

Philosophy of Quantum Mechanics

David Wallace*

February 26, 2022

Abstract

This is a general introduction to and review of the philosophy of quantum mechanics, aimed at readers with a physics background and assuming no prior exposure to philosophy. It is a draft version of an article to appear in the Oxford Research Encyclopedia of Physics.

If philosophy of physics has a central problem, it is the quantum measurement problem: the problem of how to interpret, to make sense of, perhaps even how to fix, quantum mechanics. Other theories in physics challenge our intuitions and our everyday assumptions, but only quantum theory forces us to take seriously the idea that there is no objective world at all beyond our observations — or, perhaps, that there are many. Other theories in physics leave us puzzled about aspects of how they are to be understood, but only quantum theory raises paradoxes so severe that leading physicists and leading philosophers of physics seriously consider tearing it down and rebuilding it anew. Quantum theory is both the conceptual and mathematical core of 21st-century physics, and the gaping void in our attempt to understand the worldview that 21st-century physics gives us.

Unsurprisingly, then, the philosophy of quantum mechanics is dominated by the quantum measurement problem, and to a lesser extent by the related problem of quantum nonlocality, and in this article I give an introduction to each. In section 1 I review the formalism of quantum mechanics and the quantum measurement problem. In sections 2–4 I discuss the three main classes of solution to the measurement problem: treat the formalism as representing the objective state of the system; treat it as representing only probabilities of something else; modify it or replace it entirely. In section 5 I review Bell’s inequality and the issue of nonlocality in quantum mechanics, and relate it to the interpretations discussed in sections 2–4. I make some brief concluding remarks in section 6.

A note on terminology: I use ‘quantum theory’ and ‘quantum mechanics’ interchangeably, to refer to the overall framework of quantum physics (containing quantum theories as simple as the qubit or harmonic oscillator and as complicated as the Standard Model of Particle Physics). I do not adopt the older

*Department of Philosophy/Department of History and Philosophy of Science, University of Pittsburgh; david.wallace@pitt.edu

convention (still somewhat common in philosophy of physics) that 'quantum mechanics' means only the quantum theory of particles, or perhaps even non-relativistic particles: when I want to refer to non-relativistic quantum particle mechanics I will do so explicitly.

1 The quantum formalism and the measurement problem

In general, scientific models of systems tell us what those systems are and how they behave; they explain the observed phenomena in terms of the dynamical processes of those systems; through doing so, they make predictions about the system's properties, and those properties can then be tested through observation and experiment. The quantum measurement problem arises because quantum mechanics appears to be different: it is not obvious how to understand the formalism at all, let alone how to understand it as telling us about the dynamics and properties of systems prior to those systems being observed, and although there is a rough and tacit consensus among physicists about how to use the theory in practice, it is notoriously difficult how to make clear theoretical sense of, far less justify, that consensus.

The problem is most obviously associated to what happens when a system is measured in the lab — hence 'measurement problem' — but equally apt names would be 'the problem of quantum-classical transition' (since it is through that transition that the theory makes observable predictions, whether or not a formal experiment is carried out) or 'the problem of interpretation' (since at issue is what the quantum formalism means and what kind of account it gives us of quantum systems).

1.1 The quantum formalism

To see how the measurement problem arises, let's briefly review the quantum formalism, which can be thought of as having four parts.

1. To each quantum system is assigned a *state space*: a complete, complex inner-product space (i.e. a Hilbert space). The states are represented by rays in this space: that is, by normalised vectors, up to a phase factor, representing the fact that $|\psi\rangle$ and $e^{i\theta}|\psi\rangle$ are physically equivalent. (So strictly the state space might better be identified as *projective Hilbert space*, the space of rays.)
2. The self-adjoint operators on the Hilbert space are taken to represent the physical quantities (sometimes called *observables*) that characterise the system, and at least some of those operators are explicitly treated as representing specific physically meaningful quantities, such as particle positions or the values of conserved quantities like momentum.

3. One dynamical variable, the *Hamiltonian* (identified with the system's energy) has a particular physical role: it determines the evolution of the system's state via the *Schrödinger equation*:

$$\frac{d}{dt} |\psi\rangle = -\frac{i}{\hbar} \widehat{H} |\psi\rangle. \quad (1)$$

4. An explicit rule — the *Born probability rule* — relates the quantum states and the observables to the result of experiments. The rule has two aspects, each of which applies to a measurement of a physical quantity represented by self-adjoint operator \widehat{O} for a system with state $|\psi\rangle$. Any such operator (ignoring some infinite-dimensional technicalities) can be written as

$$\widehat{O} = \sum_o o \widehat{\Pi}_o \quad (2)$$

where the sum is over all eigenvalues of \widehat{O} and $\widehat{\Pi}_o$ is the projector onto eigenstates with eigenvalue o . (In the simplifying special case where \widehat{O} is nondegenerate, this becomes

$$\widehat{O} = \sum_o o |o\rangle \langle o| \quad (3)$$

where $|o\rangle$ is the unique eigenstate of \widehat{O} with eigenvalue o .) Then:

- (a) The measurement process is in general indeterministic, and the *possible outcomes* of the measurement are the eigenvalues of \widehat{O} ;
- (b) The *probability* of outcome o obtaining is given by $\text{Pr}(o|\psi) = \langle \psi | \widehat{\Pi}_o | \psi \rangle$. In the nondegenerate case this simplifies to $\text{Pr}(o|\psi) = |\langle \psi | o \rangle|^2$.

The Born rule can equally be expressed as a rule for the *expectation value* of a measurement: If the dynamical quantity corresponding to \widehat{O} is measured, the expected value of the outcome is

$$\langle O \rangle_{|\psi\rangle} = \langle \psi | \widehat{O} | \psi \rangle.$$

(The reason that this is equivalent is that for any eigenvalue o of \widehat{O} , there is a function f_o such that $f_o(o) = 1$, $f_o(o') = 0$ for any other eigenvalue; the expected value of a measurement of $f_o(O)$ equals the probability of measuring \widehat{O} and obtaining o .)

It is through the Born rule — and only through the Born rule — that we connect the formalism to experiment and observation. Only via the Born rule can we connect the abstract formalism of state and observable with any claims about what values of observables will actually be measured when a quantum system is observed. The form of this connection is entirely novel, in that it builds the notion of measurement into the very formalism of the theory, and it is that novelty that gives rise to the measurement problem. We can see this clearly by comparing the quantum formalism with two classical formalisms.

1.2 Comparison: The formalism of classical mechanics

Firstly, consider the Hamiltonian formulation of classical mechanics. Here, there are clear analogs of the first three parts of the quantum formalism:

1. A classical system is assigned a state space (*phase space*), with states represented by points in this space.
2. Certain functions on phase space are assigned physical significance by taking them to represent the physical quantities that characterise the system (in particle mechanics, for instance, we identify certain coordinate functions with the components of position and momentum of each particle).
3. One dynamical variable, again called the Hamiltonian and again identified with the system's energy, determines the evolution of the system's state via Hamilton's equations.

But there is no need for an analog of the Born rule, since the the interpretation of the system is simple, transparent, and makes no mention of probabilities: the classical state assigns a unique value to each physical quantity, such that given a physical quantity represented by function O , and a system in state x , the value of the physical quantity is just $O(x)$. Indeed, since the physical quantities coordinatize the phase space, the classical state just is an assignment of values to all the physical quantities. In classical mechanics, whether some further physical process accurately measures the value of a physical quantity depends on the dynamics of that process, and is a matter for the designers of experimental apparatus: the formalism of the theory just says what values the quantities actually have, and 'measurements' are just processes that attempt to report those actually-possessed values.

In classical Hamiltonian mechanics, the system's state is *representational*: the state represents the physical properties of the system, and different states correspond to different ways the system might be. (An alternate name in the literature is *ontic*, borrowing from the philosophers' term *ontology*, meaning what exists, what there is.)

The representational notion is not the only notion of 'state' in classical mechanics. Another arises in classical *statistical* mechanics, the formalism of which might be summed up as:

1. A classical statistical state is not a point in phase space, but a probability distribution over it — so that the space of statistical states is not phase space, but the much larger space of such distributions.
2. Just as in Hamiltonian classical mechanics, certain functions on phase space are assigned physical significance by taking them to represent the physical quantities that characterise the system.
3. The statistical state evolves by the Liouville equation, which is just the equation for how a probability distribution over phase-space points changes with time if those points themselves evolve under Hamilton's equations.

4. Given a physical quantity represented by function O , and a statistical state ρ , the expectation value of O is

$$\langle O \rangle_\rho = \int d\mathbf{p} d\mathbf{q} \rho(\mathbf{p}, \mathbf{q}) O(\mathbf{p}, \mathbf{q}).$$

(This can readily be reformulated as a rule for the probability density of the system having a given value of O .)

Again, the first three parts of the formalism match the quantum case, and the Hamiltonian case: a state space, the representation of physical quantities by mathematical objects defined on that space, and a dynamical equation for states. The fourth part imposes the interpretation of the statistical state as a probability distribution over the physical quantities. In classical statistical mechanics, the state is not representational: it is *probabilistic*, encoding the fact that the system might have many possible values of its physical quantities, and assigning a probability to each possibility. (An alternate name in the literature is *epistemic*, borrowing from the philosophers's term *epistemology*, meaning the study of knowledge and reflecting the common interpretation of statistical-mechanical probability as ignorance of the true state.)

We can now ask: which version of state dynamics, mechanical or statistical-mechanical, applies to quantum mechanics? Or, put another way: is the quantum state representational or probabilistic? And the measurement problem arises because neither conception of state straightforwardly works in quantum mechanics.

1.3 Why the state doesn't seem to be representational

The problem with a representational concept of the quantum state is that *it is apparently incompatible with the Born rule*. Consider a measurement of some quantity represented by operator \hat{O} . It is reasonably simple to make the representational concept work if the system's state is an eigenstate of \hat{O} , say $|o\rangle$ with eigenvalue o . In that case, a measurement of \hat{O} returns o with 100% certainty, and so we could suppose that $|o\rangle$ represents a system with value o of that quantity. But what about a superposition $\alpha_1 |o_1\rangle + \alpha_2 |o_2\rangle$, where $|o_1\rangle, |o_2\rangle$ are eigenstates of \hat{O} with different eigenvalues o_1, o_2 — what value of the quantity does *that* state represent? It is common to say that it represents the system having an 'indefinite' value of the quantity, but it is not at all clear what that means — and more importantly, the Born rule tells us that if we measure that quantity, the result of the measurement is not 'the quantity has an indefinite value', but either 'the quantity has value o_1 ' or 'the quantity has value o_2 ', with some probability of either. So 'measurements', whatever they are, do not seem to return the actually-possessed value.

The problem can be sharpened by considering a measurement as a physical process (measurement devices, after all, are made of atoms, and atoms are subject to the Schrödinger equation). For definiteness, suppose we have a spin-half particle, and measure its spin along the z-axis. A very simple model of a

measurement device would have the device begin in some ‘ready’ state, $|\text{Ready}\rangle$, and evolve into some state $|\text{Records ‘Up’}\rangle$ if the particle initially has spin up in the z direction, and $|\text{records ‘Down’}\rangle$ if it has spin down. That is, the combined system of device and particle implements this dynamics:

$$\begin{aligned} |+_z\rangle \otimes |\text{Ready}\rangle &\longrightarrow |+_z\rangle \otimes |\text{Records ‘Up’}\rangle \\ |-_z\rangle \otimes |\text{Ready}\rangle &\longrightarrow |-_z\rangle \otimes |\text{Records ‘Down’}\rangle. \end{aligned} \quad (4)$$

But if the device is now used to measure a spin-half particle with an indefinite value of z -spin, from the linearity of the Schrödinger equation we know that the result is

$$\begin{aligned} &(\alpha_+ |+_z\rangle + \alpha_- |-_z\rangle) \otimes |\text{Ready}\rangle \\ &\longrightarrow \alpha_+ |+_z\rangle \otimes |\text{Records ‘Up’}\rangle + \alpha_- |-_z\rangle \otimes |\text{Records ‘Down’}\rangle. \end{aligned} \quad (5)$$

The measurement device (more precisely: the combined system of measurement device + particle) is in an indefinite state, neither definitely recording ‘Up’ nor definitely recording ‘Down’. It is extremely hard to understand what a state like this could represent, but more to the point, it flatly contradicts what the Born rule would predict, which is that the measurement interaction in this case should be indeterministic:

$$\begin{aligned} &(\alpha_+ |+_z\rangle + \alpha_- |-_z\rangle) \otimes |\text{Ready}\rangle \longrightarrow |+_z\rangle \otimes |\text{Records ‘Up’}\rangle \text{ (probability } |\alpha_+|^2) \\ &(\alpha_+ |+_z\rangle + \alpha_- |-_z\rangle) \otimes |\text{Ready}\rangle \longrightarrow |-_z\rangle \otimes |\text{Records ‘Down’}\rangle \text{ (probability } |\alpha_-|^2). \end{aligned}$$

(It is not difficult to see that this indefiniteness arises simply from the linearity of the Schrödinger equation: a more realistic model of the measurement would allow for the possibility of macroscopically many post-measurement states but would not change the basic result.)

A particularly vivid realisation of the problem arises when the device ‘records’ the measurement outcome via Schrödinger’s unfortunate cat: the cat is killed if the measurement outcome is ‘Up’, spared if it is ‘Down’. If the input state is a superposition of z -spin states, then according to the representational concept of state, the cat is in an indefinite state of alive and dead; according to the Born Rule, it is alive with some probability, dead with some other probability.


1.4 Why the state doesn’t seem to be probabilistic

The probabilistic conception of the quantum state, by contrast, handles macroscopic superpositions very naturally. On that conception, a state like

$$\alpha_+ |+_z\rangle \otimes |\text{Records ‘Up’}\rangle + \alpha_- |-_z\rangle \otimes |\text{Records ‘Down’}\rangle \quad (6)$$

just describes a situation where the measurement device has probability $|\alpha_+|^2$ of recording ‘up’ and probability $|\alpha_-|^2$ of recording ‘down’; similarly, a Schrödinger-cat state just represents a cat that might be alive and might be dead, not some weirdly indefinitely-alive creature. The problem with a probabilistic conception

Figure 1: Mach-Zender interferometer



of the quantum state is, instead, that *it is apparently incompatible with microscopic superposition and interference*. We can begin seeing this by asking how, on that conception, we should interpret a state like

$$\frac{1}{\sqrt{2}} (|+z\rangle + |-z\rangle) \tag{7}$$

— is it a system that is equally likely to have spin up or spin down in the z -direction, or a system that definitely has spin-up in the x direction? Quantum mechanics does not distinguish the two, but at least on the face of it, the probabilistic conception has to do so.

The problem gets sharper when we consider how an interference experiment plays out on this conception. Consider a Mach-Zender interferometer like the one in Figure 1, tuned so that detector ‘A’ always fires (that is: the two paths to A constructively interfere, while the two paths to B destructively interfere and cancel out). Schematically, we can represent the quantum physics as

$$|\text{input}\rangle \rightarrow \frac{1}{\sqrt{2}} (|\text{Path 1}\rangle + |\text{Path 2}\rangle)$$

$$\begin{aligned}
|\text{Path 1}\rangle &\rightarrow \frac{1}{\sqrt{2}} (|\text{Approaching A}\rangle + |\text{Approaching B}\rangle) \\
|\text{Path 2}\rangle &\rightarrow \frac{1}{\sqrt{2}} (|\text{Approaching A}\rangle - |\text{Approaching B}\rangle). \tag{8}
\end{aligned}$$

So if we block one path so that all photons have to go down the other, there will be an equal chance of detecting the photon at each of A and B, but if we allow it to go down either path, interference means that the photon is certainly detected at A. On (at least the most straightforward version of) the probabilistic conception, a superposition like

$$\frac{1}{\sqrt{2}} (|\text{Path 1}\rangle + |\text{Path 2}\rangle) \tag{9}$$

represents a photon that is either in Path 1, or in Path 2. But now we can reason as follows:

Either the photon is in Path 1, or in Path 2.

If it is in Path 1, it is 50% likely to be detected at each detector.

