The Everett Interpretation

David Wallace

The Oxford Handbook of Philosophy of Physics

Edited by Robert Batterman

Abstract and Keywords

This chapter examines the Everett interpretation of quantum mechanics, or the many-worlds interpretation. It analyzes problems that have been raised for the Everett interpretation, including the problem of providing a preferred basis and the probability problem. The chapter argues for a straightforward, fully realist interpretation of the bare mathematical formalism of quantum mechanics, which, it explain, can make sense of superposed cats without changing the theory and without changing our overall view of science.

Keywords: Everett interpretation, quantum mechanics, many-worlds interpretation, preferred basis, probability problem, mathematical formalism

1. Introduction

The Everett interpretation of quantum mechanics—better known as the Many-Worlds Theory—has had a rather uneven reception. Mainstream philosophers have scarcely heard of it, save as science fiction. In the philosophy of physics it is well known but has historically been fairly widely rejected.[1] Among physicists (at least, among those concerned with the interpretation of quantum mechanics in the first place), it is taken very seriously indeed, arguably tied for first place in popularity with more traditional operationalist views of quantum mechanics.[2]

For this reason, my task in this chapter is twofold. Primarily I wish to provide a clear introduction to the Everett interpretation in its contemporary form; in addition, though, I aim to give some insight into just *why* it is so popular among physicists. For that reason, I begin in section 2 by briefly reprising the measurement problem in a way that (I hope) gives some insight into just why Everett's idea, if workable, is so attractive. In section 3 I introduce that idea and state “the Everett interpretation”—which, I argue in that section, is really just quantum mechanics itself understood in a conventionally realist fashion. In sections 4–10 I explore the consequences of the Everett interpretation via considerations of its two traditional difficulties: the “preferred basis problem” (sections 4–6) and the “probability problem” (sections 8–10). I conclude (sections 11–12) with a brief introduction to other issues in the Everett interpretation and with some further reading.

Little about the Everett interpretation is uncontroversial, but I deal with the controversy rather unevenly. The concepts of *decoherence theory*, as I note in sections 5–6, have significantly changed the debate about the preferred-basis problem, but these insights have only entered philosophy of physics relatively recently, and relatively little in the way of criticism of a decoherence-based approach to Everettian quantum mechanics has appeared as yet (recent exceptions are Hawthorne (2010), Maudlin (2010), and Kent (2010)). By contrast (perhaps because the salient issues are closer to mainstream topics in metaphysics and philosophy of science) the probability problem has been vigorously discussed in the last decade. As such, my discussion of the former fairly uncritically lays out what I see as the correct approach to the definition of the Everett interpretation and to the preferred basis problem. Readers will, I suspect, be better served by forming their own criticisms, and seeking them elsewhere, than by any imperfect attempt of mine to pre-empt criticisms. My discussion of the latter is (somewhat) less opinionated and attempts to give an introduction to the shape of the debate on probability.

I use little technical machinery, but I assume that the reader has at least encountered quantum theory and the measurement problem, at about the level of Albert (1992) or Penrose, (1989, ch. 6).

2. The Measurement Problem

There are philosophical puzzles, perhaps, in how physical theories other than quantum mechanics represent the world, but it is generally agreed that there is no paradox. States of any such theory—be it Newtonian particle mechanics or classical electrodynamics or general relativity—are mathematical objects of some kind: perhaps functions from one space to another, perhaps N-tuples of points in a three-dimensional space, perhaps single points in a high-dimensional, highly-structured space. And, insofar as the theory is correct in a given situation, these states represent the physical world, in the sense that different mathematically defined states correspond to different ways the world can be.[3] There is space for debate as to the nature of this representation—is it directly a relationship between mathematics and the world, or should it be understood as proceeding via some linguistic description of the mathematics?[4]— but these details cause no problems for the straightforward (naive, if you like) view that a theory in physics is a description, or a representation, of the world.

Quantum mechanics, it is widely held, cannot be understood this way. To be sure, it has a clean mathematical formalism—most commonly presented as the evolution of a vector in a highly structured, high-dimensional complex vector space. To be sure, *some* of the states in that space seem at least structurally suited to represent ordinary macroscopic systems: physicists, at least, seem relaxed about regarding so-called “wave packet” states of macroscopic systems as representing situations where those systems have conventional, classically describable characteristics. But central to quantum mechanics is the superposition principle, and it tells us (to borrow a famous example) that if x is a state representing my cat as alive, and ϕ is a state representing my cat as dead, then the “superposition state” (1)

$$\psi = \alpha\chi + \beta\phi$$

is also a legitimate state of the system (where α and β are complex numbers satisfying $|\alpha|^2 + |\beta|^2 = 1$)—and what can *it* represent? A cat that is alive and dead at the same time? An undead cat, in an indefinite state of aliveness? These don't seem coherent ways for the world to be; they certainly don't seem to be ways we observe the world to be.

Nor does the practice of physics seem to treat such states as representing the state of the physical world. Confronted with a calculation that says that the final state of a system after some process has occurred is some superposition like ψ, a theoretician instead declares that the state of the system after the process cannot be known with certainty, but that it has probability $|\alpha|^2$ of being in the macroscopic physical state corresponding to x, and probability $|\beta|^2$ of being in the macroscopic physical state corresponding to ϕ. (If he is more cautious, he may claim only that it has probabilities $|\alpha|^2$, $|\beta|^2$ of being *observed* to be in those macroscopic physical states.) That is, the theoretician treats the mathematical state of the system less like the states of classical mechanics, more like those of classical *statistical* mechanics, which represent not the way the world *is* but a probability distribution over possible ways it *might be*.

But quantum mechanics cannot truly be understood that way either. The most straightforward way to understand why is via quantum interference—the α and β coefficients in ψ can be real or imaginary or complex, positive or negative or neither, and can reinforce and cancel out. Ordinary probability doesn't do that. Put more physically: if some particle is fired at a screen containing two slots, and if conditional on it going through slot 1 it's detected by detector A half the time and detector B half the time, and if conditional on it going through slot 2 it's likewise detected by each detector half the time, then we shouldn't need to know how likely it is to go through slot 1 to predict that it will have a 50% chance of being detected by A and a 50% chance of being detected by B. But a particle in an appropriately weighted superposition of going through each slot can be 100% likely to be detected at A, or 0% likely to be, or anything in between.

So it seems that our standard approach to understanding the content of a scientific theory fails in the quantum case. That in turn suggests a dilemma: either that standard approach is wrong or incomplete, and we need to understand quantum mechanics in a quite different way; or that approach is just fine, but quantum mechanics itself is wrong or incomplete, and needs to be modified or augmented. Call these strategies “change the philosophy” and “change the physics,” respectively.

Famous examples of the change-the-philosophy strategy are the original Copenhagen interpretation, as espoused by Niels Bohr, and its various more-or-less operationalist descendants. Many physicists are attracted to this strategy: they recognize the virtues of leaving quantum mechanics—a profoundly successful scientific theory—unmodified at the mathematical level. Few philosophers share the attraction: mostly they see the philosophical difficulties of the strategy as prohibitive. In particular, attempts to promote terms like “observer” or “measurement” to some privileged position in the formulation of a scientific theory are widely held to have proved untenable.

Famous examples of the change-the-physics strategy are de Broglie and Bohm's pilot-wave hidden variable theory, and Ghirardi, Rimini, and Weber's dynamical-collapse theory. Many philosophers are attracted to this strategy: they recognize the virtue of holding on to our standard picture of scientific theories as representations of an objective reality. Few physicists share the attraction: mostly they see the scientific difficulties of the strategy as prohibitive. In particular, the task of constructing alternative theories that can reproduce the empirical successes not just of nonrelativistic particle mechanics but of Lorentz-covariant quantum field theory has proved extremely challenging.[5]

But for all that both strategies seem to have profound difficulties, it seems nonetheless that one or the other is unavoidable. For we have seen (haven't we?) that if neither the physics of quantum mechanics nor the standard philosophical approach to a scientific theory is to be modified, we do not end up with a theory that makes any sense, far less one that makes correct empirical predictions.

3. Everett's Insight

It was Hugh Everett's great insight to recognize that the apparent dilemma is false— that, *contra* the arguments of section 2, we can after all interpret the bare quantum formalism in a straightforwardly realist way, without either changing our general conception of science or modifying quantum mechanics.

How is this possible? Haven't we just seen that the linearity of quantum mechanics commits us to macroscopic objects being in superpositions, in indefinite states? Actually, no. We have indeed seen that states like ψ—a superposition of states representing macroscopically different objects—are generic in unitary quantum mechanics, but it is actually a non sequitur to go from this to the claim that macroscopic objects are in indefinite states.

An analogy may help here. In electromagnetism, a certain configuration of the field—say, $\mathbf{F}_1(x, t)$ (here $\mathbf{F}$ is the electromagnetic 2-form) might represent a pulse of ultraviolet light zipping between Earth and the Moon. Another configuration, say $\mathbf{F}_2(x, t)$, might represent a different pulse of ultraviolet light zipping between Venus and Mars. What then of the state of affairs represented by (2)

$$\mathbf{F}(x,t) = 0.5\mathbf{F}_1(x,t) + 0.5\mathbf{F}_2(x,t)?$$

What weird sort of thing is this? Must it not represent a pulse of ultraviolet light that is *in a superposition* of traveling between Earth and Moon, and of traveling between Mars and Venus? How can a single pulse of ultraviolet light be in two places at once? Doesn't the existence of superpositions of macroscopically distinct light pulses mean that any attempt to give a realist interpretation of classical electromagnetism is doomed?

