4, 5, or 6.

Who is right?

The image of dice suggest that they might all be right —

— and all wrong.

Chaos teaches us that anybody, God or cat, can play dice deterministically, while the naive onlooker imagines that something random is going on. The Copenhagenists and Bohm do not notice the dynamical twists and turns of the rolling die as it bounces erratically but deterministically across the table top. They don't even see the table top. Bohm does think that what the die is doing is real, even if unobservable; the Copenhagenists don't even think that. Proponents of decoherence notice the erratic jiggles of the die, and characterise them statistically as a diffusion processs, not realising that undeneath they are actually deterministic.

Nobody tries to write down the equations for a *rolling* die.

Why not?

One good reason is that they think it can't be done.

Bell's Inequality

Perhaps we can explain the strange behaviour of fundamental particles without recourse to irreducible randomess. Why not equip each article with its own deterministic 'internal dynamic'? This should not affect how the particles interact with one another, but it *should* affect how the particle itself behaves. Instead of rolling the quantum dice to decide when to decay, a radioactive atom might monitor its internal dynamic, and decay when that dynamic attains some particular state. Before the advent of chaos we could not even contemplate playing that kind of trick, because the only known behaviour for the internal dynamic was too regular — steady, periodic, or quasiperiodic. The statistics of radioactive decay simply wouldn't fit such a model, let alone the more subtle aspects of interfering wave functions and the like. But chaos gets over that particular difficulty very neatly, and it suggests that what matters is not *whether* God plays dice — but *how.*

Physicists refer to such theories as 'hidden variable' theories, because the internal dynamic is not directly observable and the variables that define its phase space are in effect concealed from the observed reality. Bohm's theory is a kind of hidden variable theory, with the unobservable details of the *real* wave function as hidden variables. There is a celebrated proof that no 'hidden variable' theory can be consistent with quantum mechanics. (How does this affect Bohm's attempt? See below.) Its theoretical aspects were devised by John Bell in 1964. Preliminary experimental confirmation came in in 1972, and the last plausible experimental loopholes were closed in 1982.

Bell's argument can be formulated in a number of different ways: I will choose one that is close to the spirit of our discussion so far, introduced by Bohm. Bohm's scenario requires a source of spin-$^1/_2$ particles (such as electrons) produced in pairs and moving in opposite directions: one stream travels north, the other south, at the same speed. You can therefore keep track of which particles started off close together: call these 'corresponding' particles in the two streams. A spinometer measures the spin of the northbound particles in the 'up' direction; another one measures the spin of the southbound particles in a direction inclined at an angle A to 'up' (in the up/east plane). By combining these measurements it is possible to work out a 'correlation function' C(A) which measures how closely the spins in one stream match the spins of corresponding particles in the other. If C(A) is 0 then the spins (measured in the up direction) of northbound particles are statistically independent of the spins (measured at angle A) of the corresponding southbound particles. If C(A) = 1 then both spins are the same for corresponding particles: if the northbound particle has spin +$^1/_2$ then so does the corresponding southbound one; and similarly for spin -$^1/_2$. If C(A) = -1, the spins are perfectly anticorrelated: the spin of any northbound particle is the exact opposite of that of the corresponding southbound one.

For the sake of argument, Bell assumes that the observed values of the spin are not random, but are determined by 'hidden variables' — some *deterministic* dynamical system whose variables are not observed. (In my analogy with a rolling deterministic die, the dynamical state of the die while it is rolling involves hidden variables such as angular velocities, which you simply do not see if you observe only the final steady state of the die.) Suppose that you run the experiment twice: once with the spinometer in the second stream set to angle A, and once with it set to a different angle B. Bell did a calculation to see how the correlation function — which now depends in a deterministic manner on the dynamics of the hidden variables, a fact that you can exploit in the calculation — changes when angle A is replaced by angle B. And what falls out of the mathematics is *Bell's inequality*:

$$|C(\mathrm{A}) - C(\mathrm{B})| \leq C(\mathrm{A-B})+1.$$

In this kind of experiment, any system whose apparent randomness is driven by a hidden deterministic dynamic must reveal its status by satisfying Bell's inequality.

Subsequent experiments showed that the observed correlation function does *not* satisfy Bell's inequality. This was widely held to be definitive proof that quantum mechanics is *unavoidably* probabilistic. You *can't* plug the masurement gap by making God's dice deterministic.

However, there are loopholes. Bell's proof involves a number of assumptions, most of which he stated rather carefully. In particular, the proof assumes the principle of locality — that no information can travel faster than light. So Bell's inequality does not rule out a Bohm-type theory.

Dice and Determinism

Associated with any deterministic dynamical system, there is a probabilistic system that offers a kind of 'coarse-grained' representation. Instead of telling us exactly which point in phase space the system occupies at a given instant, it tells us just the probability that the point lies in a given region at some instant. The study of such probabilities, called invariant measures, goes back to the early days of statistical mechanics, when mathematicians and physicists were trying to understand gases as complex collections of molecules, bouncing madly off each other. Invariant measures explain why gases have well defined average properties like density and pressure. Before we understood the molecular basis of matter, the only things we knew about the dynamics of gases were probabilistic. Afterwards, we realised that the probabilities are derived from a deterministic — but incredibly complicated — underlying dynamic. So statistical mechanics does have a hidden variable theory, whose variables are the positions and velocities of the gas's component molecules.