If it is in Path 2, it is also 50% likely to be detected at each detector.

—

So whichever path it is in, it is 50% likely to be detected at each detector.

We seemed to have used no more than elementary logic¹ and the basics of the probabilistic conception — and yet we reach a contradiction with experiment. We can reach a similar conclusion by noting that the detections can be shifted from ‘A with certainty’ to ‘B with certainty’ by changing the path length by 1/2 wavelength in either Path 1, or Path 2 — which seems to imply that something physically relevant is going on on both paths, on every run of the experiment.

1.5 The measurement problem and its solutions

We can now give a clean statement (following [5]) of the quantum measurement problem:

The quantum measurement problem: How do we understand quantum theory in general, and the quantum state in particular, given that we are led to paradox if we think of it as *either* representational *or* probabilistic?

This is far from the only statement in the literature. (Influential alternatives include [6, 7, 8].) Philosophers of physics often take for granted that the quantum

¹There is a long tradition of considering whether ‘elementary logic’ itself needs to be revised in light of quantum mechanics. The idea was mooted in philosophy by Quine [1] and Putnam [2], and led to the field of quantum logic; despite interesting technical results, though, that field has largely fallen from favor in contemporary discussions of the measurement problem. Recent reviews are [3, 4].

state is supposed to represent the world, and phrase the measurement problem as: how can we reconcile this with macroscopic superpositions? Less commonly, one finds versions that take for granted that the quantum state is probabilistic, and phrase the measurement problem as: what is the probability space on which quantum probabilities are defined? The virtue of the statement I give here is that it avoids confusing the problem ('how can we understand the quantum state?') with aspects of the solution ('given that the state represents the world, how do we make sense of macroscopic superpositions?')

So: how do we *solve* the problem? Physical practice finesses it by moving inconsistently between the representational and probabilistic conceptions: in interference experiments we treat the state as representational; at macroscopic scales, or when we make a measurement, we shift to treating it as probabilistic. We could try to resolve the inconsistency by defending either the representational or probabilistic conception of the quantum state against the criticisms I have given — and as we will see, there are substantive defenses available for both the conceptions. For the second in particular, a popular move — going back to the founders of the field — is the *instrumental* (or *pragmatic*) conception of the quantum state, which I treat here as a subcategory of the probabilistic conception. On this conception, we abandon any attempt to extract understanding of the underlying physics from quantum dynamics and the quantum state, and instead treat the state as an instrumental tool to predict and manipulate other physical systems. The urgent issue for such approaches is to tell us how those 'other physical systems' are to be described, if not using quantum physics.

The remaining conception of quantum mechanics is *revisionary* (or *modificatory*). Adherents of this approach take at face value the apparent contradictions embedded in the representational and probabilistic conceptions of the state, and conclude that quantum mechanics is wrong, or at least not completely right; they seek modified versions of quantum mechanics that avoid the measurement problem. Ideas of this kind again go back to the dawn of quantum mechanics, to Dirac and von Neumann's postulate of wavefunction collapse upon measurement (indeed, sometimes this postulate is actually counted as part of the quantum formalism), but this approach uses an unanalysed concept of 'measurement' that most have found unsatisfactory or even in contradiction with the known microphysics of measurement devices, and contemporary advocates of the revisionary conception have looked to develop more microphysically principled modifications. Several are known, at least in the domain of nonrelativistic particle mechanics — extending any such approach to quantum field theory, and carrying out decisive empirical tests that distinguish the revisions from unmodified quantum mechanics, remain open problems.

In the next three sections I will consider each of these three conceptions of quantum mechanics: representational, probabilistic (including instrumental/pragmatic), and revisionary. A great deal has been learned about each over the years, even if the community remains far from consensus on which is preferable.

2 The representational conception of the quantum state: decoherence, Everett, shut-up-and-calculate

The distinctive feature of representational approaches to quantum theory is that the unitary dynamics of quantum theory suffices to analyze any quantum problem, including those in which measurement and observation occur. The most pressing problem for these approaches is then how to provide this unitary analysis of measurement, and there is a fairly wide consensus that this requires decoherence theory, which I review briefly in section 2.1 (see also [ORE decoherence article] and [9]). In the remainder of this section I consider whether decoherence suffices by itself, and if not then what else is needed.

2.1 Decoherence and the quantum-to-classical transition

As I observed above, physical practice attempts to avoid the measurement problem by inconsistently moving between different conceptions of the state: representational when we need to think about interference or entanglement, probabilistic when we need to think about measurements and observations. It is not obvious why this works: why couldn't interference or entanglement effects persist at the scale of macroscopic observation and lead to outright contradiction in our physics?

The answer — known in outline since the dawn of quantum theory, but enormously clarified since the development of decoherence theory in the late 20th century² — is that it is difficult to place large physical systems in superpositions without them becoming uncontrollably entangled with their environment, in a way which makes superpositions impossible to distinguish from probabilistic mixtures. As an extreme example, suppose that we somehow placed Jupiter in an equally-weighted superposition of two different positions X and Y, millions of miles apart. To distinguish this quantum state from the merely probabilistic state ‘Jupiter has a 50% chance of being at X and a 50% chance of being at Y’, we would need to carry out some kind of interference experiment, some experiment that would give different results for

$$|J+\rangle = \frac{1}{\sqrt{2}} (|\text{Jupiter at X}\rangle + |\text{Jupiter at Y}\rangle) \quad (10)$$

than for

$$|J-\rangle = \frac{1}{\sqrt{2}} (|\text{Jupiter at X}\rangle - |\text{Jupiter at Y}\rangle) \quad (11)$$

But now consider a single photon, coming in from interstellar space and passing through X. If Jupiter is at X, the photon will be affected (reflected, let's assume). If it is at Y, the photon will pass through X undetected. So the joint quantum

²See, e.g., [10, 11] for detailed reviews.

state of photon and Jupiter is now

$$\begin{aligned}
 |J\pm; \text{photon}\rangle &= \frac{1}{\sqrt{2}} |\text{Jupiter at X}\rangle \otimes |\text{photon scattered}\rangle \\
 &\pm \frac{1}{\sqrt{2}} |\text{Jupiter at Y}\rangle \otimes |\text{photon unaffected}\rangle. \quad (12)
 \end{aligned}$$

And it is now impossible to distinguish $|J+; \text{photon}\rangle$ from $|J-; \text{photon}\rangle$ — and hence, impossible to distinguish Jupiter-in-superposition from Jupiter-in-probabilistic mixture — by any physical process involving Jupiter alone. We could see this formally by constructing the density operator for Jupiter in the absence of the photon and noting that it is the same for both states, but it is simpler just to see that a phase transformation made on the photon alone (flipping the phase of the deflected photon, leaving the unaffected photon alone) is enough to move us from one state to the other. So an interference experiment that confirmed that Jupiter is really in a superposition would have to be an experiment operating jointly on Jupiter and the photon.

But of course, interstellar space contains more than one photon. Countless numbers scatter off Jupiter every second, along with dust motes, cosmic rays, and all manner of other detritus, and every single one of them will get entangled with Jupiter’s position just as our first photon does. The only way to carry out that interference experiment on Jupiter is to make it a joint interference experiment in which we manipulate both Jupiter and the vast number of degrees of freedom of Jupiter’s environment. And even if, impossibly, we were actually to capture that environment, it would no more be possible to perform the experiment than it would be to turn back time in a complex macroscopic system by reversing the velocities of all the constituents.

This is the process of *environment-induced decoherence*. Its hallmarks are that a physical system’s environment is constantly and redundantly measuring it with respect to some basis (usually — as in the case of Jupiter — the position basis, or in some coarse-grained combination of position and momentum), and that as a consequence:

- no remotely realistic interference experiment could distinguish between a superposition with respect to that basis, and a mere probabilistic mixture of basis terms;
- no remotely realistic preparation process could *prepare* the system in a superposition with respect to that basis, without it instantly and irreversibly becoming entangled with its environment;
- the internal dynamics of the system itself (e.g., Jupiter’s movement through space) cannot involve interference between terms in that basis, and so will be formally indistinguishable from some deterministic or stochastic process.

Of course, systems don’t need to be as big as Jupiter to decohere. Even dust motes are very rapidly decohered by the atmosphere or by ambient radiation, and the process only gets quicker for larger systems. Nor does the

‘environment’ doing the decohering have to be external to the system: the low-wavelength vibrational modes in an iron bar can be decohered by the bath of shorter-wavelength phonons in that same bar; the various cells and tissues of Schrödinger’s poor cat swiftly and redundantly record whether it lived or died.

Decoherence has by now been studied extensively by a variety of theoretical methods and in a variety of experimental contexts. Much remains to be learned about it, but there is not much disagreement that it is a ubiquitous feature of unitary quantum dynamics for large, complex systems, and that at least ‘for all practical purposes’ (to use John Bell’s [8] intentionally-derisive phrase) it explains why we can get away with ignoring interference and treating the quantum state probabilistically when we apply quantum theory to those systems.

2.2 The Everett interpretation: many worlds?

Nonetheless decoherence alone does not appear to solve the measurement problem. It tells us that a probabilistic interpretation of the state can work in certain large-scale situations, but that interpretation is still unavailable at the microscopic level. And decoherence is by its nature an emergent, approximate process. It does not define a sharp boundary beyond which interference disappears, but only a fuzzy boundary beyond which it is extremely small. And although it tells us that macroscopic superpositions like Schrödinger’s cat will rapidly get entangled with their environment, that just seems to make the superposition still more macroscopic (now cat *and* environment are in a superposition). The problem remains: we don’t observe systems in superpositions *at all*.

Or do we? Following Hugh Everett [12], let’s see what happens if we model observation itself — say, observation of the cat — within unitary quantum mechanics. If I observe a dead cat, I go into a quantum state (more precisely, one of a large class of quantum states) that could be schematically written as $|I \text{ see dead cat}\rangle$. Similarly, my observation of a live cat puts me in a quantum state $|I \text{ see live cat}\rangle$. But then the linearity of quantum mechanics tells us that if I observe a cat in a superposition of alive and dead, the joint state of me and the cat after the observation looks something like

$$\begin{aligned} |\psi\rangle &= \frac{1}{\sqrt{2}} |\text{Dead cat}\rangle \otimes |I \text{ see dead cat}\rangle \\ &+ \frac{1}{\sqrt{2}} |\text{Alive cat}\rangle \otimes |I \text{ see alive cat}\rangle. \end{aligned} \quad (13)$$

Importantly, there is no term in the superposition that we would write schematically as

$$|I \text{ see weird indefinite cat}\rangle.$$

According to unitary quantum mechanics itself, the result of measuring a macroscopic superposition would not be to experience something bizarre; it would be to end up in a superposition of having two mundane, non-bizarre experiences.

So how *are* we to understand a state like (13)? The idea which de Witt [13] and many others took from Everett (it’s a matter of controversy [14, 15] whether

it was Everett’s own interpretation) is that (13) is to be interpreted as *multiplicity*: as two (or two families of) observers, one seeing a dead cat, one seeing a live cat — and indeed, as two cats, one dead and one alive. On a broadly physicalist interpretation of the human mind, experiences and thoughts depend on physical processes in the brain, and these processes are multiply present, so the quantum state describes multiple, independent observers, each with their own definite experiences. And as more and more observers and other physical systems get entangled with the original cat-observer pair, the quantum state quickly ends up as something like

$$\begin{aligned}
 |\psi\rangle &= \frac{1}{\sqrt{2}} |\text{Dead cat}\rangle \otimes |\text{Everyone on Earth records ‘dead cat’}\rangle \\
 &+ \frac{1}{\sqrt{2}} |\text{Alive cat}\rangle \otimes |\text{Everyone on Earth records ‘live cat’}\rangle. \quad (14)
 \end{aligned}$$

In short order we have not just two observers, but two Earths, and before long two Solar systems. Hence the more popular name for Everett’s interpretation of quantum mechanics: the ‘many-worlds theory’.

The basic content of the Everett interpretation³ is:

1. The quantum state is representational: different states correspond to different physical systems, not to different levels of knowledge about the system.
2. Unitary quantum mechanics is complete.
3. Correctly understood, unitary quantum mechanics describes a branching multiplicity of approximately-classical worlds.

Given the frequency with which quantum theory magnifies microscopic superpositions up to classical scales (any flickering fluorescent light-bulb does so; so do classically chaotic processes), there really are many, *many* worlds — the Everett interpretation describes a Universe vastly more complex and highly structured than what we see around us. (Although of course this complexity was not added by hand: mathematically it is present in the unitarily-evolving state, however we interpret that state.) Many critics have been regarded such claims as grandiose, radically violating Ockham’s Razor and committing us unscientifically to the existence of unobservable entities. Defenders have responded that Ockham’s Razor is concerned with the simplicity of theories rather than crudely with the number of entities they posit, that we don’t regard the vastness of the astronomical Universe as a reason to reject astrophysics and cosmology, that we are justified in accepting other worlds because of the essential explanatory role they play in well-confirmed science, and that in any case their unobservability is a matter of degree and that degree changes every time we test the superposition principle in new circumstances.

³For detailed presentations see [16, 17].

More productive criticisms⁴ of the Everett interpretation have mostly fallen into two classes, known collectively as the ‘preferred basis problem’ and the ‘probability problem’. The preferred basis problem asks: if superpositions are to be understood as parallel worlds, with respect to which basis does this take place? Why is a state like (13) correctly described as two worlds in each of which the cat is definite, rather than one world with an indefinite cat in? Doesn’t unitary quantum mechanics need to be supplemented with some explicit rule that tells us which the preferred basis is — and in doing so, doesn’t it after all become a revisionary theory? The probability problem asks: since unitary quantum mechanics is deterministic, how is it to be reconciled with the probabilistic nature of our evidence for quantum theory? And how do we understand unequally weighted superpositions: if

$$\begin{aligned}
 |\psi\rangle &= \alpha |\text{Dead cat}\rangle \otimes |\text{I see dead cat}\rangle \\
 &+ \beta |\text{Alive cat}\rangle \otimes |\text{I see alive cat}\rangle
 \end{aligned}
 \tag{15}$$

describes a living cat and a dead cat irrespective of α and β , what becomes of the Born rule interpretation of $|\alpha|^2$ and $|\beta|^2$ as probabilities? Shouldn’t the probabilities be 50% irrespective of α and β , and in contradiction of the evidence for the theory?

It is fairly widely accepted that decoherence provides a solution to the preferred-basis problem. It does not provide a precise definition of the preferred basis, nor should it be expected to, because the ‘worlds’ or ‘branches’ in the Everett interpretation are not part of its fundamental structure but are high-level, emergent entities. The core physical idea here [21, 22] is that the decoherence-preferred basis is the one in which we can find autonomous, robust higher-level dynamics - fluid dynamics and the like in the first instance, but ultimately the still-higher-level processes of biochemistry and biology. What makes the decomposition into ‘dead’ and ‘alive’ branches is that with respect to that decomposition - but not with respect to wildly nonclassical ones — we find those autonomous high-level processes playing out multiply and independently. More philosophical work in the 2000s [23, 24] incorporated this approach to quantum branches into a more general framework of emergence and higher-level science in which emergent entities in general — phonons, tigers, mountains — are to be understood in this structuralist way (so that the Everett interpretation just arises from quantum theory when we apply already-understood principles of emergence to this specific case).

(The preferred-basis problem continues to concern some authors who object that no *emergent* account of the macroscopic world can be understood without a clear account of the *microscopic* world, and that unitary quantum mechanics fails to provide this account. According to this position (often called ‘primitive ontology’ [25] or ‘primary ontology’ [26]) a physical theory should be specified by first giving an ontology (a list of those entities which the theory describes, preferably understandable as local entities like particles or classical fields) and only then specifying a dynamics for those entities. Unitary quantum mechanics

⁴For details see [17, 18, 19, 20] and references therein.

certainly fails this requirement, as do most physical theories of the 20th and 21st century; whether this is an indictment of physics or of the primitive-ontology strategy lies beyond the scope of this review.)