Of course, this is nonsense. There is a perfectly prosaic description of $\mathbf{F}$: it does not describe a single ultraviolet pulse in a weird superposition, it just describes two pulses, in different places. And *this*, in a nutshell, is what the Everett interpretation claims about macroscopic quantum superpositions: they are just states of the world in which more than one macroscopically definite thing is happening at once. Macroscopic superpositions do not describe indefiniteness, they describe multiplicity.

The standard terminology of quantum mechanics can be unhelpful here. It is often tempting to say of a given macroscopic system—like a cat, say—that its possible states are all the states in some “cat Hilbert space,” $\mathscr{H}^{\mathrm{cat}}$. Some states in $\mathscr{H}^{\mathrm{cat}}$ are “macroscopically definite” (states where the cat is alive or dead, say); most are “macroscopically indefinite.” From this perspective, it is a very small step to the incoherence of unitary quantum mechanics: quantum mechanics predicts that cats often end up in macroscopically indefinite states; even if it makes sense to imagine a cat in a macroscopically indefinite state, we have certainly never seen one in such a state; so quantum mechanics (taken literally) makes claims about the world that are contradicted by observation.

From an Everettian perspective this is a badly misguided way of thinking about quantum mechanics. This $\mathscr{H}^{\mathrm{cat}}$ is presumably (at least in the nonrelativistic approximation) some sort of tensor product of the Hilbert spaces of the electrons and atomic nuclei that make up the cat. Some states in this box certainly look like they can represent live cats, or dead cats. Others look like smallish dogs. Others look like the Mona Lisa. There is an awful lot that can be made out of the atomic constituents cat of a cat, and all such things can be represented by states in $\mathscr{H}^{\mathrm{cat}}$, and so calling it a "cat Hilbert space" is very misleading.

But if so, it is equally misleading to describe a macroscopically indefinite state cat of $\mathscr{H}^{\mathrm{cat}}$ as representing (say) "a cat in a superposed state of being alive and being dead." It is far more accurate to say that such a state is a superposition of a live cat and a dead cat.

One might still be tempted to object: very well, but we don't observe the universe as being in superpositions of containing live cats and containing dead cats, any more than we observe cats as being in superpositions of alive and dead. But it is not at all clear that we *don't* observe the universe in such superpositions. After all, *cats* are the sort of perfectly ordinary objects that we seem to see around us all the time—a theory that claims that *they* are normally in macroscopically indefinite states seems to make a nonsense of our everyday lives. But the *universe* is a very big place, as physics has continually reminded us, and we inhabit only a very small part of it, and it will not do to claim that it is just "obvious" that it is not in a superposition.

This becomes clearer when we consider what actually happens, dynamically, cat to $\mathscr{H}^{\mathrm{cat}}$, to its surroundings, and to those observing it, when it is prepared in a superposition of a live-cat and a dead-cat state. In outline, the answer is that the system's surroundings will rapidly become entangled with it, so that we do not just have a superposition of live and dead cat, but a superposition of extended quasi-classical regions—"worlds," if you like—some of which contain live cats and some of which contain dead cats. If the correct way to understand such superpositions is as some sort of multiplicity, then our failure to observe that multiplicity is explained quite simply by the fact that we live in one of the "worlds" and the other ones don't interact with ours strongly enough for us to detect them.

This, in short, is the Everett interpretation. It consists of two very different parts: a contingent physical postulate, that the state of the Universe is faithfully represented by a unitarily evolving quantum state; and an a priori claim about that quantum state, that if it is interpreted realistically it must be understood as describing a multiplicity of approximately classical, approximately noninteracting regions that look very much like the "classical world."

And this is *all* that the Everett interpretation consists of. There are no additional physical *postulates* introduced to describe the division into "worlds," there is just unitary quantum mechanics. For this reason, it makes sense to talk about *the* Everett interpretation, as it does not to talk about *the* hidden-variables interpretation or *the* dynamical-collapse interpretation. The "Everett interpretation of quantum mechanics" is just quantum mechanics *itself*, "interpreted" the same way we have always interpreted scientific theories in the past: as modeling the world. Someone might be right or wrong *about* the Everett interpretation—they might be right or wrong about whether it succeeds in explaining the experimental results of quantum mechanics, or in describing our world of macroscopically definite objects, or even in making sense—but there cannot be multiple logically possible Everett interpretations any more than there are multiple logically possible interpretations of molecular biology or classical electrodynamics.[6]

This in turn makes the study of the Everett interpretation a rather tightly constrained activity (a rare and welcome sight in philosophy!). For it is not possible to solve problems with the Everett interpretation by changing the interpretative rules or changing the physics: if there are problems with solving the measurement problem Everett-style, they can be addressed only by hard study—mathematical and conceptual—of the quantum theory we have.

Two main problems of this kind have been identified:

1. the *preferred basis problem* (which might better be called the *problem of branching*)—what actually justifies our interpretation of quantum superpositions in terms of multiplicity?
2. The *probability problem*—how is the Everett interpretation, which treats the Schrödinger equation as deterministic, to be reconciled with the probabilistic nature of quantum theory?

My main task in the remainder of this chapter is to flesh out these problems and the contemporary Everettian response to each.

4. The Preferred Basis Problem

If the preferred basis problem is a question ("how can quantum superpositions be understood as multiplicities?"), then there is a traditional answer, more or less explicit in much criticism of the Everett interpretation (Barrett (1999), Kent (1990), Butterfield (1996)): they cannot. That is: it is no good just *stating* that a state like (1) describes multiple worlds: the formalism must be explicitly *modified* to incorporate them. Adrian Kent put it very clearly in an influential criticism of Everett-type interpretations:

> one can perhaps intuitively view the corresponding components [of the wave function] as describing a pair of independent worlds. But this intuitive interpretation goes beyond what the axioms justify: the axioms say nothing about the existence of multiple physical worlds corresponding to wave function components. (Kent, 1990)

This position dominated discussion of the Everett interpretation in the 1980s and early 1990s: even advocates like Deutsch (1985) accepted the criticism and rose to the challenge of providing such a modification.

Modificatory strategies can be divided into two categories. *Many-exact-worlds theories* augment the quantum formalism by adding an ensemble of "worlds" to the state vector. The "worlds" are each represented by an element in some particular choice of "world basis" $|\psi_i(t)\rangle$ at each time t: the proportion of worlds in state $|\psi_i(t))\rangle$ at time t is $|\langle\psi(t)|\psi_i(t)\rangle|^2$, where $|\psi(t)\rangle$ is the (unitarily evolving) universal state. Our own world is just one element of this ensemble. Examples of many-exact-worlds theories are given by the early Deutsch (1985, 1986), who tried to use the tensor-product structure of Hilbert space to define the world basis,[7] and Barbour (1994, 1999) who chooses the position basis.

In *many-minds theories*, by contrast, the multiplicity is to be understood as illusory. A state like (1) really is indefinite, and when an observer looks at the cat and thus enters an entangled state like (3)

$$\alpha|\text{Live cat}\rangle \otimes |\text{Observer sees live cat}\rangle + \beta|\text{Dead cat}\rangle \otimes |\text{Observer sees dead cat}\rangle$$

then the observer too has an indefinite state. However: to each physical observer is associated not one mental state, but an ensemble of them: each mental state has a definite experience, and the proportion of mental states where the observer sees the cat alive is $|\alpha|^2$. Effectively, this means that in place of a global "world-defining basis" (as in the many-exact-worlds theories) we have a "consciousness basis" for each observer.[8] When an observer's state is an element of the consciousness basis, all the minds associated with that observer have the same experience and so we might as well say that the observer is having that experience. But in all realistic situations the observer will be in some superposition of consciousness-basis states, and the ensemble of minds associated with that observer will be having a wide variety of distinct experiences. Examples of many-minds theories are Albert and Loewer (1988), Lockwood (1989, 1996), Page (1996), and Donald (1990, 1992, 2002). It can be helpful to see the many-exact-worlds and many-minds approaches as embodying two horns of a dilemma: either the many worlds really exist at a fundamental level (in which case they had better be included in the formalism), or they do not (in which case they need to be explained away as somehow illusory).

Both approaches have largely fallen from favor. Partly, this is on internal, philosophical grounds. Many-minds theories, at least, are explicitly committed to a rather unfashionable anti-functionalism—probably even some kind of dualism—about the philosophy of mind, with the relation between mental and physical states being postulated to fit the interests of quantum mechanics rather than being deduced at the level of neuroscience or psychology. If it is just a *fundamental law* that consciousness is associated with some given basis, clearly there is no hope of a functional *explanation* of how consciousness emerges from basic physics (and hence much, perhaps all, of modern AI, cognitive science, and neuroscience is a waste of time[9]). And on closer inspection, many-exact-worlds theories seem to be committed to something as strong or stronger: if "worlds" are to be the kind of thing we see around us, the kind of thing that ordinary macroscopic objects inhabit, then the relation between those ordinary macroscopic objects and the world will likewise have to be postulated rather than derived.