Could quantum theory be similar? Our experience to date makes us think that it is irreducibly probabilistic, but where do the probabilities come from? Probabilities are patterns, of a kind, and it is actually rather bizarre to think of probability as a primary physical concept when in every case where we understand the deep structure, probabilities arise from a deterministic dynamic as invariant measures. Indeed the existence of well-defined statistical patterns is evidence for a kind of order that becomes apparent only when we averge over long timescales. What makes the system's distant future resemble its past, even on average? If it's *really* random, why aren't they just different? If a radioactive atom decays in a manner that has well defined statistical regularities, where do those regularities come from? To say they are fundamentally probabilistic, and leave it at that, is simply to postulate a pattern that ought to be explained.

For deterministic chaos, in contrast, there is a clear mathematical explanation of the associated probabilities and their statistical regularities. *We know where they come from*: they arise as invariant measures. It is the determinism of the dynamics that makes the future look similar to the past. The existence of statistical regularities in quantum-level matter needs to be explained, not simply assumed; and some kind of chaotic hidden variable theory would fit the bill— if only it weren't for Bell's inequality.

However, there are more ways to Bell a cat than choking it with correlations. In principle, we can get round Bell's inequality. It tells us only that *certain kinds* of 'hidden variable' extension of conventional quantum theory can't possibly work, but it doesn't rule out every conceivable extension or alternative. It is a constraint that tells us a little bit about what kind of hidden variable model we can introduce.

Riddled Basins

In 1995 Tim Palmer, a meteorologist with a physics background and an abiding interest in chaos, discovered a very subtle potential loophole in the derivation of Bell's inequality. Basically, it is the assumption that correlation functions are (theoretically) computable. Palmer's conclusion is that quantum indeterminacy may perhaps be replaced by certain kinds of 'hidden variable' chaotic dynamic, provided that the chaos is sufficiently nasty. Nasty enough to wreck the derivation of Bell's inequalities, nice enough to remain deterministic.

The first step is to appreciate just how nasty deterministic chaos can be. A sufficiently nsty case for our purposes is is when there are at least two attractors and one of them has a 'riddled basin'. The basin of an attractor is the set of all points in phase space that are attracted to it. In 1992 Jay Alexander, I.Kan, James Yorke and Z.You discovered that the basin may be riddled with holes — more accurately long thin swirly streaks that reach right in to the attractor. Points in those streaks do *not* move towards the attractor: instead, they are repelled away. Riddled basins are not at all exotic: they show up reliably when the dynamical system obeys a few perfectly reasonable conditions. Indeed riddling can take a more extreme form, in which there are two competing attractors, both have riddled basins, and each basin fills up the holes in the other. Such *intertwined basins* are also perfectly commonplace in the world of nonlinear dynamics.

A system with two attractors whose basins are intertwined is *seriously* unpredictable. You can predict that eventually any chosen initial point will end up on either one attractor or the other, but you can't predict which. As close as you like to initial points that end up on attractor 1 there exist intial points that go to attractor 2, and vice versa. It's a bit like predicting that a rolled die will either come up 1, 2, 3, 4, 5, or 6, without saying which. Indeed a deterministic die behaves very much as if it has *six* attractors, the steady states corresponding to its six faces, all of whose basins are intertwined. For technical reasons that can't quite be true, but it is true that deterministic systems with intertwined basins are wonderful substitutes for dice; in fact they're super-dice, behaving even more 'randomly' — apparently — than ordinary dice. Super-dice are so chaotic that they are uncomputable. Even if you know the equations for the system perfectly, then given an initial state, you cannot calculate which attractor it will end up on. The tiniest error of approximation — and there will always be such an error — will change the answer completely.

You can, however, calculate probabilities. The probability that an initial state lying in some small region ends up on a particular attractor is a computable number associated with that chosen region. Systems with intertwined basins are in principle fully deterministic, so they can be represented by mathematical equations without any explicit random terms. They are practically uncomputable: given an initial state, you cannot calculate with any confidence where it will go. But they are statistically computable: given an ensemble of initial states, you can calculate the probability of ending up on any particular attractor. Palmer's idea is to use this kind of system to provide hidden variables that determine how a quantum state changes when you observe it. Think of an initial point in the hidden variables' phase space as the quantum state before you start to observe it, and the attractors for the hidden variables as representing the possible eigenstates. Because of the statistical computability of such systems, you get well-defined quantum probabilities, consistent with experiments. But the mathematics used to derive the Bell inequality relies on writing down several expressions that use the *actual values* of the hidden variables, and comparing them. Since the dynamics of the hidden variables is uncomputable, those expressions make no sense — so you can't compare them. This is the loophole through which the Bell inequality escapes.

World of If

I often wonder how different today's science would have been if chaos had been discovered *before* quantum mechanics. (Actually that's not likely, because the marvellous computers that make chaos so obvious to everybody rely upon quantum effects to make their circuitry function — but let's pretend. You can *find* chaos without computers, that's what mathematicians did in the sixties. It just required computers to convince everybody else) Now, instead of Einstein protesting that God doesn't play dice, he would probably have suggested that God *does* play dice. Nice, classical, deterministic dice. But — of course — chaotic dice. The mechanism of chaos provides a wonderful opportunity for God to run His universe with deterministic laws, yet simultaneously to make fundamental particles seem probabilistic.

Whether this approach would have established itself is, naturally, debatable. Palmer's work shows that it could at least have got started, and if we can get round the Panda Principle, we may yet find out whether it can deliver the full goods. I can't help thinking that physicists — with Einstein and Schrödinger shouting enthusiastic encouragment — would have tried long and hard to construct a deterministic but chaotic theory of the microscopic world, and that they would have abandoned this idea for a merely probabilistic theory with the utmost reluctance.

Further Reading

John Gribbin, *In Search of Schrödinger's Cat*, Black Swan.
Ian Stewart, *Does God Play Dice?* (2nd edition), Penguin, in press.

© Ian Stewart