The probability problem remains largely intact when decoherence is considered, and has been the main locus of controversy about the Everett interpretation in 21st-century discussions. But decoherence does clarify what the conditions are by which the mod-squared amplitudes have the right *formal* properties to act as probabilities, in particular when they can be used to define transition-probabilities in a consistent way. (In the ‘consistent histories’ [27, 28] or ‘decoherent histories’ [21, 29] approach to decoherence, this so-called ‘consistency condition’ is taken as a formal definition of decoherence.) And so it shows that the problem is essentially philosophical: given that some quantity in our mathematically-stated theory has the right mathematical properties to count as probability, what more needs to be said to justify the claim that it *emph*is probability. A range of arguments has been advanced as to why there is indeed a problem here; defenders of the Everett interpretation have variously argued (i) that formal properties are sufficient and that any further issue is part of the general problem of understanding probability and is not specific to this context [30, 31, 32],[16, ch.4], or (ii) that the Born rule can be derived in unitary quantum mechanics by various frequency-based [33], symmetry-based [16, ch.4], decision-theory-based [34, 35, 16], entanglement-based [36] or other [37, 38, 39] methods. (For criticism see, e.g., [40, 41, 42]).

2.3 Everett without the worlds

Although decoherence theory has been associated with Everett’s ideas since its inception (see, e.g., [43]), also since its inception there have been attempts (mostly, though not exclusively, by physicists) to get the advantages of decoherence, and of Everett’s insights about measurement processes in unitary quantum mechanics, without the (arguably) unattractive commitment to a plethora of unobservable ‘worlds’. Prominent examples of approaches along these lines include (but are by no means limited to):

- Griffiths’ and Omnes’ ‘consistent-histories’ approaches [27, 44, 45], in which (roughly) each precisification of the decoherence basis provides a legitimate probability space but it is impermissible to describe physics using more than one such space at a time.
- Gell-Mann and Hartle’s decoherent-histories approach [29, 22] (and Halliwell’s related approach [21, 46]) which is very closely related to the Everett interpretation as I described it above but is at least officially not committed to a multiplicity of really-existing worlds.
- Zurek’s decoherence-based interpretation [47], which attempts to combine aspects of the Copenhagen interpretation with Everettian quantum mechanics.

- Rovelli’s relational quantum mechanics [48, 49], which abandons the idea of giving any description of a physical system in isolation and instead incorporates as a central principle that any system can be described only from the perspective of another system (without any requirement, however, that this second system be a conscious observer or similar).

These approaches have had less attention in the philosophical and foundational literature than the ‘mainstream’ Everett interpretation (for critical discussion of consistent- and decoherent-histories approach, see [50, 51, 52, 53, 54]; for critical discussion of relationalism, see [55, 56]). The challenge for them is to state a view that (a) is sufficiently clear and unambiguous to be an intelligible solution to the measurement problem, without (b) collapsing into to Everett interpretation. (So, for instance, Gell-Mann and Hartle officially disavow many worlds, but it is difficult to find any actual difference between their discussions of decoherence, quasi-classicality, and emergence and those of self-confessed ‘many-worlders’, raising the suspicion that the difference is not much more than verbal). Those approaches that do seem to meet this challenge mostly do so at the price of abandoning an unequivocal, objective, third-party-available description of a physical system; of course, their advocates (like Griffiths or Rovelli) would see this not as a price but as a substantive discovery about the nature of science and/or the world.

3 The probabilistic conception of the quantum state: toy models and no-go theorems

There are really two different — though related — probabilistic conceptions of the quantum state, which share the idea that quantum states simply code probabilities but differ as to what the probabilities are probabilities of. In what we might call ‘objective’ or ‘epistemic’ versions, there is an underlying sub-quantum reality, and an as-yet-unknown physical theory which describes that reality; quantum theory encodes our partial information about the underlying facts. On the more ‘operationalist’ or ‘measurement-first’ versions, quantum states directly encode outcomes of measurements or operations, or perhaps our beliefs or expectations about them, or the best advice available to us as to what to expect from them. I discuss the two versions in order.

3.1 The ontic-model framework and ψ -epistemic hidden variable theories

The *ontic models* framework [57] provides a very general way to fill in the details of the more ‘realist’ probabilistic conception of the quantum state.. There is an underlying state space — \mathcal{S} , say — each point of which represents a possible actual, underlying state of the system, and a dynamics on \mathcal{S} . On any particular run of an experiment, the results of the experiment are determined by the underlying state (this determination could itself be probabilistic, say

if the underlying dynamics is stochastic). Concretely, we can assume a set of measurement probabilities \mathcal{M}_A , indexed by quantum observables A : $\mathcal{M}_A(a|x)$ is the probability that a measurement of A gives result a , given that the underlying state is x .

There is then a probability distribution on \mathcal{S} , to be interpreted in the same way we interpret statistical-mechanical probability distributions,⁵ In an experimental context, this probability distribution will be determined by the physics of the state-preparation process. For each quantum state $|\psi\rangle$, we can define a probability distribution $\rho_{|\psi\rangle}$, so that $\rho_{|\psi\rangle}$ is the probability that the system's underlying state is x , given that the system is prepared by some method which, according to quantum mechanics, would give it quantum state $|\psi\rangle$.

If we then carry out an experiment whose description in the language of quantum mechanics would be ‘prepare system in state $|\psi\rangle$, measure A ’ then the probability of getting an outcome a is fully determined by the two distributions $\rho_{|\psi\rangle}$ and \mathcal{M}_A :

$$\Pr(A = a) = \sum_{o \in \mathcal{S}} \mathcal{M}_A(o) \rho_{|\psi\rangle}(o).$$

And to reproduce the predictions of quantum mechanics, this will need to equal the Born-rule probability of that same outcome: that is, an ontic model for a quantum-mechanical system must satisfy

$$\sum_{o \in \mathcal{S}} \mathcal{M}_A(o) \rho_{|\psi\rangle}(o) = |\langle \psi | \hat{\Pi}_a | \psi \rangle| \quad (16)$$

for all $|\psi\rangle$, A and a . (The summations would need to be generalized to integrals in case of a continuum state space.)

Theories of this kind were once called *hidden-variable theories*, the idea being that the underlying — ‘hidden’ — physics is the physics of the variables represented by points in \mathcal{S} , and quantum mechanics is obtained from the hidden-variable theories in a manner analogous to how statistical mechanics is obtained from classical mechanics. More recently, ‘hidden-variable theory’ has come also to be used for theories which retain the quantum state in their microphysics but supplement it with additional degrees of freedom (this is an example of the revisionary approach; we will discuss it more in sections 4.3.–4.4). Theories of the current type are now called *ψ -epistemic* hidden-variable theories, in reference to the idea that the quantum state is a probabilistic summary of our partial knowledge and does not represent some independent part of reality.

3.2 Observable-based hidden variable theories

An extremely natural choice for the hidden variables is to suppose that they assign actual values to the various quantum observables (and then measurements of observables just return those actual values). In an *observable-based* hidden-variable theory like this, we suppose that at least some of the observables are

⁵Not that it there is any consensus on how that should be done; see [reference ORE top-level article on statistical mechanics].

represented by functions on the state space \mathcal{S} of the hidden variables. I write $\mathcal{V}_A : \mathcal{S} \rightarrow \mathbf{R}$ for the function corresponding to observable A , so that if the underlying state is x , the value of A is $\mathcal{V}_A(x)$. In a *complete* observable-based hidden variable theory, a point in \mathcal{S} encodes the values of all the quantum observables, and then condition (16) for the viability of an ontic model reduces to

$$\int dx \mathcal{V}_A(x) \rho_{|\psi\rangle}(x) = |\langle \psi | \hat{A} | \psi \rangle|^2 \quad (17)$$

for all $|\psi\rangle$ and A . This appears to have been Einstein's own conception of quantum mechanics (see [57]); it is extensively discussed in [58].

There are concrete models that reproduce some fragments of realistic quantum theories in this way. A good example is *Gaussian quantum mechanics* [59], the sector of nonrelativistic quantum mechanics we get by (i) restricting attention to Gaussian wave-packet states (or incoherent mixtures of those states), and (ii) requiring the potential function in the theory to be quadratic (so that Gaussian states remain Gaussian). (With further restrictions, the theory can also incorporate at least some classes of repeated measurements.) It is reasonably straightforward to interpret this sector of the theory probabilistically, so that there are underlying classical trajectories in phase space and the quantum state just represents our ignorance of those trajectories; it is also possible to construct Gaussian quantum mechanics directly from classical mechanics by adding an epistemic restriction on how much we can know about the classical microstate. And the resultant model can reproduce a number of quantum-mechanical phenomena, such as teleportation and the no-cloning theorem which at first sight might have appeared intrinsically quantum-mechanical.

However, there are strong no-go theorems which make it quite unlikely that an observable-based hidden-variable theory can be found for the whole of non-relativistic quantum mechanics, or indeed for any quantum-mechanical system with a Hilbert space of three or more dimensions. (Concrete models are known in two dimensions.) These results arise from two classes of assumptions: that the hidden-variable theory is *complete*, which is to say that \mathcal{V}_A exists, and (17) holds, for *any* observable A , and that the algebraic relations between observables are mirrored at least to some degree by the functions \mathcal{V}_A .

The strongest such relation that could be assumed is

$$\mathcal{V}_{f(A,B)}(x) = f(\mathcal{V}_A(x), \mathcal{V}_B(x)). \quad (18)$$

For instance, if the Hamiltonian is given by $H = p^2/2m + U(q)$, then this would require that

$$\mathcal{V}_H(x) = \mathcal{V}_p(x)^2/2m + U(\mathcal{V}_q(x)). \quad (19)$$

von Neumann demonstrated [60] that no complete observable-based hidden-variable theory can satisfy this relation (even in two dimensions); indeed, that none can satisfy the linear restriction of this relation,

$$\mathcal{V}_{\lambda A + \mu B} = \lambda \mathcal{V}_A + \mu \mathcal{V}_B. \quad (20)$$

John Bell [61] sharply criticized relations of this kind, however, in part on the grounds that they relate *non-commuting* observables, which according to the rules of quantum mechanics cannot be measured simultaneously: we cannot then operationalize the claim that measured values of (e.g.) position, momentum and energy should be related via (19). Whether or not this would be a reasonable objection to von Neumann's proof in the context of observable-based hidden variable theories is moot (Bell raised the objection in the context of the de Broglie-Bohm theory, which is not a hidden-variable theory of this kind) since in any case he, and Kochen and Specker [62], demonstrated that even a much weaker relation suffices to ground a no-go result.

Specifically, let us describe a hidden-variable theory as *non-contextual* if it satisfies (18) *whenever \hat{A} and \hat{B} commute*. In this circumstance, it *is* possible to measure A and B simultaneously, and indeed one standard method of measuring $f(A, B)$ is just to measure A and B and then calculate $f(A, B)$. In fact, since whenever \hat{A} and \hat{B} commute there is some \hat{C} , and some functions f_1, f_2 , such that $\hat{A} = f_1(\hat{C}), \hat{B} = f_2(\hat{C})$, we can rewrite the requirement of non-contextuality as

$$\mathcal{V}_{f(A)} = f(\mathcal{V}(A)). \quad (21)$$

And now non-contextuality looks almost trivial: what is it, after all, to measure A^2 , if not to measure A and square the result? Yet the Bell-Kochen-Specker theorem demonstrates that in a Hilbert space with 3 or more dimensions, there can be no complete and non-contextual observable-based hidden-variable theory.

If we tighten the conditions of the theorem so as to require a 4-dimensional Hilbert space, it becomes simple enough to demonstrate here (following [63]). Consider two spin-half particles, and let $\hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z$ be the usual sigma-matrices for a spin-half particle, satisfying $\sigma_i \sigma_j = i \epsilon_{ijk} \sigma_k$ (ϵ_{ijk} is the usual completely-antisymmetric 3-tensor). Then we can set up the following grid of observables for the pair of particles:

$\hat{\sigma}_x \otimes \hat{1}$	$\hat{1} \otimes \hat{\sigma}_x$	$\hat{\sigma}_x \otimes \hat{\sigma}_x$
$\hat{1} \otimes \hat{\sigma}_y$	$\hat{\sigma}_y \otimes \hat{1}$	$\hat{\sigma}_y \otimes \hat{\sigma}_y$
$\hat{\sigma}_x \otimes \hat{\sigma}_y$	$\hat{\sigma}_y \otimes \hat{\sigma}_x$	

We can easily verify that in each fully-filled-in row, and each fully-filled-in column, the first two operators commute and their product is the third operator.

Suppose there is a complete, non-contextual hidden-variable theory for the two spin-half particles. Then each hidden variable must assign some value to each operator in the table, and that value must be an eigenvalue of the operator: that is, it must be ± 1 . Fix some choice of hidden variables, and let the values they assign to the operators in the top-left four boxes be a, b, c, d , like so:

a	b	
c	d	

Then by non-contextuality, the remaining four values are fixed, so we have

a	b	ab
c	d	cd
ac	bd	

Now, the two operators on the bottom row also commute. So by non-contextuality, the hidden variables must assign value $abcd$ to their product. And the two operators in the right-hand column also commute, so the hidden variables must assign $abcd$ to *their* product. But

$$(\hat{\sigma}_x \otimes \hat{\sigma}_y)(\hat{\sigma}_y \otimes \hat{\sigma}_x) = +\hat{\sigma}_z \hat{\sigma}_z \quad (22)$$

and

$$(\hat{\sigma}_x \otimes \hat{\sigma}_x)(\hat{\sigma}_y \otimes \hat{\sigma}_y) = -\hat{\sigma}_z \hat{\sigma}_z. \quad (23)$$

So we are assigning value $abcd$ both to $\hat{\sigma}_z \hat{\sigma}_z$ and to $-\hat{\sigma}_z \hat{\sigma}_z$. Again using non-contextuality, we get $abcd = -abcd$. But each of a, b, c, d must equal ± 1 , so this is impossible.

What do we learn from the Bell-Kochen-Specker theorem? That at most a subset of the quantum-mechanical observables (the ‘non-contextual observables’) of the system can correspond to objective features of the system in the usual sense. For the others, their measured value depends on the context of measurement, which is to say that what is being measured is not a property of the system alone but some joint feature of the system and the measurement process. If we want to hold on to the idea of an observable-based hidden-variable theory as simply encoding already-possessed properties of the system, we instead need to drop completeness, so that only some observables actually correspond to possessed values.

In this sense, incompleteness and contextuality are two sides of the same coin: at most a subset of observables have non-contextual values; the processes that quantum physicists call ‘measurements’, when applied to the other observables, do not in the literal sense ‘measure’ anything. And in fact, the restrictions on the size of that subset are severe [64, 65]: if \hat{O} is a nondegenerate observable, then for a generic state $|\psi\rangle$ the only noncontextual observables are functions of \hat{O} .

Of course, a probability distribution over the outcomes of measurements of any one observable severely underdetermines the quantum state: if $\{|o\rangle\}$ are the eigenstates of \hat{O} , then

$$|\psi_1\rangle = \sum_o \lambda_o |o\rangle$$

and

$$|\psi_1\rangle = \sum_o \lambda_o e^{i\theta_o} |o\rangle$$

determine the same probability distribution. But since the phase information encoded in the θ_o may be dynamically relevant — even to the probability distribution over measurements of \hat{O} at later times — this tells us that the non-contextual hidden variables are not the whole story, and hence (in the terminology of section 3.1) that no observable-based hidden variable theory can be fully ψ -epistemic. The quantum state is not *simply* probabilistic; at least some features of it encode features of the system directly.

In fact, one can go further than this. The Pusey-Barrett-Rudolph (PBR) theorem establishes that under weak assumptions, essentially *all* features of the quantum state directly encode features of the system: no part of the state can be understood purely probabilistically, in terms of the probabilities of underlying properties of the system.

3.3 The PBR theorem and related results

The PBR theorem [66] is proved in the very general ontic-models framework of section 3.1. PBR (following [57]) assert that for the quantum state to be at least partially epistemic (that is, probabilistic), there ought to exist distinct quantum states $|\psi_1\rangle, |\psi_2\rangle$ which are compatible with the same ontic state — that is, there should be some ontic state x , or set X of such, that is given nonzero probability by both $|\psi_1\rangle$ and $|\psi_2\rangle$: $\rho_{|\psi\rangle}(X) > 0$ for $|\psi\rangle = |\psi_1\rangle$ and for $|\psi\rangle = |\psi_2\rangle$. If this does not hold, then (modulo some measure-theoretic subtleties if there are continuously many ontic states), there would be a unique way to recover the *quantum* state from the *ontic* state.