But more important, both approaches undermine the basic motivation for the Everett interpretation. For suppose that a wholly satisfactory many-exact-worlds or many-minds theory were to be developed, specifying an exact "preferred basis" of worlds or minds. Nothing would then stop us from taking that theory, discarding all but one of the worlds/minds[10] and obtaining an equally empirically effective theory without any of the ontological excess that makes Everett-type interpretations so unappealing. Put another way: an Everett-type theory developed along the lines that I have sketched would really just be a hidden-variables theory with the additional assumption that a continuum of many noninteracting sets of hidden variables exists, each defining a different classical world. (This point is made with some clarity by Bell (1981b) in his classic attack on the Everett interpretation.)

At time of writing, almost no advocate of "the many-worlds Interpretation" actually advocates anything like the many-exact-worlds approach[11] (Deutsch, for instance, clearly abandoned it some years ago) and many-minds strategies that elevate consciousness to a preferred role continue to find favor mostly in the small group of philosophers of physics strongly committed for independent reasons to a nonfunctionalist philosophy of mind. Advocates of the Everett interpretation among physicists (almost exclusively) and philosophers (for the most part) have returned to Everett's original conception of the Everett interpretation as a pure interpretation: something that emerges simply from a realist attitude to the unitarily evolving quantum state.

How is this possible? The crucial step occurred in physics: it was the development of *decoherence theory*.

5. The Role of Decoherence

A detailed review of decoherence theory lies beyond the scope of this chapter, but in essence, decoherence theory explores the dynamics of systems that are coupled to some environment with a high number of degrees of freedom. In the most common models of decoherence, the "system" is something like a massive particle and the "environment" is an external environment like a gas or a heat bath, but it is equally valid to take the "system" to be the macroscopic degrees of freedom of some large system and to take the "environment" to be the residual degrees of freedom of that same system. For instance, the large system might be a solid body, in which case the "system" degrees of freedom would be its centre-of-mass position and its orientation and its "environment" degrees of freedom would be all the residual degrees of freedom of its constituents; or it might be a fluid, in which case the "system" degrees of freedom might be the fluid density and velocity averaged over regions a few microns across.

Whatever the system-environment split, "decoherence" refers to the tendency of states of the system to become entangled with states of the environment. Typically no system state is entirely immune to such entanglement, but certain states—normally the wave-packet states, which have fairly definite positions and momentums— get entangled fairly slowly. Superpositions of such states, on the other hand, get entangled with the environment extremely quickly, for straightforward physical reasons: if, say, some stray photon in the environment is on a path that will take it through point q, then its future evolution will be very different according to whether or not there is a wave-packet localized at q. So if the system is in a superposition of being localized at q and being localized somewhere else, pretty soon system-plus-environment will be in a superposition of (system localized at q, photon scattered) and (system localized somewhere else, photon not scattered). Intuitively, we can think of this as the system being constantly measured by the environment, though this "measurement" is just one more unitary quantum-mechanical process.

Mathematically, this looks something like the following. If $|q,p\rangle$ represents a wave-packet state of our macroscopic system with position q and momentum p, then an arbitrary nonentangled state of the system will have state (4)

$$\int \mathrm{d}q\mathrm{d}p\, \alpha(q,p)|q,p\rangle,$$

so that if the environment state is initially $|env_0\rangle$, the combined system-plus-environment state is (5)

$$\left(\int \mathrm{d}q\mathrm{d}p\, \alpha(q,p)|q,p\rangle\right) \otimes |env_0\rangle.$$

But very rapidly (very rapidly, that is, compared to the typical timescales on which the system evolves), this state evolves into something like (6)

$$\int \mathrm{d}q\mathrm{d}p\, \alpha(q,p)|q,p\rangle \otimes |env(q,p)\rangle,$$

where $\langle env(q,p)|env(q',p')\rangle \simeq 0$ unless $q \simeq q'$ and $p \simeq p'$. In this way, the environment records the state of the system, and it does so quickly, repeatedly, and effectively irreversibly (more accurately, it is reversible only in the sense that other macroscopic-scale processes, like the melting of ice, are reversible).

Why does this matter? Because as long as the environment is constantly recording the state of the system in the wave-packet basis, interference experiments cannot be performed on the system: any attempt to create a superposition of wave-packet states will rapidly be undone by decoherence. The overall quantum system (that is, the system-plus-environment) remains in a superposition, but this has no dynamical significance (and, in particular, cannot be empirically detected) without carrying out in-practice-impossible experiments on an indefinitely large region of the universe in the system's vicinity.

And *this* matters, in turn, because it is interference phenomena that allow the different structures represented by a quantum state in a superposition to interact with one another, so as to influence each other and even to cancel out with one another. If interference is suppressed with respect to a given basis, then evolving entangled superpositions of elements of that basis can be regarded as instantiating multiple independently evolving, independently existing structures. As such, if macroscopic superpositions are decohered—as they inevitably will be—then such superpositions really can be taken to represent multiple, dynamically isolated, macroscopic states of affairs.

For this reason, by the mid-1990s decoherence was widely held in the physics community to have solved the preferred basis problem, by providing a definition of Everett's worlds. (It was just as widely held to have solved the measurement problem entirely, independent of the Everett interpretation; since decoherence does not actually remove macroscopic superpositions, though, it was never clear how decoherence alone was supposed to help.) Philosophers of physics were rather more skeptical (Simon Saunders was a notable exception; cf. Saunders Saunders (1993), 1995), essentially because decoherence seems to fall foul of Kent's criticism: however suggestive it might be, it does not seem to succeed in defining an “explicit, precise rule” (Kent 1990) for what the worlds actually are. For decoherence is by its nature an approximate process: the wave-packet states that it picks out are approximately defined; the division between system and environment cannot be taken as fundamental; interference processes may be suppressed far below the limit of experimental detection but they never quite vanish. The previous dilemma remains (it seems): either worlds are part of our fundamental ontology (in which case decoher-ence, being merely a dynamical process within unitary quantum mechanics, and an approximate one at that, seems incapable of defining them), or they do not really exist (in which case decoherence theory seems beside the point).

Outside the philosophy of physics, though (notably in the philosophy of mind, and in the philosophy of the special sciences more broadly), it has long been recognized that this dilemma is mistaken, and that something need not be *fundamental* to be *real*. In the last decade, this insight was carried over to the philosophy of physics.

6. Higher-Order Ontology and the Role of Structure

On even cursory examination, we find that science is replete with perfectly respectable entities that are nowhere to be found in the underlying microphysics.

Douglas Hofstader and Daniel Dennett make this point very clearly:

> Our world is filled with things that are neither mysterious and ghostly nor simply constructed out of the building blocks of physics. Do you believe in voices? How about haircuts? Are there such things? What are they? What, in the language of the physicist, is a hole—not an exotic black hole, but just a hole in a piece of cheese, for instance? Is it a physical thing? What is a symphony? Where in space and time does “The Star-Spangled Banner” exist? Is it nothing but some ink trails in the Library of Congress? Destroy that paper and the anthem would still exist. Latin still *exists* but it is no longer a living language. The language of the cavepeople of France no longer exists at all. The game of bridge is less than a hundred years old. What sort of a thing is it? It is not animal, vegetable, or mineral.
>
> These things are not physical objects with mass, or a chemical composition, but they are not purely abstract objects either—objects like the number pi, which is immutable and cannot be located in space and

> time. These things have birthplaces and histories. They can change, and things can happen to them. They can move about—much the way a species, a disease, or an epidemic can. We must not suppose that science teaches us that every *thing* anyone would want to take seriously is identifiable as a collection of particles moving about in space and time. (Hofstadter and Dennett 1981, 6–7)

The generic philosophy-of-science term for entities such as these is *emergent*: they are not directly definable in the language of microphysics (try defining a haircut within the Standard Model!) but that does not mean that they are somehow independent of that underlying microphysics.

To look in more detail at a particularly vivid example, consider tigers, which are (I take it!) unquestionably real, objective physical objects, even though the Standard Model contains quarks, electrons, and the like, but no tigers. Instead, tigers should be understood as patterns, or structures, *within* the states of that microphysical theory.

To see how this works in practice, consider how we could go about studying, say, tiger hunting patterns. In principle—and only in principle — the most reliable way to make predictions about these would be in terms of atoms and electrons, applying molecular dynamics directly to the swirl of molecules that make up, say, the Kanha National Park (one of the sadly diminishing places where Bengal tigers can be found). In practice, however (even ignoring the measurement problem itself!), this is clearly insane: no remotely imaginable computer would be able to solve the 10^{35} or so simultaneous dynamical equations that would be needed to predict what the tigers would do.

Actually, the problem is even worse than this. For in a sense, we *do* have a computer capable of telling us how the positions and momentums of all the molecules in the Kanha National Park change over time. It is called the Kanha National Park. (And it runs in real time!) Even if, *per impossibile*, we managed to build a computer simulation of the Park accurate down to the last electron, it would tell us no more than what the Park itself tells us. It would provide no explanation of any of its complexity. (It would, of course, be a superb vindication of our extant microphysics.)

If we want to understand the complex phenomena of the Park, and not just reproduce them, a more effective strategy can be found by studying the structures observable at the multi-trillion-molecule level of description of this “swirl of molecules.” At this level, we will observe robust—though not 100% reliable—regularities, which will give us an alternative description of the tiger in a language of cell membranes, organelles, and internal fluids. The principles by which these interact will be deducible from the underlying microphysics (in principle at least; in practice there are usually many gaps in our understanding), and will involve various assumptions and approximations; hence very occasionally they will be found to fail. Nonetheless, this slight riskiness in our description is overwhelmingly worthwhile given the enormous gain in usefulness of this new description: the language of cell biology is both explanatorily far more powerful, and practically far more useful, than the language of physics for describing tiger behavior.

Nonetheless it is still ludicrously hard work to study tigers in this way. To reach a really practical level of description, we again look for patterns and regularities, this time in the behavior of the cells that make up individual tigers (and other living creatures that interact with them). In doing so we will reach yet another language, that of zoology and evolutionary adaptationism, which describes the system in terms of tigers, deer, grass, camouflage, and so on. This language is, of course, the norm in studying tiger hunting patterns, and another (in practice very modest) increase in the riskiness of our description is happily accepted in exchange for another phenomenal rise in explanatory power and practical utility.

The moral of the story is: there are structural facts about many microphysical systems which, although perfectly real and objective (try telling a deer that a nearby tiger is not objectively real) simply cannot be seen if we persist in analyzing those systems in purely microphysical terms. Zoology is of course grounded in cell biology, and cell biology in molecular physics, but the entities of zoology cannot be discarded in favor of the austere ontology of molecular physics alone. Rather, those entities are structures instantiated within the molecular physics, and the task of almost all science is to study structures of this kind.

Of *which* kind? (After all,“structure” and “pattern” are very broad terms: almost any arrangement of atoms might be regarded as some sort of pattern.) The tiger example suggests the following answer, which I have previously (Wallace, 2003a, 93) called “Dennett's criterion” in recognition of the very similar view proposed by Daniel Dennett (1991):

> **Dennett's criterion**: A macro-object is a pattern, and the existence of a pattern as a real thing depends on the usefulness—in particular, the explanatory power and predictive reliability—of theories which admit that pattern in their ontology.

Nor is this account restricted to the relation between physics and the rest of science: rather, it is ubiquitous within physics itself. Statistical mechanics provides perhaps the most important example of this: the temperature of bulk matter is an emergent property, salient because of its explanatory role in the behavior of that matter. (It is a common error in textbooks to suppose that statistical-mechanical methods are used only because in practice we cannot calculate what each atom is doing separately: even if we could do so, we would be missing important, objective properties of the system in question if we abstained from statistical-mechanical talk.) But it is somewhat unusual because (unlike the case of the tiger) the principles underlying statistical-mechanical claims are (relatively!) straightforwardly derivable from the underlying physics.

For an example from physics that is closer to the cases already discussed, consider the case of quasi-particles in solid-state physics. As is well known, vibrations in a (quantum-mechanical) crystal, although they can in principle be described entirely in terms of the individual crystal atoms and their quantum entanglement with one another, are in practice overwhelmingly simpler to describe in terms of “phonons”—collective excitations of the crystal that behave like “real” particles in most respects. And furthermore, this sort of thing is completely ubiquitous in solid-state physics, with different sorts of excitation described in terms of different sorts of “quasi-particle”—crystal vibrations are described in terms of phonons; waves in the magnetization direction of a ferromagnet are described in terms of magnons, collective waves in a plasma are described in terms of plasmons, and so on.[12]

Are quasi-particles real? They can be created and annihilated; they can be scattered off one another; they can be detected (by, for instance, scattering them off “real” particles like neutrons); sometimes we can even measure their time of flight; they play a crucial part in solid-state explanations. We have no more evidence than this that “real” particles exist, and indeed no more grip than this on what makes a particle “real,” and so it seems absurd to deny that quasi-particles exist—and yet, they consist only of a certain pattern within the constituents of the solid-state system in question.

When *exactly* are quasi-particles present? The question has no precise answer. It is essential in a quasi-particle formulation of a solid-state problem that the quasi-particles decay only slowly relative to other relevant timescales (such as their time of flight) and when this criterion (and similar ones) is met then quasi-particles are definitely present. When the decay rate is much too high, the quasi-particles decay too rapidly to behave in any “particulate” way, and the description becomes useless explanatorily; hence, we conclude that no quasi-particles are present. It is clearly a mistake to ask *exactly* when the decay time is short enough (2.54 × the interaction time?) for quasi-particles not to be present, but the somewhat blurred boundary between states where quasi-particles exist and states when they don't should not undermine the status of quasi-particles as real, any more than the absence of a precise boundary to a mountain undermines the existence of mountains.

What has all this got to do with decoherence and Everett? Just this: that the branches which appear in decoherence are precisely the kind of entities that special sciences in general tell us to take seriously. They are emergent, robust *structures* in the quantum state, and as such, we have (it seems) as much reason to take them ontologically seriously as we do any other such structure in science—such as those structures that we identify as chairs and tables, cats and dogs and tigers. So—on pain of rejecting the coherence of the special sciences as a whole—we should accept that unitary quantum mechanics is *already* a many-worlds theory: not a many-exact-worlds theory in which the worlds are *part of* the basic mathematical structure, but an emergent-worlds theory in which the worlds are instantiated as higher-level structures within that basic structure.

In this sense, advocacy of the Everett interpretation has come full circle: the rise and fall of many-exact-worlds and many-minds theories has returned us to Everett's original insight that unitary quantum mechanics should be understood as, not modified to become, a many-worlds theory.

7. Aspects of the Probability Problem

Concerns about probability, and attempts to resolve concerns about probability, have been part of the Everett interpretation since its inception, and the bulk of philosophical work on the interpretation continues to focus on this issue, so that I can do no more here than provide an introduction. I will do so by briefly considering three questions that might be (and indeed have been) raised by critics:

1. How can probability even make sense in the Everett interpretation, given that it is deterministic and that all possible outcomes occur?
2. What justifies the actual form of the quantum probability rule in the Everett interpretation?
3. How can the Everett interpretation make sense of the scientific process by which quantum mechanics was experimentally tested?

Before doing so, however, I make two more general observations. First, if there is a problem of probability in the Everett interpretation, then it is an essentially *philosophical* problem. There is no mystery about how probabilistic theories are mathematically represented in theoretical physics: they are represented by a space of states, a set of histories in that space of states (that is, paths through, or ordered sequences of elements drawn from, that space), and a probability measure over those histories (that is, a rule assigning a probability to each subset of histories, consistent with the probability calculus). Given decoherence, quantum mechanics provides all three (at the emergent level where branches can be defined) just fine, using the standard modulus-squared amplitude rule to define the probability of each branch; indeed, historically much of the motivation of the decoherence program was to ensure that the probability calculus was indeed satisfied by the modulus-squared amplitudes of the branches. So a physicist who objected to the rather philosophical tenor of the debates on probability in the Everett interpretation would be missing the point: insofar as he is unconcerned with *philosophical* aspects of probability, he should have no qualms about Everettian probability at all.

Second, it has frequently been the case that what appear to be philosophical problems with probability *in Everettian quantum mechanics* in fact turn out to be philosophical problems with probability *simpliciter*. Probability poses some very knotty philosophical issues, which often we forget just because we are so used to the concept in practice; sometimes it takes an unfamiliar context to remind us of how problematic it can be.

Note that it is of no use for a critic to respond that all the same we have a good practical grasp of probability in the non-Everettian context but that that grasp does not extend to Everettian quantum physics. For that is exactly the point at issue: the great majority, if not all, of the objective probabilities we encounter in science and daily life ultimately have a quantum-mechanical origin, so if the Everett interpretation is correct, then most of our practical experience of probability is with Everett-type probability.

8. Probability, Uncertainty, and Possibility

How can there be probabilities in the Everett interpretation? (asks the critic): there is nothing for them to be probabilities of! Defenders will reply that the probabilities are probabilities of branches (understood via decoherence), but the objection is that somehow it is illegitimate to assign probabilities to the branches, either because probabilities require uncertainty and it makes no sense to be uncertain of which outcome will occur in a theory like Everett's, or because somehow probabilities quantify alternative possibilities and there are no alternative possibilities in the Everett interpretation.

The conciliatory approach here would be to argue that these concepts do after all find a home in Everettian quantum mechanics; that is, to argue that people in an Everettian universe should indeed regard different branches as different alternative possibilities, and be uncertain as to which one will actually occur. To my knowledge this was first argued for by Saunders (1998), via an ingenious thought experiment related to traditional intuition pumps in the philosophy of personal identity; Saunders' goal was to make it intuitive that someone in an Everettian universe should indeed be uncertain about their future, even if they knew the relevant facts about the future branches (though see Greaves (2004) for an attempted rebuttal). Subsequent work (much of it building on Saunders') has tried to go beyond intuitive plausibility and give a positive account of what would ground uncertainty in the Everett interpretation.