Put another way, to each quantum state we can assign its ‘ontic support’: the set of ontic states to which it gives nonzero probability. (Again, there are measure-theoretic subtleties here if there are continuously many ontic states.) The theory is ψ -epistemic to the degree that quantum states have overlapping ontic support: if no two quantum states have overlapping ontic support, the theory is in no way epistemic.

Here is how the theorem works: suppose we have a collection of identical systems each of which has been *independently* prepared either in state $|\psi_1\rangle$ or in state $|\psi_2\rangle$, with $|\langle\psi_1|\psi_2\rangle| = \cos\theta$. Without loss of generality, we can choose an orthonormal pair of states $\{|0\rangle, |1\rangle\}$ for each system such that $|\psi_1\rangle = |0\rangle$ and $|\psi_2\rangle = \cos\theta|0\rangle + \sin\theta|1\rangle$, and confine our attention to the 2-dimensional subspace of each system spanned by $\{|0\rangle, |1\rangle\}$.

To illustrate the general idea (following PBR’s own account), consider the special case where there are just 2 systems and where $\cos\theta = 1/\sqrt{2}$, so that $|\psi_1\rangle = (1/\sqrt{2})(|0\rangle + |1\rangle)$, and define $|\pm\rangle = (1/\sqrt{2})(|0\rangle \pm |1\rangle)$. There are then four ways in which the two systems might have been prepared:

- I. System 1 in state $|\psi_1\rangle$, system 2 in state $|\psi_1\rangle$ — joint state $|00\rangle$.
- II. System 1 in state $|\psi_1\rangle$, system 2 in state $|\psi_2\rangle$ — joint state $|0+\rangle$.
- III. System 1 in state $|\psi_2\rangle$, system 2 in state $|\psi_1\rangle$. — joint state $|+0\rangle$.
- IV. System 1 in state $|\psi_2\rangle$, system 2 in state $|\psi_2\rangle$ — joint state $|++\rangle$.

Now suppose the joint system is measured in the following orthonormal basis of entangled states:

$$|I\rangle = \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)$$

$$\begin{aligned}
|II\rangle &= \frac{1}{\sqrt{2}}(|0-\rangle + |1+\rangle) \\
|III\rangle &= \frac{1}{\sqrt{2}}(|+1\rangle + |-0\rangle) \\
|IV\rangle &= \frac{1}{\sqrt{2}}(|+-\rangle + |-+\rangle)
\end{aligned}$$

(one can readily verify that this is indeed orthonormal). The state $|I\rangle$ is orthogonal to $|00\rangle$, so that if I is how the systems were prepared, $|I\rangle$ will never be obtained on measurement — or, put another way, if we obtain measurement result I we have ruled out preparation process I . Similarly, $|II\rangle$ is orthogonal to $|0+\rangle$, so that measurement result II rules out preparation process II, and so forth.

But now: suppose that there is *any* ontic state x given non-zero probability by both state $|\psi_1\rangle$ and $|\psi_2\rangle$. Then there is a non-zero probability for any of preparation processes I-IV to have have left the combined system in the joint ontic state (x, x) . Since preparation process I has zero probability of leading to outcome I, we know that (x, x) has zero probability of leading to outcome I (in terms of the ontic-model framework, we have

$$0 = |\langle I|0+\rangle|^2 = \sum_o \mathcal{A}_{|I\rangle\langle I|}(o)\rho_{|0+\rangle}(o) \quad (24)$$

so if $\rho_{|0+\rangle}(x, x) > 0$, $\mathcal{A}_{|I\rangle\langle I|}(o) = 0$.) But by the same argument, (x, x) has zero probability of giving results II, III or IV either. And this is impossible, because outcomes I-IV are exhaustive. So our original assumption must be false: there is no such x ; the ontic supports of $|\psi_1\rangle$ and $|\psi_2\rangle$ are non-overlapping.

The full PBR theorem generalizes this result to arbitrary states: it shows that if $|\langle\psi_1|\psi_2\rangle| = \cos\theta < 1$, there is some N such that we can consider N identical systems each prepared in $|\psi_1\rangle$ or $|\psi_2\rangle$ and then construct a joint measurement on the combined system such that each outcome rules out one of the 2^N possible preparations. It follows that any two distinct states have non-overlapping ontic support; the PBR theorem thus apparently rules out ψ -epistemic hidden variable theories altogether.

Of course, results of this kind always have assumptions. In the case of the PBR theorem, the most significant assumption is *Preparation Independence*, the assumption that if two systems X, X' are independently prepared in pure quantum states $|\psi\rangle, |\psi'\rangle$, then the joint probability over ontic states of $X \times X'$ from this preparation is just the product of the probability distributions associated to the two states:

$$\rho_{|\psi\rangle\otimes|\psi'\rangle}(o, o') = \rho_{|\psi\rangle}(o)\rho_{|\psi'\rangle}(o'). \quad (25)$$

Violation of Preparation Independence would imply some very subtle underlying correlation of the state of the Universe, such that completely disconnected preparation processes in labs on opposite sides of the world nevertheless had correlated outcomes; it is extremely hard to see how any such theory could be

constructed. (But at some level, this is just a bet on how future theory development will go: if we actually *had* any such theory, the PBR theorem would be of limited significance in assessing it, since we could study the theory itself directly. See also section 5.4.)

The PBR theorem, and the substantial literature developing and extending its results and varying its assumptions (see [67] for a review), has by no means ended attempts to develop ψ -epistemic hidden-variable theories. But it does seem to give significant support either to the idea that the quantum state is not after all to be understood probabilistically, or that any such understanding must make more explicit connection to observation, experiment and measurement, and abandon the classical conception of physics as describing an observer-independent reality; I now turn to how this latter move avoids the difficulties arising from the BKS and PBR theorems, and to its broader goals and challenges.

3.4 Measurement-first approaches to quantum mechanics

All the empirical content of quantum mechanics arises via the Born probability rule, $\langle O \rangle_{|\psi\rangle} = \langle \psi | \hat{O} | \psi \rangle$, and the starting point for the broad class of approaches we might call ‘measurement-first’ is that this rule also gives a complete understanding of the quantum state. $|\psi\rangle$ does not represent any physical features of the underlying system, either directly or through defining probabilities over underlying hidden variables; it simply defines the rule by which the probabilities of experimental outcomes are determined. On this approach to quantum mechanics, ‘measurements’ are misnamed, since they do not really ‘measure’ anything. They are simply indeterministic physical processes, and the role of quantum mechanics is simply to calculate the probabilities associated to those processes.

Measurement-first approaches in a certain sense take the quantum state as probabilistic, and so bypass the problem of Schrödinger’s cat: if the cat’s state is in a superposition of alive and dead, that just means that a measurement of whether it is alive might give either outcome. But since the probabilities are defined directly over measurement outcomes, and not over any intermediary hidden variables, these approaches also bypass the BKS and PBR theorems.

Approaches of this kind are often called ‘instrumentalist’ or ‘operationalist’ given the central role they assign to measurements. (These terms have subtly different meanings in the philosophy literature; in physics they are largely conflated.) The idea has been present since the dawn of quantum mechanics: Heisenberg’s view seems to have been fairly operationalist, and operationalism is one of the sometimes-contradictory strands in the so-called Copenhagen interpretation (see [68, 69] and references therein for more on the vexed question of what ‘the Copenhagen interpretation’ really means).

In more recent physics, operationalism has been advocated in an influential *Physics World* article by Fuchs and Peres [70], and developed by Fuchs and others into ‘quantum Bayesianism’ [71, 72] and later ‘QBism’ [73, 74] (for critical discussion see, e.g., [75, 76]). The distinctive feature of the former is that

the probabilities of quantum mechanics are simply taken to encode an agent's beliefs about measurement outcomes, without any commitment to there being anything underlying those outcomes; the further distinctive feature of the latter is its insistence that no fully-third-personal account of physics is possible, even at the level of observers: all probabilities are probabilities for some localized agent, and there is no guarantee of a consistent account of two or more agents, except insofar as the second agent is treated as one more system about which the first agent has beliefs and assigns probabilities just as they would to any other system. Turning from physicists to philosophers, Richard Healey [77, 78] has developed a 'quantum pragmatism' which treats the quantum state not as an encoding of an agent's belief, but as pragmatic advice to that agent as to how they should act; related accounts include [79, 80, 81]. (In each case these are simplified accounts of subtle positions and interested readers are referred to the references; see [82] for a general review.)

It would beg the question to criticize measurement-first approaches for failing to tell an objective, agent-independent story about the external world: their advocates typically deny that there is any such story to be told, or at any rate that quantum mechanics is supposed to tell it. A sharper criticism [83, 84] is that if quantum mechanics simply tells us the probabilities of outcomes, we need a language or a mathematical framework to say what those outcomes are which is not itself quantum-mechanical. In the heyday of the Copenhagen interpretation, that language was classical physics: Bohr's insistence in a classical description of any experiment can be seen as recognition that the classical language needed to be intelligible independent of the quantum. But it is hard to maintain that approach in modern physics, in which inherently-quantum terms like 'laser' or 'condensate' are ubiquitous even in describing experimental apparatus, in which the workings of that apparatus rely on quantum principles, and in which quantum mechanics is applied to systems like plasmas, superconductors and the early Universe which resist classical descriptions. (Put another way, treating 'measurement' (still less 'belief') as a primitive concept in our interpretation of quantum mechanics runs into the problem that the practice of physics itself does not treat measurement as primitive: measurement devices are physical systems, built according to physical principles.) The most severe challenge for operationalist approaches — the same challenge, in fact, that philosophical operationalism in general has struggled to overcome — is how to make adequate contact with the way in which physics in fact uses quantum theory, outside the somewhat stylized quantum-information contexts in which they are most often developed.

4 Modifying quantum mechanics

The conceptually most straightforward way to address the paradoxes of quantum mechanics is to decide that they are not merely paradoxes by *contradictions*, demonstrations that quantum mechanics is wrong. Any theory that predicts that cats are alive and dead at the same time, when manifestly they are not,

might be said to have refuted itself; perhaps the issue is not how to understand quantum mechanics but how to modify it — or replace it — so that it is not in flat contradiction with the facts. Given how fantastically successful quantum mechanics is, of course, such modifications have to be done delicately, to preserve those successes — and this is easier said than done.

There are a vast number of proposed strategies for how quantum mechanics can be modified, but most proposals — including almost all those which have been developed far enough to reproduce the predictions of a non-trivial part of quantum mechanics — can be classified as either *hidden-variable* or *dynamical-collapse* theories. This classification follows a famous observation by John Bell [85, p.201] that “either the Schrödinger equation is not everything, or it is not right” — that is, either the quantum formalism needs to be *supplemented* with additional ‘hidden’ variables which describe the actual outcomes of experiments, or else the linear quantum dynamics need to be modified so that macroscopic superpositions are suppressed.

4.1 Dynamical collapse theories

To understand how dynamical-collapse theories work, consider again a Schrödinger-cat state,

$$|\psi\rangle = \alpha |\text{Alive}\rangle + \beta |\text{Dead}\rangle \quad (26)$$

Since (putting aside Everett’s move, and the possibility of a probabilistic reading of the state) this is not what we find when we look at the cat, the theory needs to be modified so that states like this do not arise, or at any rate do not persist when observed. This amounts to changing the equations of quantum mechanics, to introduce a new evolution that can be written as

$$\begin{aligned} |\psi\rangle &\longrightarrow |\text{Alive}\rangle && \text{with probability } |\alpha|^2 \\ |\psi\rangle &\longrightarrow |\text{Dead}\rangle && \text{with probability } |\beta|^2. \end{aligned} \quad (27)$$

If this ‘quantum state collapse’ has always occurred by the time that we actually observe the cat, it resolves the measurement problem (we find the cat either alive or dead (and not in a weird superposition of both), and the probability of each matches what the Born rule predicts).

In the infancy of quantum theory, Dirac [86] and von Neumann [60] proposed that it should occur exactly when a system is measured. That is, there should be two different sorts of quantum-mechanical dynamics: unitary (Schrödinger) dynamics, which applies whenever a system is not being measured, and the collapse rule, which occurs at the point of measurement. If we adopted the probabilistic reading of the quantum state, collapse would be trivial, corresponding only to our update of information when we actually discovered whether the cat survived. But as a proposed modification of quantum theory to solve the measurement problem, it cannot be thought of that way: instead, it is an instant, random change of the actual state the cat is in. (As such, if we introduce quantum state collapse to solve the measurement problem, we do so as part of a physical interpretation of the quantum state — the probabilities now occur because of

the randomness in the collapse rule, not as part of the very interpretation of the state.)

This way of presenting quantum mechanics is still found in introductory textbooks, but it has largely been abandoned [5] in the actual practice of quantum mechanics. (It treats measurement as a primitive, unanalyzed notion, and thus shares in the problems of ‘measurement-first’ interpretations, but it lacks their advantages.) But there is an alternative way to think of quantum state collapse: instead of including a fundamental posit that collapse occurs upon measurement, we could imagine a theory where collapse occurs for some other reason, in some other circumstances, that can be described and defined sharply in terms of microscopic physics — and yet those circumstances are in fact such as to ensure that collapse has occurred well before actual measurements are completed. This alternative would be a *dynamical* collapse theory — dynamical referring to some microscopically-defined, bona fide dynamical mechanism, instead of collapse by definition being triggered by the macroscopic concept of measurement.

The constraints on any such theory are strict. If collapse happens too soon, it will suppress the interference effects which quantum theory relies on for its predictions and explanations, and so will falsify itself. If it happens too late, it will fail in its duty to suppress Schrodinger-cat states. But in the non-relativistic regime at least, such theories can be constructed. Probably the best-known, the GRW theory (for Ghirardi, Rimini and Weber) will be briefly described in the next section; other important examples are the CSL (continuous state localization) theory [87] and the proposals of Diosi, Penrose, Stamp and co-workers [88, 89, 90, 91, 92] that attempt to link dynamical collapse to the problem of quantum gravity.

4.2 The GRW theory

The GRW theory [93] applies to a finite number (say, N) of distinguishable, non-relativistic particles, and is most straightforwardly presented in the position representation. Its basic assumptions are:

1. each particle has some small, independent probability of undergoing spontaneous collapse in any small time interval, so that the probability of particle i collapsing over a time δt is $\delta t/\tau$.
2. If collapse occurs for particle i , it occurs at some ‘collapse center’ a , and is represented mathematically by

$$\psi(x_1, \dots, x_n; t) \rightarrow \psi'(x_1, \dots, x_n; t) = \mathcal{N}\psi(x_1, \dots, x_n)\rho(x_i - a) \quad (28)$$

Here ρ is some reasonably-narrowly-peaked wave-packet centered at 0, so that $\rho(x - a)$ is a function of x centered at a : standardly it is taken to be a Gaussian,

$$\rho(x) \propto e^{-x^2/2\Lambda^2}. \quad (29)$$

\mathcal{N} is a normalization constant, chosen to preserve the standard normalization condition for the wavefunction.

3. If (without loss of generality) particle 1 undergoes a collapse, the probability of that collapse having center a (conditional of collapse having happened in the first place) is

$$\Pr(\text{collapse at } a) = \int dx_2 \cdots dx_N |\psi(a, x_2, \dots, x_N)|^2. \quad (30)$$

In other words, the probability of the collapse center being at a is numerically equal to the Born-rule probability that a measurement of the position of particle 1 would give result a . (And similarly for the other particles.)

4. Collapses are instantaneous, and between collapses the wavefunction evolves unitarily according to the Schrödinger equation.

The parameters τ and Λ are to be thought of as new constants of nature, ultimately to be determined empirically: in the original GRW paper, the sample values were $\tau = 10^{16}$ s and $\lambda = 10^{-7}$ m.