I am aware of three broad strategies of this kind. First, and most directly, Lev Vaidman points out (Vaidman 2002) that someone who carried out a quantum measurement *but did not observe the result* would be in a state of genuine (albeit indexical) uncertainty. (There would be multiple copies of the experimenter, some in branches with one result and some in branches with another, but each would be in subjectively identical states.) It is unclear whether this notion of uncertainty (which does not appear to apply to pre-measurement situations) is sufficient to assuage concerns.

An alternative approach via indexical uncertainty—this time also applying to the pre-measurement situation—is to think about branches as four-dimensional rather than three-dimensional entities (thus entailing that branches overlap in some sense[13] prior to whatever quantum event causes them to diverge. Uncertainty is then to be understood as uncertainty as to which four-dimensional branch an observer is part of. For exploration and defense of this position, see Saunders and Wallace (2008a, 2008b), Saunders (2010), and Wilson (2010a, 2010b); for criticism, see Lewis (2007b) and Tappenden (2008).

The third strategy is closely related to the second, but takes its cue from semantics rather than from metaphysics: namely, consider the way in which words like “uncertainty” would function in an Everettian universe (possibly given some theory of semantic content along the “charity” lines advocated by Lewis (1974), Davidson (1973), and others) and argue that they would in fact function in such a way as to make claims like “one or other outcome of the measurement will occur, but not both” actually turn out correct. I explore this idea in Wallace (2005, 2006) and in chapter 7 of Wallace (2012); see also Ismael (2003) for a position that combines aspects of the second and third strategies. Whether such semantical considerations are metaphysically (let alone physically) relevant depends on one's view of metaphysics; Albert (2010), for instance, argues that they are irrelevant.

A conciliatory approach of a rather different kind is to concede that probability has no place in an Everettian world and to show how one can do without it; typically, this is done by arguing that human activity in general, and science in particular, would proceed as if quantum-mechanical mod-squared amplitude was probability, even if “really” it was not. Deutsch (1999) and Greaves (2004) advocate positions of this kind; both regard “probability” as something to be understood decision-theoretically, via an agent's actions. If it can be argued that (rational) agents in an Everettian world would act as if each branch has a certain probability, then (Deutsch and Greaves argue) this is sufficient.

Of course, there is also a decidedly nonconciliatory response available: just to deny the claim that genuine probability requires either alternative probabilities or any form of uncertainty. One seldom hears actual *arguments* for these requirements; typically they are just stated as if they were obvious. And perhaps they are *intuitively* obvious, but it is not clear that this has any particular bearing on anything. Someone who adopts the (hopelessly unmotivated) epistemological strategy of regarding intuitive obviousness as a guide to truth in theoretical physics will presumably have given up on the Everett interpretation long ago in any case.

This response is actually fairly close to Deutsch's and Greaves's position: if it can be argued that mod-squared amplitude functions exactly like probability but lacks certain standardly required philosophical features that probability has, it is open to us just to deny that those philosophical features are required, and to adopt the position that insofar as mod-squared amplitude functions *exactly like* probability, that's all that's required to establish that it *is* probability. This is my own view on the problem,[14] developed *in extenso* in Wallace (2012).

9. The Quantitative Problem

Grant, if only for the sake of argument, that it is somehow legitimate to attach probabilities to branches. There is a further question: Why should those probabilities be required to equal those given by quantum mechanics?

One version of this objection—going right back to Graham (1973)—is that the quantum probabilities *cannot* be the right probabilities, because the right probabilities must give each branch equal probability. There is generally no positive *argument* given for this claim, beyond some gesture to the effect that the versions of me on the different branches are all “equally me”; still, it has a strong intuitive plausibility.

It can, however, be swiftly dismissed. It is possible to argue that the rule is actually inconsistent when branching events at multiple times are considered,[15] but more crucially, decoherence just does not license any notion of branch count. It makes sense, in the presence of decoherence, to say that the quantum state (or some part of it) branches into a part in which measurement outcome X occurs and a part in which it does not occur, but it makes no sense at all to say *how many branches* comprise the part in which X occurs. Study of a given branch at a finer level of detail will inevitably show it to consist of many sub-branches; eventually this will cease to be the case as decoherence ceases to be applicable and interference between branches becomes nonnegligible; but there is no well-defined point at which this occurs and different levels of tolerance—as well as small changes in other details of how we define "branch"—lead to wildly differing answers as to how many branches there are. Put plainly,"branch number," insofar as it is defined at all in a given decoherence formalism, is an artifact of the details of that formalism. (And it is not by any means defined in all such formalisms; many use a continuum framework in which the concept makes no sense even inside the formalism. For more details on this and on the general question of branch counting, see chapter 3 of Wallace (2012).)

So much for branch counting. The question remains: What positive justification can be given for identifying mod-squared amplitude with probability? One might answer, as did Simon Saunders in the 1990s (Saunders 1995, 1997, 1998), by rejecting the idea that any "positive justification" is needed: after all,in general we do not argue that the probabilities in a physical theory are what they are (nor indeed, in general, that the other physical magnitudes in a theory have the interpretation they have); we just postulate it. It is not immediately clear why this response is any less justified in the Everett interpretation than in non-Everettian physics; indeed, arguably it works rather better as a postulate, since it is at least clear what categorical, previously understood magnitude is to be identified with probability. By contrast, in *classical physics* the only real candidate seems to be long-run relative frequencies or some related concept, and even establishing that those have the *formal* properties required of probability has proven fraught. The most promising candidate so far is Lewis's "best systems analysis" (Lewis 1986, 55, 128–131), which constructs probabilities indirectly from relative frequencies and related categorical data; even if that analysis succeeded fully, though, it would deliver no more than quantum physics (together with decoherence) has already delivered, namely a set of quantities with the right formal properties to be identified with probability but no further justification for making such an identification.[16]

Papineau (1996, 2010) puts essentially the same point in a more pessimistic way. He identifies two criteria that a theory of probability must satisfy: a "decision-theoretic link" (why do we use probability as a guide to action?) and an "inferential link" (why do we learn about probabilities from observed relative frequencies?) and concedes that Everettian quantum mechanics has no good explanation of why either is satisfied—but, he continues, neither does any other physical theory, nor any other extant philosophical theory of probability. The Everett interpretation (Papineau argues) therefore has no *special* problem of probability.

In fact, in recent years the possibility has arisen that probability may actually be in *better* shape in Everettian quantum mechanics than in non-Everettian physics. Arguments originally given by David Deutsch (1999) and developed in Wallace (2003b, 2007) suggest that it may be possible to derive the quantum probability rule from general principles of decision theory, together with the mathematical structure of quantum mechanics shorn of its probabilistic interpretation. A fully formalized version of this argument can be found in Wallace (2010) and in chapters 5 and 6 of Wallace (2012).

In philosophical terms, what such arguments attempt to do is to show that rational agents, cognizant of the facts about quantum mechanics and conditional on believing those facts to be true, are required to treat mod-squared amplitude operationally as probability. Specifically, they are required to use observed relative frequencies as a guide to working out what the unknown mod-squared amplitudes are (Papineau's inferential link), and to use known mod-squared amplitudes as a guide to action (his decision-theoretic link).[17]

Space does not permit detailed discussion of this approach to probability, but at essence it relies on the symmetries of quantum mechanics. There is a long tradition of deriving probability from considerations of symmetry, but in the classical case these approaches ultimately struggle with the fact that *something* must break the symmetry, simply to explain why one outcome occurs rather than another. This is, of course, not an issue for Everettian quantum mechanics! From this perspective, the role of decision theory is less central in the arguments than it might appear: its main function is to justify the applicability of probabilistic concepts to Everettian branches at all. (And conversely, if one were concerned *purely* with the question of what the probabilities of each branch were, and prepared to grant that branches *did* have probabilities and that they satisfy normal synchronic and diachronic properties, it is possible to prove the quantum probability rule without any mention of decision theory; cf. Wallace (2012, ch. 4).)

If this last approach to probability works (and fairly obviously, I believe it does), it marks a rather remarkable shift in the debate; probability, far from being something that makes the Everett interpretation unintelligible, becomes something that can be understood in Everettian quantum mechanics in a way which does not seem available otherwise. (See Saunders (2010) for further development of this theme.) I feel obliged to note that it is highly controversial whether the approach does indeed work; for recent criticism, see Albert (2010), Price (2010), and other articles in Saunders et al. (2010).

10. Epistemic Puzzles

The rise of decision-theoretic approaches to Everettian probability (whether to make sense of probability or to derive the probability rule) has led to a new worry about probability in the Everett interpretation. Suppose for the sake of argument that it really can be shown, or legitimately postulated, that someone who accepts the Everett interpretation as correct should behave, at least for all practical purposes, as if mod-squared amplitude were probability. What has that to do with the question of why we should believe the Everett interpretation in the first place? Put another way, how would it license us to interpret the usual evidence for quantum mechanics as evidence for *Everettian* quantum mechanics?

This suggests a division of the probability problem into *practical* and *epistemic* problems (Greaves 2007a), where the former concerns how rational agents should act *given that Everettian quantum theory is correct*, and the latter concerns how evidence bears on the truth of quantum theory in the first place, given that it is to be interpreted à la Everett. Arguably, Deutsch's decision-theoretic program (and my development of it) speaks only to the practical problem; indeed, arguably most of the tradition in thinking about Everettian probability speaks only to the practical problem.