To see how the theory works, consider two scenarios. In the first, we have (say) 10^{20} particles, all unentangled with one another. Every second, thousands will undergo collapse, and if each is delocalized over a region $\gtrsim 10^{-7}$ m, the collapse will significantly alter its quantum state; however, since the particles are unentangled, the collapse *only* affects the particle which collapses. Hence any empirical data is highly unlikely to detect the collapse, because the vast majority of particles will not undergo collapse in any reasonable period, and any given particle will typically evolve unitarily for $\sim 10^9$ years before collapsing.

In the second scenario, we still have 10^{20} particles, but they are bound together into a dust mote, and the dust mote is in a superposition of two macroscopically distinct locations, say with centers of mass at $\pm X/2$, so that the combined wavefunction is something like

$$\psi(x_1, \dots, x_N) = \alpha \psi_0(x_1 + X/2, \dots, x_n + X/2) + \beta \psi_0(x_1 - X/2, \dots, x_N - X/2) \quad (31)$$

where $\psi_0(x_1, \dots, x_n)$ is the wavefunction of a dust mote centered at 0. Writing

$$\psi_{\pm}(x_1, \dots, x_N) = \psi_0(x_1 \mp X/2, \dots, x_N \mp X/2) \quad (32)$$

we can express this more compactly as

$$\psi = \alpha \psi_- + \beta \psi_+. \quad (33)$$

Within $\sim 10^{-4}$ s one of the particles will collapse, and with probability ~ 1 the collapse centre will be in the vicinity of either $X/2$ or $-X/2$. The new wavefunction, to a very good approximation, will be either

$$\psi' \simeq \psi_- + (\beta/\alpha) e^{-\frac{1}{2}(X/\Lambda)^2} \tilde{\psi}_+ \quad (\text{probability } |\alpha|^2) \quad (34)$$

or

$$\psi' \simeq (\alpha/\beta)e^{-\frac{1}{2}(X/\Lambda)^2}\tilde{\psi}_- + \psi_+ \quad (\text{probability } |\beta|^2) \quad (35)$$

where $\tilde{\psi}_\pm$ are distorted [94] versions of the original ψ_\pm . Assuming that $X \gg \Lambda$, the wavefunction goes from being a macroscopic superposition to being (extremely close to) a macroscopically definite state. Furthermore, the probability of the system collapsing to a state localized at $-X/2$ is exactly the ordinary quantum probability ($|\alpha|^2$) of a position measurement finding the system at $-X/2$. And any subsequent collapse is almost certain also to be at that location.

So: as long as we are considering small systems, GRW is empirically indistinguishable from unitary quantum mechanics; once we start trying to construct macroscopic superpositions, GRW very rapidly collapses them, and does so with the correct probabilities. These are the two central requirements for a collapse theory.

There has been a certain amount of foundational discussion of GRW as though it was a candidate for a fundamental theory, in which context some philosophical problems become salient, notably the ‘problem of tails’ ([95, 96, 97]; for a review see [98]): state (34) is still, technically, a macroscopic superposition, just one with very uneven amplitudes, so can we really take it to solve the measurement problem? But it is probably better to think of GRW as a phenomenological parameterization of what a more fundamental modification to quantum theory might look like in the non-relativistic regime. And from this point of view, probably the most important feature of GRW is that in principle it is a *testable* alternative to quantum mechanics.

More precisely, for any given choice of Λ and τ , GRW makes concrete predictions that disagree with quantum mechanics; testing GRW, then, means constructing superpositions of many particles and either (i) detecting dynamical collapse or (ii) not doing so, and so constraining the choice of Λ and τ . No experiment can rule out all possible values (if τ is so large that only one particle in the Universe would have collapsed by now, we would have no chance of detecting that collapse) but sufficiently large values of τ or Λ will fail to achieve GRW’s basic goals of suppressing macroscopic superpositions and so solving the measurement problem.

The difficulty in carrying out these tests (and this applies to all collapse theories of which I am aware, not just to GRW) is that superpositions of many particles tend to be very prone to decoherence, and decoherence is very hard to distinguish from genuine non-unitary dynamics (decoherence provides an *effective* rather than a *real* collapse of the quantum state, but we can only tell the two apart by implausibly precise measurements of the environment). So the trick is to somehow construct a system which is *macroscopic* (hence prone to GRW collapses) but with very few effective degrees of freedom (hence protected from internal decoherence) and very isolated from its environment (hence protected from decoherence from the external environment). A very cold, very isolated metal bar, for instance, is a potentially promising system for testing dynamical collapse. Some experiments of this kind have been carried out (by far the most

impressive is the LIGO gravity-wave observatory, which of course was not built to test collapse theories but in fact tests them quite stringently); so far there has been no evidence of dynamical collapse (GRW-style or otherwise), but there are still regions of parameter space not ruled out (see [99] for a review).

For general reviews of GRW theory, and the related CSL theory, see [100, 101].

4.3 Hidden-variable theories

In section 3.1, we met ψ -epistemic hidden-variable theories. In these theories, the real physical content of quantum mechanics was supposed to be expressed in terms of some new degrees of freedom — the hidden variables — and their dynamics. The quantum state was to be understood as an indirect representation of a probability distribution over those new degrees of freedom, to be understood in the same way that the probability distributions of statistical mechanics are to be understood. We also saw that — thanks to no-go results like the Bell-Kochen-Specker and Pusey-Barrett-Rudolph theorems — the prospects for finding any such theory for realistic quantum systems are fairly dim.

But there is an alternative conception of hidden-variable theory which embraces the dual nature of the quantum state as both representational and probabilistic. In these theories, the hidden variables are added to the quantum formalism, physical measurements are supposed to return information about those variables, and the probabilities of quantum mechanics are taken to arise from some probability distribution over the true, unknown, values of the variables. But the quantum state is not interpreted probabilistically: it remains in the theory, evolves unitarily, and serves to determine the dynamics of the hidden variables. Theories of this kind are variably called *ψ -ontic* or *dualist* hidden-variable theories — *ψ -ontic* because the quantum state represents part of the ontology of the theory and not just our ignorance or some other measure of probability, *dualist* because the theory’s ontology contains both whatever the quantum state represents and whatever entities are described by the hidden variables.

Constructing a hidden-variable theory normally requires selecting some subset of quantum observables and taking the hidden variables to be the actually-possessed values of those observables (so that measurements of the observables return those actually-possessed values). In accordance with the no-go results of section 3.2, that subset needs to be taken to be (at most) a complete set of commuting observables, and ‘measurements’ of observables not in the set are determined contextually and do not simply return already-possessed values.

There is then a choice to make: is the preferred set of observables to be determined by the quantum state, or is it fixed, state-independently, as part of the specification of the dynamics? The first option leads to so-called *modal interpretations* [102, 103, 104], which were popular in the 1980s and 1990s but have somewhat fallen out of favor due to fairly severe technical problems ([105, 106]; see also [98] and references therein). Most hidden-variable theories discussed at present choose the second option; at least in the nonrelativistic context, by

far the most common choice of preferred observable is particle position, and by far the most commonly-discussed hidden-variable theory is the *de Broglie-Bohm theory* (also called *Bohmian mechanics* or the *pilot-wave theory*), developed by Bohm [107] in the 1950s following earlier work by de Broglie. In the following section I briefly present the de Broglie-Bohm theory, both for its own interest and as an illustration of the general structure of hidden-variable theories of this kind.

4.4 The de Broglie-Bohm theory

The de Broglie-Bohm theory is a ψ -ontic hidden-variable theory for a finite number N of distinguishable spinless nonrelativistic particles (generalization to spin and to indistinguishable particles is relatively straightforward). Physically it consists of N point particles, together with the wavefunction (or more accurately: together with whatever the wavefunction is taken to represent). It is specified in three steps:

1. The *state* of the theory at a time t is given by:
 - The wave-function, a complex, differentiable, square-integrable function $\psi(x_1, \dots, x_N; t)$ with square-integral 1. (This is basically the normal wavefunction in position representation, though the de Broglie-Bohm theory places stricter constraints of differentiability on the wavefunction than ordinary quantum theory does.)
 - N points $q_1(t), \dots, q_N(t)$ in three-dimensional space, representing the positions at time t of the N particles.
2. The *dynamics* of the theory are given by:
 - The normal Schrödinger dynamics for the wave-function,

$$i\hbar \frac{\partial \psi}{\partial t} = - \sum_I \frac{\hbar^2}{2m_I} \nabla_I^2 \psi + V\psi \quad (36)$$

where ∇_I is the gradient with respect to the i th coordinate, and m_I is the quantum-mechanical mass of the i th particle.

- The *guidance equation*,

$$m_J \frac{dq_J}{dt} = \hbar \mathbf{Im} \left(\frac{\nabla_I \psi}{\psi} \right). \quad (37)$$

Here the right-hand side of the equation is evaluated at the point $(q_1(t), \dots, q_N(t))$ determined by the joint positions of all of the particles; the guidance equation thus makes the velocity of each particle depend on the positions of all of the particles, via its dependence on the wavefunction in the vicinity of the configuration-space point picked out by their joint locations.

3. The *probabilities* of the theory are given by a probability function $\rho(x_1, \dots, x_n; t)$, representing the probability density to find the N particles at time t in the vicinity of (x_1, \dots, x_N) . At some *fixed* time t_0 , the probability distribution is required to satisfy

$$\rho(x_1, \dots, x_N; t_0) = |\psi(x_1, \dots, x_N; t_0)|^2. \quad (38)$$

It is easy to verify that if (38) holds at time t_0 , and the guidance equation holds for all times, then $\rho(x_1, \dots, x_N; t) = |\psi(x_1, \dots, x_N; t)|^2 = 0$ for all times t , not just the specific time t_0 (a property called *equivariance*).

The de Broglie-Bohm theory connects to experimental results through a 2-part interpretative assumption, logically independent of the formal specification of the theory:

Interpretative assumption 1: Those quantum measurements that physicists treat as ‘position measurements’ in fact return the positions of the de Broglie-Bohm particles.

Interpretative assumption 2: All quantum measurements are ultimately re-describable as measurements of position; or, put another way: the results of any measurement are ultimately encoded in the positions of particles.

Given assumption 1, the probability distribution over measurements of position is in fact given by ρ — but, by construction, it is numerically equal to the Born probability distribution, and so reproduces the empirical predictions of quantum theory as far as position measurements are concerned. Given assumption 2, this can be extended to all quantum measurements. In addition, the guidance equation implies that when the quantum state is in a macroscopic superposition, the particles jointly respond to only that part of the wavefunction corresponding to one term in the superposition; other parts of the wavefunction can in practice be discarded. The de Broglie-Bohm theory thus effectively reproduces wavefunction collapse despite its unitary quantum dynamics, ensuring that repeated measurements give the same outcomes.

(To see the importance of this last condition, contrast the de Broglie-Bohm theory with the so-called ‘Everett(?)’ theory described by Bell [108] (its connection to the actual Everett interpretation is tenuous). In that theory, the guidance equation is dropped and the probability rule is just imposed independently at each instant of time, describing a wildly stochastic dynamics in which the particles constantly jump around in space. The Everett(?) theory reproduces the instantaneous probabilities, but implies that Schrödinger’s cat is constantly fluctuating from alive to dead and back again, with our memories and records constantly being rewritten; theories of this kind are normally deemed unacceptable — though see [109] for a defense.)

Interpretative assumption 2 (originally stated in [110]) seems fairly unproblematic; as a last resort, one can always fall back on the fact that lab results can be written down on paper and discussed by scientists, and both these processes encode the results in position data. Assumption 1 is more controversial.

Advocates of the de Broglie-Bohm theory have generally [111, 25, 112] claimed that a core presumption of the theory is that ordinary macroscopic objects are comprised of the de Broglie-Bohm particles (so that, for instance, the needle on a measurement device is a swarm of such particles); as such, measurements just are correlations of particle positions with one another, and indeed the definite world we observe around us is simply a vast collection of these particles. From this perspective it is the wave-function, not the de Broglie-Bohm particles, that is ‘hidden’: it affects our observations only indirectly, through its dynamical influence on the particles. (Some advocates of the de Broglie-Bohm theory (e.g. [113]) go so far as to suggest the wave-function might be akin to a law of nature rather than representing any physical quantities.) Critics have countered that the claim that the observable world is just the particles needs to be argued for and not stipulated, and that since the structural features of the world are encoded in (one branch of) the wave-function, there is no clear reason to regard measurements as recording that structure rather than the particle positions. The ‘Everett-in-denial’ version of this argument [114, 115, 116, 117, 118] asserts that since the de Broglie-Bohm theory is unitary quantum theory plus some particles that do not interact with the wavefunction, it should be regarded as a version of the Everett interpretation equipped with additional, redundant, ontology.

Any resolution of this debate turns on fairly philosophical issues of how theories can and should be interpreted — and this is typical of debates about the de Broglie-Bohm theory since — unlike collapse theories — the theory is generally taken as empirically equivalent to ordinary quantum theory. This fact, together with the clear *mathematical* observation that the theory is unitary quantum mechanics plus additional particles (however this is to be interpreted physically or metaphysically) probably accounts for the relative lack of interest in the theory in the mainstream physics community. (Its advocates would respond that its practical — and pedagogical significance is the much greater level of explanatory clarity it offers.)

That said, it is not *uncontroversial* that the theory cannot be tested empirically against standard quantum mechanics: there is a research program [119, 120, 121] that seeks to understand the probability assumption (38) as having the character of statistical equilibrium, so that one might imagine exotic circumstances (the early Universe) in which the probability rule does not hold. In this circumstance, the de Broglie-Bohm theory’s predictions would deviate — perhaps wildly, from standard quantum predictions. No empirical evidence has been found to support this idea to date, but work continues.

For reviews of the de Broglie-Bohm theory, see [122, 123, 110, 124, 125].

4.5 Beyond non-relativistic quantum theory

The most severe problem for modificatory theories is, of course, the enormous empirical success of standard quantum theory and the corresponding challenge for any research program that seeks to replace it. The de Broglie-Bohm theory and various developments of the GRW theory are the furthest-developed of

any such modifications so far, and neither has really succeeded in going far beyond nonrelativistic particle mechanics into relativistic quantum field theory: specifically, to the best of my knowledge, neither has developed to the stage of being able to reproduce a prediction in (say) quantum electrodynamics beyond leading order, or reproduce the renormalization analysis at the core of modern quantum field theory. (There are developments of both that incorporate some features of quantum field theory short of this test: in the de Broglie-Bohm theory, see [126, 127, 128, 129, 130, 131]; for the GRW theory, see [132, 133, 134]. There are also versions of dynamical collapse [92] which appear to match quantum field theory at the microscopic level but as yet lack a demonstration that the collapse mechanism actually solves the measurement problem.)

This is a very substantial limitation on the scope of these theories, going well beyond an inability to apply them to high-energy physics [84]. A proper quantum description of light requires quantum field theory; hence, at present neither the GRW nor the de Broglie-Bohm theory can provide a full account of classic quantum experiments like the (photon) two-slit experiment, or the spectral lines of hydrogen. Furthermore, there is at least some tension between the strategy of these theories (which rely on a relatively direct encoding of macroscopic observables in the fundamental structure of the theory) and the modern understanding of quantum field theory (in which the theory is taken as an effective description of the world at certain scales, and in which the relation between ‘fundamental’ ontology and phenomena is opaque and indirect). At any rate, extending the theory to quantum field theory in a reasonably systematic manner remains the overwhelmingly most substantial challenge for any proposal to modify quantum mechanics.

5 Quantum nonlocality

Since Einstein, Podolsky and Rosen’s famous EPR thought-experiment [135] it has been clear that there are at least subtleties, and perhaps even tensions or outright contradictions, between quantum entanglement and the idea that physics is local. Philosophical discussion of this issue mostly focuses around Bell’s inequality; in this section I present a modern version of that inequality and consider its implications for the various approaches to quantum mechanics which we have discussed so far.

5.1 Bell’s inequality

Put aside quantum mechanics for the moment and consider this simple coordination game (adapted from [136]). Alice and Bob are in separate rooms, and each has a card with ‘+’ on one side and ‘-’ on the other, along with a randomizing device that can read either ‘Heads’ or ‘Tails’. The device might be a simple tossed coin, or a pseudo-random-number generator, or it might generate the random numbers by consulting the ambient background noise or the New York Stock Exchange — all that matters is that it gives each outcome with

Figure 2: Scoring in the Bell-inequality game

Randomizer results	Score for ++ or --	Score for +- or -+
HH	1	0
HT	1	0
TH	1	0
TT	0	1

probability 50% and that Alice’s device and Bob’s are uncorrelated.