The last decade has seen the development of a small, but complex, literature on this subject. In essence, there are two strategies that have been developed for solving the epistemic problem. The first is highly philosophical: if it can be established that mod-squared amplitude *is probability*, then (it is claimed) no more is required of the Everett interpretation than of any other physical theory as regards showing why probability plugs into our epistemology in the way it does. Strategies of this form rely on a mixture of solutions to the practical problem (cf. section 9), arguments that branching leads to genuine uncertainty about the future and/or genuine probabilities (cf. section 8), and appeal to the no-double-standards principle I mentioned in section 7. The strategy is tacit in Saunders (1998); I defended an explicit version in Wallace (2006) (and, in less developed form, in Wallace (2002)); Wilson (2010b) defends a similar thesis.

The other strategy is significantly more technical and formal: namely, construct a formal decision-theoretic framework to model the epistemic situation of agents who are unsure whether the Everett interpretation is correct, and show that in that situation (perhaps contingent on a solution to the practical problem), agents regard “ordinary” evidence as confirmatory of quantum mechanics in a standard way, even when quantum mechanics is understood according to the Everett interpretation. This strategy was pioneered by Greaves (2004) and brought to a mature state in Greaves (2007a) and Greaves and Myrvold (2010). The latter two papers, on slightly different starting assumptions (including in both cases the Bayesian approach to statistical inference) take it as given that conditional on the Everett interpretation being true, mod-squared amplitude functions as probability in decision-making contexts, and derive that agents will update their personal probability in quantum mechanics via standard update procedures, whether or not quantum probabilities are to be understood in Everettian terms. As such, these arguments take as input a solution to the practical problem (whether postulated or derived via Deutsch's and/or my arguments) and give as output a solution to the epistemic problem. It is also possible (cf. Wallace, 2012, ch. 6) to combine the two strategies into one theorem, which makes standard decision-theoretic assumptions and derives solutions to the epistemic and practical problems in a unified fashion.

11. Other Topics

While the bulk of contemporary work on the Everett interpretation has been concerned with the preferred-basis and probability problems (and, more generally, has been concerned with whether the interpretation is viable, rather than with its philosophical implications if viable), there are a goodly number of other areas of interest within the Everett interpretation (or, as I would prefer to put it: within quantum mechanics, once it is understood that it should be interpreted Everett-style), and I briefly mention some of these here.

- If Everettian quantum mechanics is only *emergently* a theory of branching universes, what is its *fundamental* ontology, insofar as that question has meaning? That is: what kind of physical entity is represented by the quantum state? Of course, this question can be asked of any approach to quantum theory that takes the state as representing something physically real, but it takes on a particular urgency in the Everett interpretation given that the theory is supposed to be *pure* quantum mechanics, shorn of any additional mathematical structure. For various approaches to the problem, see Deutsch and Hayden (2000), Deutsch (2002), Wallace and Timpson (2007, 2010), Maudlin (2010) (who argues that there is *no* coherent understanding of the Everett interpretation's ontology), Hawthorne (2010) (who is at least sympathetic to Maudlin), Allori et al. (2009), and (in the general context of the ontology of the quantum state) Albert (1996) and Lewis (2004b).
  (I should add one cautionary note: it is very common in the literature to phrase the question as, what is the ontology of the wave-function? But recall that the wave-function is only one of a great many ways to represent the quantum state, and one which is much more natural in nonrelativistic physics than in quantum field theory.)
- It is generally (and in my view correctly) held that the experimental violation of Bell's inequalities[18] shows not just that hidden variable theories must involve superluminal dynamics, but that *any* empirically adequate theory must involve superluminal dynamics.[19] But the Everett interpretation is generally (and again correctly, in my view) viewed as an exception, essentially because it violates a tacit premise of Bell's derivation, that only one outcome actually occurs.[20] There has, however, been rather little exploration of this issue; Bacciagaluppi (2002) is a notable exception.[21]
- There is an ongoing (and somewhat sensationalist) discussion in the literature about so-called “quantum suicide”: the idea that an agent in an Everettian universe should expect with certainty to survive any process which third-party observers regard him as having nonzero probability of surviving. The idea has been around in the physics community for a long time (see, e.g., Tegmark (1998); it was first introduced to philosophers by David Lewis, in his only paper on the Everett interpretation (Lewis 2004a) and has been discussed further by Lewis (2000) and Papineau (2003).
- Everett was originally motivated in part by a desire for an interpretation of quantum mechanics that was suitable for cosmology in that it did not assume an external observer. The Everett interpretation has been widely influential in quantum cosmology ever since: for an introduction, see Hartle (2010). It is not *universally* acknowledged that quantum cosmology does require the Everett interpretation, though; for dissenting views (from widely differing perspectives), see Fuchs and Peres (2000), Smolin (1997: 240–266), and Rovelli (2004: 209–222).
- The de Broglie-Bohm “pilot wave” theory (aka Bohmian mechanics) has sometimes been criticized for being “Everett in denial”: that is, being the Everett interpretation with some additional epiphenomenal structure. For examples of this criticism, see Deutsch (1996) and Brown and Wallace (2004); for responses, see Lewis (2007a) and Valentini (2010) (see also Brown's (2010) response to Valentini). Allori et al. (2008) can also be read as a response, insofar as it advocates a position on the ontology of a physical theory far removed from that of section 6 and from which the Everett-in-denial objection cannot be made.

12. Further Reading

Saunders et al. (2010) is an up-to-date and edited collection of articles for and against the Everett interpretation, including contributions from a large fraction of the physicists and philosophers involved in the contemporary debate; Saunders's introduction to the book provides an overview of the Everett interpretation complementary to this chapter. Barrett (1999) is a comprehensive guide to discussions of the Everett interpretation in (mostly) the philosophy of physics literature, up to the late 1990s. DeWitt and Graham (1973) is a classic collection of original papers. Wallace (2012) is my own book-length defense of the Everett interpretation; Wallace (2008) is a review of the measurement problem more generally, focused on the role of decoherence theory. Greaves (2007b) reviews work in the probability problem.

Afterword

I have left undiscussed the often-unspoken, often-felt objection to the Everett interpretation: that it is simply unbelievable. This is because there is little to discuss: that a *scientific theory* is wildly unintuitive is no argument at all against it, as twentieth-century physics proved time and again. David Lewis is memorably reported to have said that he did not know how to refute an incredulous stare; had he been less charitable, he might have said explicitly that an incredulous stare is not an argument, and that if someone says that they are incapable of believing a given theory—philosophical or scientific—they are but reporting on their psychology.

References

Abrikosov, A. A., L. P. Gorkov, and I. E. Dzyalohinski (1963). *Methods of quantum field theory in statistical physics*. New York: Dover. Revised English edition; translated and edited by R. A. Silverman.

Albert, D. Z. (1992). *Quantum mechanics and experience*. Cambridge, MA: Harvard University Press.

———. (1996). Elementary quantum metaphysics. In J. T. Cushing, A. Fine, and S. Goldstein (Eds.), *Bohmian mechanics and quantum theory: An appraisal*, 277–284. Dordrecht: Kluwer Academic Publishers.

———. (2010). Probability in the Everett picture. In S. Saunders, J. Barrett, A. Kent, and D. Wallace (Eds.), *Many worlds? Everett, quantum theory and reality*. Oxford: Oxford University Press.

Albert, D. Z., and B. Loewer (1988). Interpreting the many worlds interpretation. *Synthese* 77: 195–213.

Allori, V., S. Goldstein, R. Tumulka, and N. Zanghi (2008). On the common structure of Bohmian mechanics and the Ghirardi–Rimini–Weber theory. *British Journal for the Philosophy of Science* 59: 353–389.

———. (2009). Many-worlds and Schrödinger's first quantum theory. Forthcoming in *British Journal for the Philosophy of Science*; available online at **http://arxiv.org/abs/0903.2211**.

Bacciagaluppi, G. (2002). Remarks on space-time and locality in Everett's interpretation. In J. Butterfield and T. Placek (Eds.), *Non-locality and modality*, Vol. 64 of Nato Science Series II. Mathematics, Physics and Chemistry. 105–122. Dordrecht: Kluwer.

Barbour, J. B. (1994). The timelessness of quantum gravity: II. The appearance of dynamics in static configurations. *Classical and Quantum Gravity* 11: 2875–2897.

———. (1999). *The end of time*. London: Weidenfeld and Nicholson.

Barrett, J. A. (1999). *The quantum mechanics of minds and worlds*. Oxford: Oxford University Press.

Bell, J. S. (1981a). Bertlmann's socks and the nature of reality. *Journal de Physique* 42, C2: 41–61. Reprinted in Bell (1987), 139–158.

———. (1981b). Quantum mechanics for cosmologists. In C. J. Isham, R. Penrose, and D. Sciama (Eds.), *Quantum gravity 2: A second Oxford Symposium*. Oxford: Clarendon Press. Reprinted in Bell (1987), 117–138; page references are to that version.

———. (1987). *Speakable and unspeakable in quantum mechanics*. Cambridge: Cambridge University Press.

Brown, H. (2010). Reply to Valentini: “de Broglie-Bohm theory: Many worlds in denial?”. In *Many worlds? Everett, quantum theory, and reality*. Oxford: Oxford University Press. 510–517.

Brown, H. R., and D. Wallace (2004). Solving the measurement problem: de Broglie-Bohm loses out to Everett. Forthcoming in *Foundations of Physics*; available online at **http://arxiv.org/abs/quant-ph/0403094**.

Butterfield, J. N. (1996). Whither the minds? *British Journal for the Philosophy of Science* 47: 200–221.

Colin, S. (2003). Beables for quantum electrodynamics. Available online at **http://arxiv.org/abs/quant-ph/0310056**.

Colin, S., and W. Struyve (2007). A Dirac sea pilot-wave model for quantum field theory. *Journal of Physics A* 40: 7309–7342.

Davidson, D. (1973). Radical interpretation. *Dialectica* 27: 313–328.

Dennett, D. C. (1991). Real patterns. *Journal of Philosophy 87*, 27–51. Reprinted in D. Dennett, *Brainchildren* (London: Penguin, 1998), 95–120.

Deutsch, D. (1985). Quantum theory as a universal physical theory. *International Journal of Theoretical Physics* 24(1): 1–41.

———. (1986). Interview. In P. Davies and J. Brown (Eds.), *The ghost in the atom*, 83–105. Cambridge: Cambridge University Press.

———. (1996). Comment on Lockwood. *British Journal for the Philosophy of Science* 47: 222–228.

———. (1999). Quantum theory of probability and decisions. *Proceedings of the Royal Society of London* A455: 3129–3137. Available online at **http://arxiv.org/abs/quant-ph/9906015**.

———. (2002). The structure of the multiverse. *Proceedings of the Royal Society of London* A458, 2911–2923.

Deutsch, D., and P. Hayden (2000). Information flow in entangled quantum systems. *Proceedings of the Royal Society of London* A456: 1759–1774. Available online at **http://arxiv.org/abs/quant-ph/9906007**.

DeWitt, B., and N. Graham, eds. (1973). *The many-worlds interpretation of quantum mechanics*. Princeton: Princeton University Press.

Donald, M. J. (1990). Quantum theory and the brain. *Proceedings of the Royal Society of London* A427: 43–93.

———. (1992). A priori probability and localized observers. *Foundations of Physics* 22: 1111–1172.

———. (2002). Neural unpredictability, the interpretation of quantum theory, and the mind–body problem. Available online at **http://arxiv.org/abs/quant-ph/0208033**.

Dürr, D., S. Goldstein, R. Tumulka, and N. Zanghi (2004). Bohmian mechanics and quantum field theory. *Physical Review Letters* 93: 090402.

———. (2005). Bell-type quantum field theories. *Journal of Physics* A38: R1.

Everett, H. I. (1957). Relative state formulation of quantum mechanics. *Review of Modern Physics* 29: 454–462. Reprinted in DeWitt and Graham (1973).

Farhi, E., J. Goldstone, and S. Gutmann (1989). How probability arises in quantum-mechanics. *Annals of Physics* 192: 368–382.

Foster, S., and H. R. Brown (1988). On a recent attempt to define the interpretation basis in the many worlds interpretation of quantum mechanics. *International Journal of Theoretical Physics* 27: 1507–1531.

Fuchs, C., and A. Peres (2000). Quantum theory needs no "interpretation." *Physics Today* 53(3): 70–71.

Graham, N. (1973). The measurement of relative frequency. In B. DeWitt and N. Graham (Eds.), *The many-worlds interpretation of quantum mechanics*. Princeton: Princeton University Press: 229–252.

Greaves, H. (2004). Understanding Deutsch's probability in a deterministic multiverse. *Studies in the History and Philosophy of Modern Physics* 35: 423–456.

———. (2007a). On the Everettian epistemic problem. *Studies in the History and Philosophy of Modern Physics* 38: 120–152.

———. (2007b). Probability in the Everett interpretation. *Philosophy Compass* 38: 120–152.

Greaves, H., and W. Myrvold (2010). Everett and evidence. In S. Saunders, J. Barrett, A. Kent, and D. Wallace (Eds.), *Many worlds? Everett, quantum theory and reality*. Oxford: Oxford University Press. 229–252.

Hartle, J. (2010). Quasiclassical realms. In S. Saunders, J. Barrett, A. Kent, and D. Wallace (Eds.), *Many worlds? Everett, quantum theory, and reality*. Oxford: Oxford University Press: 73–98.

Hawthorne, J. (2010). A metaphysician looks at the Everett interpretation, in Saunders et al. (2010), 144–153.

Hofstadter, D. R., and D. C. Dennett, eds. (1981). *The mind's I: Fantasies and reflections on self and soul*. London: Penguin.

Ismael, J. (2003). How to combine chance and determinism: Thinking about the future in an Everett universe. *Philosophy of Science* 70: 776–790.

Kent, A. (1990). Against many-worlds interpretations. *International Journal of Theoretical Physics* A5: 1745–1762. Expanded and updated version at **http://www.arxiv.org/abs/gr-qc/9703089**.

———. (2010). One world versus many: The inadequacy of Everettian accounts of evolution, probability, and scientific confirmation. In Saunders, Barrett, Kent, and Wallace (2012), 307–354.

Kittel, C. (1996). *Introduction to solid state physics*. 7th ed. New York: Wiley.

Ladyman, J., and D. Ross (2007). *Every thing must go: Metaphysics naturalized*. Oxford: Oxford University Press.

Lewis, D. (1974). Radical interpretation. *Synthese* 23: 331–344. Reprinted in David Lewis, *Philosophical Papers*, Vol. I. Oxford: Oxford University Press, 1983.

———. (1980). A subjectivist's guide to objective chance. In R. C. Jeffrey (Ed.), *Studies in inductive logic and probability*. Vol. 2. Berkeley: University of California Press. Reprinted, with postscripts, in David Lewis, *Philosophical Papers*, Vol. 2. Oxford: Oxford University Press, 1986; page numbers refer to this version.

———. (1986). *Philosophical Papers*. Vol. 2. Oxford: Oxford University Press.

———. (2004a). How many lives has Schrödinger's cat? *Australasian Journal of Philosophy* 81: 3–22.

Lewis, P. (2007a). How Bohm's theory solves the measurement problem. *Philosophy of Science* 74: 749–760.

Lewis, P. J. (2000). What is it like to be Schrödinger's cat? *Analysis* 60: 22–29.

———. (2004b). Life in configuration space. *British Journal for the Philosophy of Science* 55: 713–729.

———. (2007b). Uncertainty and probability for branching selves. *Studies in the History and Philosophy of Modern Physics* 38: 1–14.

Lockwood, M. (1989). *Mind, brain and the quantum: The compound 'I'*. Oxford: Blackwell Publishers.

———. (1996). “Many minds” interpretations of quantum mechanics. *British Journal for the Philosophy of Science* 47: 159–188.

Maudlin, T. (2002). *Quantum non-locality and relativity: Metaphysical intimations of modern physics*. 2d ed. Oxford: Blackwell.

———. (2010). Can the world be only wavefunction? In Saunders, Barrett, Kent, and Wallace (2010), 121–143.

McCall, S. (2000). QM and STR: The combining of quantum mechanics and relativity theory. *Philosophy of Science* 67: S535–S548.

Myrvold, W. (2002). On peaceful coexistence: Is the collapse postulate incompatible with relativity? *Studies in the History and Philosophy of Modern Physics* 33: 435–466.

Page, D. N. (1996). Sensible quantum mechanics: Are probabilities only in the mind? *International Journal of Modern Physics* D5: 583–596.

Papineau, D. (1996). Many minds are no worse than one. *British Journal for the Philosophy of Science* 47: 233–241.

———. (2003). Why you don't want to get into the box with Schrödinger's cat. *Analysis* 63: 51–58.

———. (2010). A fair deal for Everettians. In Saunders, Barrett, Kent, and Wallace (2010), 205–226.

Penrose, R. (1989). *The emperor's new mind: Concerning computers, brains and the laws of physics*. Oxford: Oxford University Press.

Price, H. (2010). Decisions, decisions, decisions: can Savage salvage Everettian probability? In Saunders, Barrett, Kent, and Wallace (2010), 369–390.

Rovelli, C. (2004). *Quantum gravity*. Cambridge: Cambridge University Press.

Saunders, S. (1993). Decoherence, relative states, and evolutionary adaptation. *Foundations of Physics* 23: 1553–1585.

———. (1995). Time, decoherence and quantum mechanics. *Synthese* 102: 235–266.

———. (1997). Naturalizing metaphysics. *The Monist* 80(1): 44–69.

———. (1998). Time, quantum mechanics, and probability. *Synthese* 114: 373–404.

———. (2003). Physics and Leibniz's principles. In K. Brading and E. Castellani (Eds.), *Symmetries in physics: Philosophical reflections*, 289–308. Cambridge: Cambridge University Press.

———. (2010). Chance in the Everett interpretation. In Saunders, Barrett, Kent, and Wallace (2010), 181–205.

Saunders, S., and D. Wallace (2008a). Branching and uncertainty. *British Journal for the Philosophy of Science* 59: 293–305.