In each round of the game, Alice and Bob each activate their randomizer, and then place the card with either ‘+’ or ‘-’ face-up. Scoring is simple (see figure 2): they (jointly) score a point whenever their cards display the same symbol — *except* that if both randomizers display ‘tails’, they score the point whenever the cards display *different* symbols. Alice and Bob can consult beforehand, but cannot communicate with one another between the randomizers being triggered and the cards being laid down.

For the moment, suppose that Alice and Bob adopt *deterministic* strategies. In principle any such strategy might win on 0, 1, 2, 3, or all 4 or the four possible outcomes from the randomizer, and so might give an average score per round of 0, 0.25, 0.5, 0.75, or 1. For instance, Alice might choose to play ‘+’ if the randomizer reads H and ‘-’ if it reads T, while Bob decides to play ‘+’ on every run, a strategy which scores on average 0.75 per round. Or they might both choose to play ‘+’ if the randomizer reads H and ‘-’ if it reads T, in which case they will average only 0.25 per round. It is easy to see, however, that no deterministic strategy scores higher than 0.75 (or lower than 0.25). For any strategy beating 0.75 would have to win *every round*, and that cannot be done. (Suppose that Alice plays x if her randomizer reads H. To win in an HH round, Bob must also play x if his randomizer reads H; to win in an HT round, he must also play x if his randomizer reads T. In other words, he must play x always. To win in a TH round, Alice must play x if her randomizer reads T. But then both Alice and Bob will play x in a TT round, and will lose.)

The restriction to deterministic strategies can easily be lifted. Any randomness (say, Alice playing ‘+’ with 75% likelihood on H and 25% likelihood on T, or Alice and Bob randomly selecting their joint strategy before each round) just causes Alice and Bob to be playing a ‘mixed strategy’, a probabilistic mixture of deterministic strategies. The expected score from that strategy will just be a weighted average of the expected scores from the deterministic strategies; since each of those expected scores is between 0.25 and 0.75, the mixed strategy will likewise average between these bounds.

Suppose Alice and Bob manage to score outside the 0.25-0.75 range (and to do so consistently enough and over a long enough period to rule out dumb luck as an explanation). How could they have done it? There seem to be only two possibilities. The most obvious is that they can signal between their rooms: if Alice and Bob sneaked a cellphone in with them, for instance, their score becomes unmysterious. It is helpful [137, 138] to distinguish two forms

of signalling: Alice and Bob violate *parameter independence* if their choice of card played can depend on the other player's randomizer output; they violate *outcome independence* if their choice of card played can depend on the other player's choice of card. Violating either principle will allow them to exceed the 0.75 bound. We can make this form of cheating more difficult if we arrange for Alice and Bob's play in each round to be spacelike separated; in that case, violation of either principle seems to involve some sort of superluminal signalling.

There is a subtler way for them to cheat: they might have advance knowledge of what the randomizer outputs will be. If Alice and Bob can predict the randomizer results in advance to some extent, and if they can select their strategy so that they make a different choice if they anticipate different randomizer outputs, they can again get outside the 0.25-0.75 range. Alice and Bob violate the *no-conspiracy* requirement if they somehow arrange this.

All of this can be put somewhat more precisely. On a given run of the experiment, let a and b be, respectively, the outcomes of Alice and Bob's randomizers, and let A and B be, respectively, Alice and Bob's plays; let λ represent their choice of strategy. Then a given strategy is defined by a probability measure $\Pr(AB|ab\lambda)$: the joint probability that Alice plays A and Bob plays B , given randomizer outputs a and b and strategy λ . The overall setup of that run is given by a probability measure $\Pr_S(ab\lambda)$: the joint probability that the randomizers read a and b and the strategy chosen is λ .

Our requirements to avoid cheating are then:

Parameter independence: $\Pr(A|ab\lambda) = \Pr(A|ab'\lambda)$ and vice versa.

Outcome independence: $\Pr(A|Bab\lambda) = \Pr(A|B'ab\lambda)$ and vice versa.

No conspiracies: $\Pr(ab\lambda) = \Pr(a)\Pr(b)\Pr(\lambda)$.

The restriction of Alice and Bob's joint score to between 0.25 and 0.75 is called the *Bell inequality*; more precisely, it is the *CHSH inequality* (for Clauser, Horne, Shimony and Holt [139]), a form of the Bell inequality. It is often expressed using a different scoring system, where Alice and Bob score +4 if they win a round and -4 if they lose; so expressed, the inequality bounds their average score between +2 and -2, and the inequality can be re-expressed as

$$CHSH \equiv |\mathcal{C}(HH) + \mathcal{C}(HT) + \mathcal{C}(TH) - \mathcal{C}(TT)| \leq 2 \quad (39)$$

where $\mathcal{C}(ab)$ is the correlation coefficient between Alice and Bob's play, given randomizer results a, b : that is,

$$\mathcal{C}(ab) = \Pr(++|ab) + \Pr(--|ab) - \Pr(+-|ab) - \Pr(-+|ab). \quad (40)$$

Bell's theorem is then the result that if Alice and Bob conform to these requirements, the Bell inequality holds. Of course, the coordination-game way I have presented it is optional (and was not how it was originally presented; see [140, 141] for more traditional versions): we can take 'strategies' simply to be relevant facts about the physical setup of a particular run and replace Alice and Bob's choice of cards with some mechanically-realized process with two possible outcomes.

5.2 Cheating with qubits

There is a different (or, perhaps, not so different) way in which Alice and Bob might cheat so as to violate the Bell inequality. Suppose that before each round of the experiment they share a pair of qubits — spin-half particles, say — in an entangled singlet state:

$$|\psi\rangle = \frac{1}{\sqrt{2}}(|+\rangle\otimes|-\rangle - |-\rangle\otimes|+\rangle). \quad (41)$$

Depending on the outcome of their randomizers, Alice and Bob then decide to measure the spin of their particle along an axis at some chosen angle from the z axis in the x - z plane (they can coordinate their choice of axes in advance). Since singlet states are rotationally invariant, all that matters in predicting the result of their joint measurement is the angular difference θ between their two measurement axes. Taking Alice to measure (without loss of generality) in the $(|+\rangle, |-\rangle)$ basis, Bob is then measuring in the

$$|+(\theta)\rangle = \cos(\theta/2)|+\rangle + \sin(\theta/2)|-\rangle, |-(\theta)\rangle = \sin(\theta/2)|+\rangle - \cos(\theta/2)|-\rangle \quad (42)$$

basis. Elementary quantum mechanics tells us that the probability of correlation ($++$ or $--$) is $\sin^2(\theta/2)$ and the probability of anticorrelation ($+-$ or $-+$) is $\cos^2(\theta/2)$, so that the correlation coefficient is $-\cos(\theta)$.

Alice and Bob adopt the following specific strategies. If Alice's randomizer gets H , she measures spin at $-\phi$ to the $+z$ axis; if she gets T , she measures spin at $+\phi$ to that axis. Meanwhile, if Bob's randomizer gets H , he measures spin along the $-z$ axis; if he gets T , he measures spin at $+2\phi$ to the $-z$ -axis. The result is that Alice and Bob's measurements differ by an angle of $180^\circ - \phi$ in all cases except when their randomizers read TT , in which case their measurements differ by $180^\circ - 3\phi$. Using the second scoring rule, and the correlation result just established, we get

$$CHSH = 3\cos(\phi) - \cos(-3\phi). \quad (43)$$

It is easy to show that there are values of ϕ for which this violates Bell's inequality: indeed, differentiating and setting to zero, we find this expression for $CHSH$ takes its maximum for $\phi = 60^\circ$, at which

$$CHSH(\phi = 60^\circ) = 2\sqrt{2}. \quad (44)$$

Quantum theory violates the Bell inequality. And the specific setup Alice and Bob use here to violate the inequality has been experimentally realized many times (replacing Alice and Bob with simple mechanisms), including with spacelike separation between the two sets of measurements: the quantum-mechanically predicted result of $2\sqrt{2}$ has been robustly reproduced. (For a survey of the experimental situation, see [141].)

5.3 The scope of Bell's theorem

In the philosophy literature there are two main interpretations of Bell's theorem. The *narrow interpretation* takes it specifically as a result about hidden-variable theories: Bell's setup (it is claimed) tacitly assumes exactly the sort of classical physics that quantum theory rejects. For advocates of the narrow interpretation, Bell's theorem is just another nail in the coffin of the quixotic project to construct hidden-variable alternatives to quantum theory. On the narrow interpretation, experimental violation of the Bell inequality is further evidence for the supremacy of quantum mechanics (advocates of the narrow interpretation tend to be sceptical of the project of modifying quantum theory to solve the measurement problem). Careful advocates of the narrow interpretation include [142, 143]; there are many more careless advocates.

The *wide interpretation* takes Bell's theorem as a result about pretty much any physical theory at all: if a theory is local (captured by outcome independence and parameter independence) and if it satisfies no-conspiracy, it will satisfy the Bell inequality. And since the Bell inequality has been violated in the lab (and assuming the no-conspiracy rule is reasonable) it follows that any empirically adequate scientific theory is non-local. In particular, quantum mechanics, insofar as it is empirically adequate, is non-local: or, more precisely (since advocates of the wide interpretation normally take the measurement problem very seriously) any empirically adequate modification or interpretation of quantum mechanics must be non-local. On the wide interpretation, experimental violation of the Bell inequality is a result of profound significance, transcending quantum mechanics itself (Shimony called it a piece of 'experimental metaphysics' [144]). Advocates of the wide interpretation include [6, 145].

Rather than attempt directly to adjudicate this interpretative dispute, I will briefly consider some of the various interpretative approaches I have already discussed and ask what forms of nonlocality we see in them and how those relate to the Bell inequality.

The Everett interpretation: At least in the mainstream 'many-worlds' reading of the Everett interpretation, Bell's inequality fairly obviously is inapplicable, since the proof of the inequality tacitly assumes that measurements have unique outcomes. So there is no reason (at least, no reason arising from Bell's theorem) to think that Everettian quantum mechanics involves any sort of dynamical nonlocality, or any consequent clash with relativistic covariance. (This could have been expected in any case from the fact that the Schrödinger equation is exceptionless in Everettian quantum mechanics and that, at least in relativistic quantum field theory, that equation is explicitly covariant and local. See [142] for more on locality in quantum field theory.)

Relativised to a world, however, Bell's theorem can be proved; as such, it follows that any world-defining rule — insofar as it applies globally — must involve nonlocality. (Again, this might have been expected: if any measurement causes the *whole universe* to split, that seems to define

a preferred simultaneity surface on which it splits.) So if branching is an objective — albeit emergent — feature of the world and not just a convention, the branching process must after all proceed locally. (For detailed attempts to develop theories of branching along these lines see [146], [16, ch.9].)

For the same reason, Bell’s theorem seems to hold for ‘single-world’ versions of Everett, raising further questions as to whether such versions really can be coherently described without modifying quantum theory.

ψ -epistemic hidden-variable theories: If there is an underlying (and single-world) theory which stands to quantum theory as classical mechanics stands to classical statistical mechanics, Bell’s theorem must apply to that theory; since Bell’s inequality is violated, that theory must violate either outcome independence, preparation independence, or no-conspiracy. This stands as a further constraint (alongside those already imposed by the PBR theorem) on constructing any such theory.

Operationalist/pragmatist interpretations: Insofar as it is possible to describe even an operationalist-interpreted quantum theory in terms of objectively recorded measurement outcomes and parameter choices, Bell’s theorem ought to apply to that theory; hence it must violate one of the premises of Bell’s theorem (probably outcome independence). That said, it is not obvious that it *is* possible to describe all operationalist interpretations in this way. One supposed advantage of QBism, for instance (cf [73]) is that the only possible descriptions of a system are from the perspective of a single observer or agent, so that (for instance) it makes sense to discuss Alice’s belief’s about what measurement result by Bob will be *reported to her* but not about what measurement result *he obtains*. Since these reports travel subluminally, there is — arguably — no violation of locality.

Dynamical-collapse theories: Dynamical-collapse theories fairly clearly violate outcome independence: there are correlations between Alice and Bob’s measurement outcomes that are not factored out by any shared initial state. It is less clear that this nonlocality should be interpreted as action at a distance [147], or that it entails any clash with relativistic covariance: indeed, there is a known relativistically-covariant version of GRW [132], albeit only for non-interacting particles.

The de Broglie-Bohm theory: The de Broglie-Bohm theory explicitly includes action at a distance: the trajectory taken by Alice’s particle can be directly influenced by Bob’s choice of measurement parameter and vice versa, though Alice and Bob’s ignorance of the exact positions of their particles prevents them using this information to signal to one another. This is a violation of parameter independence; it also strongly suggests that any theory along these lines will have to involve a preferred choice of simultaneity surface, in violation of relativistic covariance. Bell [110]

and other advocates of the de Broglie-Bohm theory have argued that since Bell's inequality tells us that *any* empirically-adequate theory must violate locality in some fashion, it is a virtue rather than a vice of the de Broglie-Bohm theory that it makes explicit how the trick is done. Of course, this argument will have no traction with anyone who thinks unmodified quantum mechanics can evade Bell's theorem via Everettian or operationalist means.

We can usefully extract from this discussion two tacit principles that seem to be presupposed in the proof of Bell's inequality, along with its explicit requirements:

Objectivity: It must be possible to give a third-party description of an experiment, even one spread over an extended region, without having to localize that description to the first-personal experiences of a particular observer or agent.

Single outcomes: Measurements must have unique outcomes: it must actually be the case that the measurements used in Bell's theorem in the end have as outputs '+' and '-', rather than superpositions of such.

The two are even somewhat related. The Everett interpretation permits an objective description, but that description will be of an entangled superposition in a macroscopically non-classical state. If we require a *classical* description, that description is restricted to the past light cone of an observer.

5.4 The conspiracy loophole

Before leaving Bell's theorem, let's briefly consider its no-conspiracy requirement. The dominant view in foundations of physics is that no physically reasonable theory could exploit this loophole: the random processes that might be used in violations of Bell inequalities could be anything at all and it beggars belief that the world could be so conspiratorially interconnected as to make, say, stock market prices or weather patterns subtly correlated with quantum or subquantum processes here and now in the lab. Recent tests of Bell's theorem have selected the settings of the measurement devices using photons from astronomically distant sources [148], and 100,000 human subjects making choices in an online video game [149]; both robustly violated Bell's inequality. It is often claimed ([150, p.266],[145]) that the scientific method itself requires something like a no-conspiracy assumption. Discussions between the more fervent advocates of this view, and the minority (e. g. , [151, 152]) who have continued to try to develop 'conspiratorial' (or, as they are sometimes called, 'superdeterministic') hidden-variable theories, have mostly produced more heat than light.

The point to keep in mind about any such discussions is that they occur only because we do not have concrete examples of any (non-toy) version of quantum mechanics that holds on to locality by violating no-conspiracy. If we did, we could discuss that theory directly: if it was usable to do science, and

especially if it made novel predictions deviating from quantum mechanics, and *especially* if some of those predictions were confirmed, then we could discuss the theory directly without reference to Bell's theorem. In the absence of that theory, critics of the conspiracy loophole are saying in effect that it will not in fact be possible to construct any scientifically useful theory via that loophole. I *strongly* suspect they are right, but it is hard to see what a completely conclusive argument might look like. (An analogy: Einstein thought [153] that the scientific method requires something like a locality assumption, someone who agreed, and who knew Bell's theorem but not the de Broglie-Bohm theory, might conclude that searching for a single-outcome, objective (and non-conspiratorial) theory would be fruitless, since any such theory would be non-local. The de Broglie-Bohm theory — whatever its virtues or vices — proves that this prediction about future theories was wrong: it is after all possible to build a scientifically coherent non-local theory.)

6 Conclusions

Among physicists, the (more operationalist versions of the) probability-based approach, and the Everett interpretation, are roughly as popular as one another, with different sub-communities having different preferences. (The modificatory strategies are much less popular among physicists, although they are probably the most common choice among philosophers of physics.) But more popular than either is the 'shut-up-and-calculate' approach [154]: the view that we should not worry about these issues and should get on with applying quantum mechanics to concrete problems.⁶

In its place, there is much to be said for 'shut up and calculate'. Not everyone needs to be interested in the interpretation of quantum mechanics; insofar as a physicist working on, say, solar neutrinos or superfluidity can apply the quantum formalism without caring about its interpretation, they should go right ahead — just as a biochemist may be able to ignore quantum mechanics entirely, or a behavioral ecologist may be able to ignore biochemistry. Division of labor is unavoidable in science, and often desirable. But there is a more aggressive reading of 'shut up and calculate' — not just as a description of a physicist's own approach, but as an exhortation to the community to stop wasting their time. That exhortation is often accompanied by the claim that since all the 'interpretations of quantum mechanics' give the same predictions anyway, it is pointless or even unscientific to worry about which one is correct.

Philosophers of physics tend to give a high-minded response to such skepticism: quantum theory tells us about the deepest nature of reality; how could we not be interested in its nature? (This response is often (e.g., [156, p.xiv]) accompanied by a lament about physicists' own retreat from seeking conceptual clarity in their deepest theories.) But there is a more pragmatic response: the various approaches to the measurement problem are not just different verbal glosses on the same quantum formalism but constitute significantly different

⁶Part of this section is adapted from [155, pp.129–132].

strategies for using and applying — or even modifying — that formalism. This is most obvious for modificatory approaches, which are proposals for genuinely, mathematically distinct theories. In some cases, these theories already make predictions — albeit difficult to test — that differentiate them from unmodified quantum theory; in others, they are the seeds of research programs which may lead in a testably different direction from quantum theory. One can regard this as promising or unpromising science, but it is clearly, recognizably, science. (Which is not to say that all advocates of these theories treat them this way — a good test of how serious an advocate of a dynamical-collapse or hidden-variable theory is for their proposal as science is whether they welcome, or resist, the implication that their theory might have testable deviations from quantum mechanics.)

But even within those approaches that leave the quantum formalism unmodified — approaches which treat the quantum state as probabilistic, and Everett-type approaches based on decoherence and the emergence of a classical branching structure — there are major differences of scientific method. The probability-based interpretation has generally been applied in situations where the goal is to understand the interventions and manipulations we might make on a system; it is extremely well suited to the study of computability and information processing, where it has inspired a great deal of insightful work; it naturally leads us to ask why the framework of quantum theory is what it is and not something else; it is the dominant approach in quantum information theory.

The decoherence-based approach has been applied more to situations where the goal is to understand how systems evolve and develop when left to themselves. It has been central in our understanding of quantum/classical transitions, in environments ranging from the present-day laboratory to the early Universe; it provides a framework and a language to handle situations where experiment and measurement do not have a clear meaning; its language of branches and worlds has been valuable in quantum cosmology and in non-equilibrium statistical mechanics; it treats the framework of quantum theory as a given and uses it to understand and explore issues in specific quantum theories; it is the dominant approach in high-energy physics and in string theory.

(Indeed, when physicists in those communities ‘shut up and calculate’, they are normally calculating in unitary quantum mechanics, eschewing any talk of external observers, modelling measurement physically, using decoherence to mediate the quantum/classical transition, and accepting the in-principle possibility that macroscopic superpositions could be recohered, even as they stress its in-practice impossibility; the gap between this position and full-on advocacy of the Everett interpretation is thin and largely verbal.)

This is not to say that everyone who has used decoherence-based methods to study cosmology is explicitly committed to the Everett interpretation and its language of many worlds, or that everyone who was inspired by a probabilistic interpretation of quantum theory to prove a valuable theorem in quantum information is explicitly committed to one or other variety of instrumentalism. It is to say that there has been, and continues to be, a continuous flow of ideas

and inspiration between considerations of the quantum measurement problem to and more concrete issues in quantum theory.

References

- [1] Willard Van Orman Quine. Two dogmas of empiricism. *Philosophical Review*, 60:20–43, 1951.
- [2] Hilary Putnam. Is logic empirical? In R. Cohen and M. Wartofsky, editors, *Boston Studies in the Philosophy of Science*, vol. 5, pages 216–241. Reidel, Dordrecht, 1968. Reprinted as ‘The Logic of Quantum Mechanics’ in H. Putnam, *Mathematics, Matter and Method: Philosophical Papers*, vol. 1 (Cambridge: Cambridge University Press, 1975), pp. 174–197.
- [3] Michael Dickson. Quantum logic is alive \wedge (it is true \vee it is false). *Philosophy of Science*, 68:S274–S287, 2001.
- [4] Guido Bacciagaluppi. Is logic empirical? In Kurt Engesser, Dov Gabbay, and Daniel Lehmann, editors, *The Handbook of Quantum Logic and Quantum Structures*. Elsevier, 2007. Available online at <http://philsci-archive.pitt.edu/archive/00003380/>.
- [5] David Wallace. What is orthodox quantum mechanics? In Alberto Cordero, editor, *Philosophers Look at Quantum Mechanics*, pages 285–309. Springer, Switzerland, 2019.
- [6] D. Z. Albert. *Quantum mechanics and experience*. Harvard University Press, Cambridge, Massachusetts, 1992.
- [7] Tim Maudlin. Three measurement problems. *Topoi*, 14:7–15, 1995.
- [8] J.S. Bell. Against measurement. In *Speakable and Unsayable in Quantum Mechanics (Second Edition)*, pages 213–231. Cambridge University Press, 2004.
- [9] Guido Bacciagaluppi. The role of decoherence in quantum mechanics. In the Stanford Encyclopedia of Philosophy (Summer 2005 edition), Edward N. Zalta (ed.), available online at <http://plato.stanford.edu/archives/sum2005/entries/qm-decoherence>, 2005.
- [10] E. Joos, H. D. Zeh, C. Kiefer, D. Giulini, J. Kupsch, and I. O. Stamatescu. *Decoherence and the Appearance of a Classical World in Quantum Theory*. Springer, Berlin, 2nd edition edition, 2003.
- [11] Maximilian Schlosshauer. *Decoherence and the Quantum-to-Classical Transition*. Springer, Berlin, 2007.

- [12] H. III Everett. Relative State Formulation of Quantum Mechanics. *Review of Modern Physics*, 29:454–462, 1957. Reprinted in [157].
- [13] B. DeWitt. Quantum Mechanics and Reality. *Physics Today*, 23(9):30–35, 1970. Reprinted in [157].
- [14] Brett Maynard Bevers. Everett’s “many-worlds” proposal. *Studies in the History and Philosophy of Modern Physics*, 42:3–12, 2011.
- [15] Jeffrey A. Barrett. Everett’s pure wave mechanics and the notion of worlds. *European Journal for Philosophy of Science*, 1:277–302, 2011.
- [16] David Wallace. *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford University Press, Oxford, 2012.
- [17] Simon Saunders. The Everett interpretation: structure. In Eleanor Knox and Alastair Wilson, editors, *The Routledge Companion to Philosophy of Physics*. Routledge, 2021. Preprint at <http://philsci-archive.pitt.edu/18772/>.
- [18] Simon Saunders, Jon Barrett, Adrian Kent, and David Wallace, editors. *Many Worlds? Everett, Quantum Theory, and Reality*, Oxford, 2010. Oxford University Press.
- [19] Simon Saunders. The Everett interpretation: probability. In Eleanor Knox and Alastair Wilson, editors, *The Routledge Companion to Philosophy of Physics*. Routledge, 2021. Preprint at philsci-archive.pitt.edu/19362/.
- [20] Lev Vaidman. The Many-Worlds Interpretation of Quantum Mechanics. In the Stanford Encyclopedia of Philosophy (Summer 2002 edition), Edward N. Zalta (ed.), available online at <http://plato.stanford.edu/archives/sum2002/entries/qm-manyworlds> ., 2002.
- [21] J. J. Halliwell. A review of the decoherent histories approach to quantum mechanics. *Annals of the New York Academy of Sciences*, 755:726–740, 1995.
- [22] M. Gell-Mann and J. B. Hartle. Classical equations for quantum systems. *Physical Review D*, 47:3345–3382, 1993.
- [23] David Wallace. Everett and Structure. *Studies in the History and Philosophy of Modern Physics*, 34:87–105, 2003.
- [24] David Wallace. Decoherence and ontology: Or: How I learned to stop worrying and love FAPP. In Simon Saunders, Jon Barrett, Adrian Kent, and David Wallace, editors, *Many Worlds? Everett, Quantum Theory, and Reality*, pages 53–72. Oxford University Press, Oxford, 2010.

- [25] Valia Allori, Sheldon Goldstein, Roderich Tumulka, and Nino Zanghí. On the common structure of Bohmian mechanics and the Ghirardi-Rimini-Weber theory. *British Journal for the Philosophy of Science*, 59:353–389, 2008.
- [26] Tim Maudlin. Can the world be only wavefunction? In Simon Saunders, Jonathan Barrett, Adrian Kent, and David Wallace, editors, *Many Worlds? Everett, Quantum Theory, and Reality*. Oxford University Press, Oxford, 2010.
- [27] Robert B. Griffiths. Consistent histories and the interpretation of quantum mechanics. *Journal of Statistical Physics*, 36:219–272, 1984.
- [28] Roland Omnes. Consistent interpretations of quantum mechanics. *Reviews of Modern Physics*, 64:339–382, 1992.
- [29] M. Gell-Mann and J. B. Hartle. Quantum Mechanics in the Light of Quantum Cosmology. In W. H. Zurek, editor, *Complexity, Entropy and the Physics of Information*, pages 425–459. Westview Press, Redwood City, California, 1989.
- [30] D. Papineau. Many Minds are No Worse than One. *British Journal for the Philosophy of Science*, 47:233–241, 1996.
- [31] S. Saunders. Time, Quantum Mechanics, and Probability. *Synthese*, 114:373–404, 1998.
- [32] David Papineau. A fair deal for Everettians. In Simon Saunders, Jonathan Barrett, Adrian Kent, and David Wallace, editors, *Many Worlds? Everett, Quantum Theory, and Reality*. Oxford University Press, Oxford, 2010.
- [33] E. Farhi, J. Goldstone, and S. Gutmann. How probability arises in quantum-mechanics. *Annals of Physics*, 192:368–382, 1989.
- [34] David Deutsch. Quantum theory of probability and decisions. *Proceedings of the Royal Society of London*, A455:3129–3137, 1999.
- [35] David Wallace. How to prove the Born rule. In Simon Saunders, Jon Barrett, Adrian Kent, and David Wallace, editors, *Many Worlds? Everett, Quantum Theory and Reality*. Oxford University Press, Oxford, 2010. Available online at <http://arxiv.org/abs/0906.2718>.
- [36] Wojciech H. Zurek. Probabilities from entanglement, Born’s rule from envariance. *Physical Review A*, 71:052105, 2005. Available online at <http://arxiv.org/abs/quant-ph/0405161>.
- [37] Hilary Greaves. On the Everettian epistemic problem. *Studies in the History and Philosophy of Modern Physics*, 38:120–152, 2007.

- [38] H. Greaves and W. Myrvold. Everett and evidence. In Simon Saunders, Jon Barrett, Adrian Kent, and David Wallace, editors, *Many Worlds? Everett, Quantum Theory and Reality*. Oxford University Press, Oxford, 2010.
- [39] Sean M. Carroll and Charles T. Sebens. Many worlds, the born rule, and self-locating uncertainty. In D.C. Struppa and J.M. Tollaksen, editors, *Quantum Theory: a Two-Time Success Story, Yakir Aharonov Festschrift*, page 157. Springer-Verlag, Berlin, 2013.
- [40] Adrian Kent. One world versus many: the inadequacy of Everettian accounts of evolution, probability, and scientific confirmation. In Simon Saunders, Jonathan Barrett, Adrian Kent, and David Wallace, editors, *Many Worlds? -Everett, Quantum Theory, and Reality*, Oxford, 2010. Oxford University Press.
- [41] David Z. Albert. Probability in the Everett picture. In Simon Saunders, Jon Barrett, Adrian Kent, and David Wallace, editors, *Many Worlds? Everett, Quantum Theory and Reality*. Oxford University Press, Oxford, 2010.
- [42] Huw Price. Probability in the Everett picture. In Simon Saunders, Jon Barrett, Adrian Kent, and David Wallace, editors, *Many Worlds? Everett, Quantum Theory and Reality*. Oxford University Press, Oxford, 2010.
- [43] H.D. Zeh. Toward a quantum theory of observation. *Foundations of Physics*, 3:109–116, 1973.
- [44] Robert B. Griffiths. *Consistent Quantum Theory*. Cambridge University Press, Cambridge, 2002.
- [45] Roland Omnes. Logical reformulation of quantum mechanics. i. foundations. *Journal of Statistical Physics*, 53:893–932, 1988.
- [46] J. J. Halliwell. Macroscopic superpositions, decoherent histories and the emergence of hydrodynamic behaviour. In Simon Saunders, Jon Barrett, Adrian Kent, and David Wallace, editors, *Many Worlds? Everett, Quantum Theory, and Reality*. Oxford University Press, Oxford, 2010.
- [47] Wojciech H. Zurek. Decoherence, einselection, and the quantum origins of the classical: the rough guide. *Philosophical Transactions of the Royal Society of London*, A356:1793–1820, 1998.
- [48] Carlo Rovelli. Relational quantum mechanics. *International Journal of Theoretical Physics*, 35:1637, 1996. Substantially revised version at arxiv.org/abs/quant-ph/9609002.
- [49] Carlo Rovelli. Space is blue and birds fly through it. *Philosophical Transactions of the Royal Society A*, 376:2017.0312, 2018.

- [50] J. A. Barrett. *The quantum mechanics of minds and worlds*. Oxford University Press, Oxford, 1999.
- [51] A. Bassi and G. Ghirardi. Decoherent histories and realism. *Journal of Statistical Physics*, 98:457–494, 2000.
- [52] E. Okon and D. Sudarsky. Measurements according to consistent histories. *Studies in History and Philosophy of Modern Physics*, 52:217–222, 2014.
- [53] R. Griffiths. Consistent quantum measurements. *Studies in History and Philosophy of Modern Physics*, 48:7–12, 2015.
- [54] E. Okon and D. Sudarsky. The consistent histories formalism and the measurement problem. *Studies in History and Philosophy of Modern Physics*, 52:188–197, 2015.
- [55] M. Brown. Relational quantum mechanics and the determinacy problem. *British Journal for the Philosophy of Science*, 60:679, 2009.
- [56] Bas van Fraassen. Rovelli’s world. *Foundations of Physics*, 40:390–417, 2010.
- [57] Nicholas Harrigan and Robert W. Spekkens. Einstein, incompleteness, and the epistemic view of states. *Foundations of Physics*, 40:125, 2010.
- [58] Michael Redhead. *Incompleteness, Nonlocality and Realism: A Prolegomenon to the Philosophy of Quantum Mechanics*. Oxford University Press, Oxford, 1987.
- [59] S.D. Bartlett, T. Rudolph, and R. W. Spekkens. Reconstruction of Gaussian quantum mechanics from Liouville mechanics with an epistemic restriction. *Physical Review A*, 86:012103, 2012.
- [60] J. von Neumann. *Mathematical Foundations of Quantum Mechanics*. Princeton University Press, Princeton, 1955.
- [61] J. S. Bell. On the problem of hidden variables in quantum mechanics. *Reviews of Modern Physics*, 38:447–452, 1966. Reprinted in [85], pp. 1–13.
- [62] S. Kochen and E. Specker. The problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics*, 17:59–87, 1967.
- [63] N. David Mermin. Hidden variables and the two theorems of John Bell. *Reviews of Modern Physics*, 65:803–815, 1993.
- [64] Jeffrey Bub and Robert Clifton. A uniqueness theorem for “no collapse” interpretations of quantum mechanics. *Studies in the History and Philosophy of Modern Physics*, 27:181–219, 1996.