———. (2008b). Saunders and Wallace reply. *British Journal for the Philosophy of Science* 59: 315–317.

Saunders, S., J. Barrett, A. Kent, and D. Wallace, eds. (2010). *Many worlds? Everett, quantum theory, and reality*. Oxford: Oxford University Press.

Smolin, L. (1997). *The life of the cosmos*. New York: Oxford University Press.

Struyve, W., and H. Westman (2006). A new pilot-wave model for quantum field theory. *AIP Conference Proceedings* 844: 321.

Tappenden, P. (2008). Saunders and Wallace on Everett and Lewis. *British Journal for the Philosophy of Science* 59: 307–314.

Tegmark, M. (1998). The interpretation of quantum mechanics: Many worlds or many words? *Fortschrift Fur Physik* 46: 855–862. Available online at **http://arxiv.org/abs/quant-ph/9709032**.

Timpson, C., and Brown, H. R. (2002). Entanglement and relativity. In R. Lupacchini and V. Fano (Eds.), *Understanding physical knowledge*, Preprint no. 24, Departimento di Filosofia, Università di Bologna, CLUEB, 2002, 147–166. Available online at **http://arxiv.org/abs/quant-ph/0212140**.

Tsvelik, A. M. (2003). *Quantum field theory in condensed matter physics*. 2d ed. Cambridge: Cambridge University Press.

Tumulka, R. (2006). Collapse and relativity. In A. Bassi, T. Weber, and N. Zanghi (Eds.), *Quantum mechanics: Are there quantum jumps? and on the present status of quantum mechanics*, 340. American Institute of Physics Conference Proceedings. Available online at **http://arxiv.org/abs/quant-ph/0602208**.

Vaidman, L. (2002). The many-worlds interpretation of quantum Mechanics. In *the Stanford Encyclopedia of Philosophy* (Summer 2002 edition), ed. Edward N. Zalta. Available online at **http://plato.stanford.edu/archives/sum2002/entries/qm-manyworlds**.

Valentini, A. (2010). De Broglie-Bohm pilot wave theory: Many worlds in denial? In Saunders, Barrett, Kent, and Wallace (2010), 476–509.

Wallace, D. (2002). Quantum probability and decision theory, revisited. Available online at **http://arxiv.org/abs/quant-ph/0211104**. (This is a long (70-page) and not-formally-published paper, now entirely superseded by published work and included only for historical purposes.).

———. (2003a). Everett and structure. *Studies in the History and Philosophy of Modern Physics* 34: 87–105.

———. (2003b). Everettian rationality: Defending Deutsch's approach to probability in the Everett interpretation. *Studies in the History and Philosophy of Modern Physics* 34: 415–439.

———. (2005). Language use in a branching universe. Unpublished manuscript; Available online from **http://philsci-archive.pitt.edu**.

———. (2006). Epistemology quantized: Circumstances in which we should come to believe in the Everett interpretation. *British Journal for the Philosophy of Science* 57: 655–689.

———. (2007). Quantum probability from subjective likelihood: Improving on Deutsch's proof of the probability rule. *Studies in the History and Philosophy of Modern Physics* 38: 311–332.

———. (2008). The interpretation of quantum mechanics. In D. Rickles (Ed.), *The Ashgate companion to contemporary philosophy of physics*, 197–261. Burlington, VT: Ashgate.

———. (2010). How to prove the Born rule. In Saunders, Barrett, Kent, and Wallace (2010), 227–263. Available online at **http://arxiv.org/abs/0906.2718**.

———. (2012). *The emergent multiverse: Quantum theory according to the Everett interpretation*. Oxford: Oxford University Press: 227–263.

Wallace, D., and C. Timpson (2007). Non-locality and gauge freedom in Deutsch and Hayden's formulation of quantum mechanics. *Foundations of Physics* 37(6): 951–955.

———. (2010). Quantum mechanics on spacetime I: Spacetime state realism. *British Journal for the Philosophy of Science* vol. 61: 697–727. Available online at **http://philsci-archive.pitt.edu**.

Wilson, A. (2010a). Macroscopic ontology in Everettian quantum mechanics. *Philosophical Quarterly* 60: 1–20.

Wilson, A. (2010b). Modality Naturalized: The Metaphysics of the Everett Interpretation. Ph.D. diss., University of Oxford.

Notes:

(1) Arguably this has changed, but only in the last decade or so, and more so in the UK than elsewhere. (Students occasionally ask me how the Everett interpretation is perceived outside Oxford; my flippant answer is that there is a significant divide between philosophers who do and do not take it seriously, and the divide is called the Atlantic Ocean.)

(2) This is largely anecdotal; see, however, Tegmark (1998).

(3) This simplifies slightly: it is frequently convenient—notably in cases involving symmetry—to define the space of states so that the mathematics-to-physics relation is many-to-one, and it is somewhat controversial in some such cases whether it *is* many-to-one (see, e.g., Saunders (2003) and references therein.) Such concerns are orthogonal to the quantum measurement problem, though.

(4) That is: what is the correct view of scientific theories—semantic or syntactic (cf. Ladyman and Ross (2007, 111–118) and references therein).

(5) In the case of dynamical-collapse theories, Tumulka (2006) has produced a relativistically covariant theory for *non-interacting* particles, but to my knowledge there is no dynamical-collapse theory empirically equivalent to any relativistic theory with interactions. There has been rather more progress in the case of hidden variable theories (perhaps unsurprisingly, as these supplement but do not modify the already-known unitary dynamics); for three

different recent approaches, see Dürr et al. (2004, 2005) (hidden variables are particle positions), Struyve and Westman (2006) (hidden variables are bosonic field strengths), and Colin (2003) and Colin and Struyve (2007) (hidden variables are local fermion numbers). As far as I know, no such approach has yet been demonstrated to be empirically equivalent to the Standard Model to the satisfaction of the wider physics community.

(6) Perhaps in *some* sense there are multiple interpretations of classical electromagnetism: perhaps realists could agree that the electromagnetic field is physically real but might disagree about its nature. Some might think that it was a property of spacetime points; others might regard it as an entity in its own right. I am deeply skeptical as to whether this really expresses a distinction, but in any case, I take it *this* is not the problem that we have in mind when we talk about the measurement problem.

(7) A move criticized on technical grounds by Foster and Brown (1988).

(8) Given that an “observer” is represented in the quantum theory by some Hilbert space many of whose states are not conscious at all, and that conversely almost any sufficiently large agglomeration of matter can be formed into a human being, it would be more accurate to say that we have a consciousness basis for all *systems*, but one with many elements that correspond to no conscious experience at all.

(9) In fact many adherents of many-minds theories (e.g., Lockwood and Donald) embrace this conclusion, having been led to reject functionalism on independent grounds.

(10) It would actually be a case of discarding all but one *set* of minds—one for each observer.

(11) Barbour (1999) might be an exception; so might Allori et al. (2009), though it is unclear if Allori et al. are actually advocating the interpretation rather than using it to illustrate broader metaphysical themes.

(12) For an elementary introduction, see, e.g., Kittel (1996); for a more systematic treatment see, e.g., Tsvelik (2003) or (old but classic) Abrikosov, Gorkov, and Dzyalohinski (1963).

(13) In exactly *what* sense is controversial, and the debate arguably overlaps(!) with others in mainstream metaphysics; see Saunders (2010) and Wilson (2010a, 2010b) for further discussion.

(14) It represents a departure from my position in Wallace (2006).

(15) See Wallace (2012) for details; I learned the argument from David Deutsch in conversation.

(16) For reasons of space I omit detailed discussion of the parallel tradition in Everettian quantum mechanics of identifying probability via long-run relative frequency (notably by Everett himself (1957) and by Farhi, Goldstone, and Gutmann(1989). I discuss this program in detail in chapter 4 of Wallace(2012); my conclusion is that it works about as well, or as badly, as equivalent classical attempts, though there is no direct Everettian analogue to the best-systems approach.

(17) A more precise way of stating both is that the program attempts to show that agents are rationally required to act as if mod-squared amplitude played the objective-probability role in David Lewis's *Principal Principle*; cf. Lewis (1980).

(18) See, the discussions in e.g., Bell (1981a) or Maudlin (2002).

(19) That the dynamics are thereby required to violate Lorentz covariance does not uncontroversially follow; cf. Myrvold (2002), Wallace and Timpson (2010), and Tumulka (2006).

(20) For a more detailed analysis—which gives a slightly different account of why the Everett interpretation is an exception to Bell's result—see Timpson and Brown (2002).

(21) Storrs McCall also explores these issues in developing his approach to quantum mechanics (see, e.g., McCall (2000)); that approach is related to, but not identical to, the Everett interpretation (and, insofar as it relies on an explicit and precise concept of branching without offering a dynamical explication of when branching occurs, arguably fails to solve the measurement problem.)

David Wallace

David Wallace studied physics at Oxford University before moving into philosophy of physics. He is now Tutorial Fellow in Philosophy of Science at Balliol College, Oxford, and university lecturer in Philosophy at Oxford University. His research interests include the interpretation of quantum mechanics and the philosophical and conceptual problems of quantum field theory, symmetry, and statistical physics.