- [65] Jeffrey Bub, Rob Clifton, and Sheldon Goldstein. Revised proof of the uniqueness theorem for ‘no collapse’ interpretations of quantum mechanics. *Studies in the History and Philosophy of Modern Physics*, 31:95, 2000.
- [66] Matthew F. Pusey, Jonathan Barrett, and Terry Rudolph. On the reality of the quantum state. *Nature Physics*, 8:476, 2011. arXiv:1111.3328v2.
- [67] M.S. Leifer. Is the quantum state real? an extended review of ψ -ontology theorems. *Quanta*, 3:67–155, 2014.
- [68] Don Howard. Who invented the Copenhagen Interpretation? a study in mythology. *Philosophy of Science*, 71:669–682, 2004.
- [69] Simon Saunders. Complementarity and scientific rationality. *Foundations of Physics*, 35:347–372, 2005.
- [70] Chris Fuchs and Asher Peres. Quantum theory needs no “interpretation”. *Physics Today*, 53(3):70–71, 2000.
- [71] Carlton M. Caves, Christopher A. Fuchs, and Ruediger Schack. Quantum probabilities as Bayesian probabilities. *Physical Review A*, 65:022305, 2002.
- [72] Chris Fuchs. Quantum mechanics as quantum information (and only a little more). Available online at <http://arXiv.org/abs/quant-ph/0205039>, 2002.
- [73] Christopher A. Fuchs, N. David Mermin, and Ruediger Schack. An introduction to QBism with an application to the locality of quantum mechanics. *American Journal of Physics*, 82:749–754, 2014.
- [74] Christopher A. Fuchs and Ruediger Schack. QBism and the Greeks: why a quantum state does not represent an element of physical reality. *Physica Scripta*, 90:015104, 2015.
- [75] Chris Timpson. Quantum bayesianism: a study. *Studies*, 39:579–609, 2008.
- [76] Amit Hagar. A philosopher looks at quantum information theory. *Philosophy of Science*, 70:752–775, 2003.
- [77] Richard Healey. Quantum theory: A pragmatist approach. *British Journal for the Philosophy of Science*, 63:729–771, 2012.
- [78] Richard Healey. *The Quantum Revolution in Philosophy*. Oxford University Press, Oxford, 2017.
- [79] Jeff Bub and Itamar Pitowsky. Two dogmas about quantum mechanics. In Simon Saunders, Jon Barrett, Adrian Kent, and David Wallace, editors, *Many Worlds? Everett, Quantum Theory, and Reality*. Oxford University Press, Oxford, 2010.

- [80] Simon Friederich. In defense of non-ontic accounts of quantum states. *Studies in History and Philosophy of Modern Physics*, 44:77–92, 2013.
- [81] Simon Friederich. *Interpreting Quantum Theory — a Therapeutic Approach*. Palgrave Macmillan, London, 2015.
- [82] Richard Healey. Quantum-Bayesian and pragmatist views of quantum theory. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy (Spring 2017 edition)*. 2017. <https://plato.stanford.edu/archives/spr2017/entries/quantum-bayesian/>.
- [83] Harvey Brown. The reality of the wavefunction: Old arguments and new. In Alberto Cordero, editor, *Philosophers Look at Quantum Mechanics*, pages 63–86. Springer Verlag, Berlin, 2019.
- [84] David Wallace. On the plurality of quantum theories: Quantum theory as a framework, and its implications for the quantum measurement problem. In Juha Saatsi and Steven French, editors, *Scientific Realism and the Quantum*, pages 78–102. Oxford University Press, Oxford, 2020.
- [85] J. S. Bell. *Speakable and Unspeakable in Quantum Mechanics*. Cambridge University Press, Cambridge, 1987.
- [86] Paul Dirac. *The Principles of Quantum Mechanics*. Oxford University Press, 1930.
- [87] P. Pearle. Combining Stochastic Dynamical State-Vector Reduction with Spontaneous Localization. *Physical Review A*, 39(5):2277–2289, 1989.
- [88] L. Diósi. A universal master equation for the gravitational violations of quantum mechanics. *Physics Letters*, 120:377–381, 1987.
- [89] L. Diósi. Gravity related spontaneous wavefunction collapse in bulk matter. *New Journal of Physics*, 16:105006, 2014.
- [90] Roger Penrose. On gravity’s role in quantum state reduction. *General Relativity and Gravitation*, 28:581–600, 1996.
- [91] Roger Penrose. On the gravitization of quantum mechanics 1: Quantum state reduction. *Foundations of Physics*, 44:557–575, 2014.
- [92] P.C.E. Stamp. Rationale for a correlated worldline theory of quantum gravity. *New Journal of Physics*, 17:065017, 2015.
- [93] G. Ghirardi, A. Rimini, and T. Weber. Unified Dynamics for Micro and Macro Systems. *Physical Review D*, 34:470–491, 1986.
- [94] David Wallace. Life and death in the tails of the GRW wave function. <https://arxiv.org/abs/1407.4746>, 2014.

- [95] D. Z. Albert and B. Loewer. Tails of Schrödinger’s Cat. In R. Clifton, editor, *Perspectives on Quantum Reality*, pages 81–92. Kluwer Academic Publishers, Dordrecht, 1996.
- [96] Alberto Cordero. Are GRW tails as bad as they say? *Philosophy of Science*, 66:S59–S71, 1999.
- [97] Peter J. Lewis. Quantum mechanics and ordinary language: the fuzzy link. *Philosophy of Science*, 70(55):713–729, 2004.
- [98] David Wallace. The interpretation of quantum mechanics. In Dean Rickles, editor, *The Ashgate Companion to Contemporary Philosophy of Physics*, pages 197–261. Ashgate, Burlington, VT, 2008. Available online at <http://arxiv.org/> under the title “The Quantum Measurement Problem: State of Play”.
- [99] Matteo Carlesso and Sandro Donadi. Collapse models: main properties and the state of art of the experimental tests. <https://arxiv.org/abs/1907.12460>, 2019.
- [100] Angelo Bassi and GianCarlo Ghirardi. Dynamical reduction models. *Physics Reports*, 379:257, 2003.
- [101] Gian Carlo Ghiardi. Collapse theories. In the Stanford Encyclopedia of Philosophy (Summer 2002 edition), Edward N. Zalta (ed.), available online at <http://plato.stanford.edu/archives/spr2002/entries/qm-collapse> ., 2002.
- [102] B. C. van Fraassen. *Quantum Mechanics*. Oxford University Press, Oxford, 1991.
- [103] Dennis Dieks and P. E. Vermaas, editors. *The Modal Interpretation of Quantum Mechanics*, Dordrecht, 1998. Kluwer Academic Publishers.
- [104] O. Lombardi and D. Dieks. Modal interpretations of quantum mechanics. In E. Zalta, editor, *The Stanford Encyclopedia of Philosophy (Spring 2017 edition)*. 2017.
- [105] Frank Arntzenius. Curiouser and curiouser: A personal evaluation of modal interpretations. In Dieks and Vermaas [103], pages 337–377.
- [106] Guido Bacciagaluppi. Delocalized properties in the modal interpretation of quantum mechanics. *Foundations of Physics*, 30:1431–1444, 2000.
- [107] D. Bohm. A Suggested Interpretation of Quantum Theory in Terms of “Hidden” Variables. *Physical Review*, 85:166–193, 1952.
- [108] J. S. Bell. Quantum Mechanics for Cosmologists. In C. J. Isham, R. Penrose, and D. Sciama, editors, *Quantum Gravity 2: a second Oxford Symposium*, Oxford, 1981. Clarendon Press. Reprinted in [85], pp.117–138; page references are to that version.

- [109] Patrick Duerr and Alexander Ehmman. The physics and metaphysics of Tychistic Bohmian Mechanics. *Studies in History and Philosophy of Science A*, 90:168–183, 2021.
- [110] John S. Bell. On the impossible pilot wave. *Foundations of Physics*, 12:989–999, 1982. Reprinted in [85], pp. 159–168.
- [111] Tim Maudlin. The nature of the quantum state. In Alyssa Ney and David Albert, editors, *The Wave Function*, pages 126–153. Oxford University Press, 2013.
- [112] V. Allori. Primitive ontology and the structure of fundamental physical theories. In *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, pages 58–75. Oxford University Press, New York, 2013.
- [113] Detlev Dürr, Sheldon Goldstein, and Nino Zanghí. Bohmian mechanics and the meaning of the wave function. In R. S. Cohen, M. Horne, and J. Stachel, editors, *Potentiality, Entanglement and Passion-at-a-Distance — Quantum Mechanical Studies in Honor of Abner Shimony*. Kluwer, Dordrecht, 1997. Available online at <http://arxiv.org/abs/quant-ph/9512031>.
- [114] D. Deutsch. Comment on Lockwood. *British Journal for the Philosophy of Science*, 47:222–228, 1996.
- [115] Harvey R. Brown and David Wallace. Solving the measurement problem: de Broglie-Bohm loses out to Everett. *Foundations of Physics*, 35:517–540, 2005.
- [116] Antony Valentini. De Broglie-Bohm pilot wave theory: Many worlds in denial? In *Many Worlds? Everett, Quantum Theory and Reality*. Oxford University Press, Oxford, 2010.
- [117] Harvey Brown. Reply to Valentini: ‘de Broglie-Bohm theory: Many worlds in denial?’. In *Many Worlds? Everett, Quantum Theory, and Reality*. Oxford University Press, Oxford, 2010.
- [118] Craig Callender. One world, one beable. *Synthese*, 192:3153–3177, 2015.
- [119] Antony Valentini and Hans Westman. Dynamical origin of quantum probabilities. Available online at <http://arxiv.org/abs/quant-ph/0403034>, 2004.
- [120] Antony Valentini. Inflationary cosmology as a probe of primordial quantum mechanics. *Physical Review D*, 82:063513, 2010.
- [121] Nicholas G. Underwood and Antony Valentini. Quantum field theory of relic nonequilibrium systems. *Physical Review D*, 92:063531, 2015.

- [122] Detlev Dürr, Sheldon Goldstein, and Nino Zanghí. Bohmian mechanics as the foundation of quantum mechanics. In James T. Cushing, Arthur Fine, and Sheldon Goldstein, editors, *Bohmian Mechanics and Quantum Theory: An Appraisal*, pages 21–44. Kluwer Academic Publishers, Dordrecht, 1996.
- [123] Sheldon Goldstein. Bohmian mechanics. In E. Zalta, editor, *The Stanford Encyclopedia of Philosophy (Summer 2017 edition)*. 2017. <https://plato.stanford.edu/archives/sum2017/entries/qm-bohm/>.
- [124] P. Holland. *The Quantum Theory of Motion*. Cambridge University Press, Cambridge, 1993.
- [125] Jean Bricmont. *Making Sense of Quantum Mechanics*. Springer, Berlin, 2016.
- [126] Detlev Dürr, Sheldon Goldstein, Roderich Tumulka, and Nino Zanghí. Bohmian mechanics and quantum field theory. *Physical Review Letters*, 93:090402, 2004.
- [127] Detlev Dürr, Sheldon Goldstein, Roderich Tumulka, and Nino Zanghí. Bell-type quantum field theories. *Journal of Physics*, A38:R1, 2005.
- [128] Ward Struyve. Field beables for quantum field theory. Available online at <http://arxiv.org/abs/0707.3685>, 2007.
- [129] Samuel Colin and Ward Struyve. A Dirac sea pilot-wave model for quantum field theory. *Journal of Physics A*, 40:7309–7342, 2007.
- [130] Ward Struyve and Hans Westman. A new pilot-wave model for quantum field theory. *AIP CONFERENCE PROCEEDINGS*, 844:321, 2006.
- [131] Ward Struyve and Hans Westman. A minimalist pilot-wave model for quantum electrodynamics. *Proceedings of the Royal Society of London A*, 463:3115–3129, 2007.
- [132] Roderich Tumulka. Collapse and relativity. In A. Bassi, T. Weber, and N. Zanghí, editors, *Quantum Mechanics: Are There Quantum Jumps? and on the Present Status of Quantum Mechanics*, page 340. American Institute of Physics Conference Proceedings, 2006. Available online at <http://arxiv.org/abs/quant-ph/0602208>.
- [133] Daniel J. Bedingham. Relativistic state reduction dynamics. Online at <http://arxiv.org/abs/1003.2774>, 2010.
- [134] Daniel J. Bedingham. Relativistic state reduction model. Online at <http://arxiv.org/abs/1103.3974>, 2011.
- [135] A. Einstein, B. Podolsky, and N. Rosen. Can quantum-mechanical description of reality be considered complete? *Physical Review*, 47:777–780, 1935.

- [136] Scott Aaronson. *Quantum Computing since Democritus*. Cambridge University Press, Cambridge, 2013.
- [137] J.P. Jarrett. On the significance of the locality conditions in the Bell argument. *Nous*, 18:569–589, 1984.
- [138] Abner Shimony. Events and processes in the quantum world. In Roger Penrose and Chris Isham, editors, *Quantum Concepts in Space and Time*, pages 182–203. Oxford University Press, 1986.
- [139] J.F. Clauser, M.A. Horne, A. Shimony, and R.A. Holt. Proposed experiment to test local hidden-variable theories. *Physical Review Letters*, 23:880–884, 1969.
- [140] J. S. Bell. Bertlmann’s socks and the nature of reality. *Journal de Physique*, 42:C2 41–61, 1981. Reprinted in [85], pp. 139–158.
- [141] Wayne Myrvold, Marco Genovese, and Abner Shimony. Bell’s theorem. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy (Fall 2021 edition)*. 2021.
- [142] Laura Ruetsche. Locality in (axiomatic) quantum field theory: A minority report. In Eleanor Knox and Alastair Wilson, editors, *The Routledge Companion to Philosophy of Physics*. Routledge, London, 2021.
- [143] Harvey Brown and Chris Timpson. Bell on Bell’s theorem: The changing face of nonlocality. In M. Bell and S. Gao, editors, *Quantum Nonlocality and Reality - 50 years of Bell’s Theorem*. Cambridge University Press, Cambridge, 2016.
- [144] Robert S. Cohen, M. Horne, and J.J. Stachel, editors. *Experimental Metaphysics: Quantum Mechanical Studies for Abner Shimony, Volume 1*. Kluwer, 1997.
- [145] Tim Maudlin. What Bell did. *Journal of Physics A*, 47:424010, 2014.
- [146] Guido Bacciagaluppi. Remarks on space-time and locality in Everett’s interpretation. In Jeremy Butterfield and Tomasz Placek, editors, *Nonlocality and Modality*, volume 64 of *Nato Science Series II. Mathematics, Physics and Chemistry*. Kluwer, 2002.
- [147] Wayne Myrvold. On peaceful coexistence: is the collapse postulate incompatible with relativity? *Studies in the History and Philosophy of Modern Physics*, 33:435–466, 2002.
- [148] Johannes *et al* Handsteiner. Cosmic Bell test: Measurement settings from milky way stars. *Physical Review Letters*, 118:060401, 2017.
- [149] Abellan *et al*. Challenging local realism with human choices. *Nature*, 557:212–216, 2018.

- [150] Anton Zeilinger. *Dance of the Photons: From Einstein to Quantum Teleportation*. Farrar, Straus and Giroux, New York, 2010.
- [151] Gerard 't Hooft. The cellular automaton interpretation of quantum mechanics. <https://arxiv.org/abs/1405.1548>, 2014.
- [152] Tim Palmer and Sabine Hossenfelder. Rethinking superdeterminism. *Frontiers in Physics*, 8:139, 2020.
- [153] Don A. Howard and Marco Giovanelli. Einstein's philosophy of science. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy (Fall 2019 Edition)*. 2019.
- [154] N. David Mermin. Could Feynman have said this? *Physics Today*, 57:10, 2004.
- [155] David Wallace. *Philosophy of Physics: a Very Short Introduction*. Oxford University Press, Oxford, 2021.
- [156] Tim Maudlin. *Philosophy of Physics: Space and Time*. Princeton University Press, Princeton, 2012.
- [157] B. DeWitt and N. Graham, editors. *The many-worlds interpretation of quantum mechanics*. Princeton University Press, Princeton, 1973.