Philosophical Issues raised by Quantum Theory and its Interpretations

Wayne C. Myrvold
Department of Philosophy
The University of Western Ontario
wmyrvold@uwo.ca

In Olival Friere, Jr., ed., *The Oxford Handbook of the History of Quantum Interpretations*, Oxford University Press, 2022.

Abstract

This chapter serves as an introduction to the philosophical issues raised by quantum theory. It begins with a brief overview of the formalism of quantum theory. The so-called "measurement problem" is introduced, and the main approaches to it surveyed. We then discuss the implications of quantum theory for metaphysics. One question concerns the implications of quantum nonlocality for our understanding of spacetime and of causality. Another has to do with the ontology of quantum states. Should these be regarded as physically real, and, if so, what sort of reality should be ascribed to them?

1 Introduction

The philosophical questions surrounding quantum theory revolve around the question: What, if anything, does the empirical success of quantum theory tell us about the physical world? Since the key papers formulating what we now call quantum mechanics were published in the years 1925–27, we are only a few years away from the centennial of the theory's inception. As we approach the centennial, there is more intense discussion than ever about its import.

Why is this? At the heart of the discussions is the following situation. We have in our textbooks, and teach to our students, what amounts to an operational recipe sufficiently precise for most applications. We learn to associate quantum states with various physical situations, and use them to calculate probabilities of outcomes of experiments. This is enormously important, as it is the basis both for the experimental testing of the theory and for its application. The success of the theory in these contexts is the reason we are taking it seriously at all. But the operational recipe does not, without further ado, yield anything like a clear description of what the physical systems to which it is applied are like, or what they are doing in between experiments.

The various approaches to the question of description of physical systems and processes (among which are those that hold that we should *refrain* from describing physical systems) are sometimes referred to as *interpretations* of quantum mechanics. This terminology is potentially misleading, for two reasons. The first is that it might suggest that the task of interpreting quantum theory is akin to supplying a model for an uninterpreted formal system. This is nothing at all like the task at hand. We don't have an uninterpreted formal system, or a mathematical theory devoid of physical significance. What we have is a formalism with an agreed-upon operational significance (or, at least, sufficiently close to agreed-upon for most applications). The question is what, if anything, is to be added to this operational core.

The second reason that the phrase "interpretations of quantum mechanics" is potentially misleading is that some of the avenues of approach involve formulation of a physical theory distinct from standard quantum theory, in some cases differing in empirical content. These are not merely different interpretations of a common theory, but alternate physical theories.

2 What is a quantum theory?

In this section quantum theories are briefly described, with an emphasis on the agreed-upon operational core, which every interpretational project must take into account. Quantum theories can be expressed in a number of different mathematical forms that are equivalent as far as the operational core is concerned. To avoid the pitfall of tying interpretational matters too closely to any particular formulation, we focus on what all the formulations have in common. This means eschewing, in the first instance, talk of Hilbert spaces or wave functions. Readers who find the presentation disconcertingly unfamiliar may be reassured that these can be introduced when desired.

A *quantum-mechanical theory* is a quantum theory of a system (such as a finite number of particles) having finitely many degrees of freedom. A quantum theory of a system with infinitely many degrees of freedom is a *quantum field theory*. We use the term *quantum theory* to embrace both quantum-mechanical theories and quantum field theories.

2.1 Constructing quantum theories

To construct a quantum theory, one identifies a system or systems of interest, and the dynamical variables that are to be modelled. The system could, for example, be the familiar textbook example of a hydrogen atom, with the variables to be modelled being the positions and momenta of a proton and an electron, and perhaps also their spins.

It is a peculiarity of quantum theories that, in order to formulate one, we begin with a classical theory and subject it to a procedure known as *quantization*. This typically begins with a Lagrangian or Hamiltonian formulation of a classical theory. In these formulations, the configuration of a system is represented by variables $\{q_1, \ldots, q_n\}$. These could, for example, be positions of a number of point particles, or they could specify the positions and angular orientations of a number of rigid bodies. One associates with each of these configuration variables q_i its *conjugate momentum* p_i. The complete physical state of a system is given by a specification of the values of these *canonical variables*, $\{(q_i, p_i) \,|\, i = 1, \ldots, n\}$. The set of all possible states is called the *phase space* of the system. Any dynamical variable of the system is a function of the canonical variables $\{(q_i, p_i)\}$.

A quantum theory is constructed by taking the canonical variables of a system and associating with them elements of an algebra $\mathcal{A}_Q$ with a non-commutative multiplication.[1] We will refer to the elements of this algebra as *operators*. For any two operators A, B, we define the

[1] It is an algebra over the complex numbers. This means that any operator can be multiplied by any complex number, that operators can be added and multiplied, and that multiplication distributes over addition.

commutator,

$$[A, B] = AB - BA. \tag{1}$$

When AB is equal to BA, A and B are said to *commute*.

The distinctive quantum relations are the *canonical commutation relations*. These specify commutators for the operators $\{(Q_i, P_i)\}$ that correspond to canonical variables $\{(q_i, p_i)\}$. The rules are,

- Operators corresponding to different degrees of freedom commute. This means that, for distinct i, j, Q_i and P_i commute with Q_j and P_j.
- $[Q_i, P_i] = i\hbar\,\mathbb{1}$, where $\mathbb{1}$ is the identity operator and $\hbar = h/2\pi$, where h is Planck's constant. The special operator $\mathbb{1}$ is the multiplicative identity; the result of multiplying any operator A by $\mathbb{1}$ is just A itself.

For any operator A, there is an operator $A^\dagger$, called the *adjoint* of A. These satisfy,

- $(A^\dagger)^\dagger = A$.
- For any operators A, B, and any complex numbers a, b,

  i). $(a\,A + b\,B)^\dagger = a^* A^\dagger + b^* B^\dagger$.

  ii). $(AB)^\dagger = B^\dagger A^\dagger$.

An operator that is its own adjoint is said to be *self-adjoint*. We can associate with any operator a set of real or complex numbers called its *spectrum*. If the operator is self-adjoint, its spectrum consists of real numbers only.

We associate with any experiment a self-adjoint operator whose spectrum is the set of possible values of the outcome variable. For an experiment that, classically, would be regarded as a measurement of a given quantity that is a function of the canonical variables $\{(q_i, p_i)\}$, the associated operator is the corresponding function of the operators $\{(Q_i, P_i)\}$ (this does not yield a unique prescription, but this is not a matter we will go into in this chapter).

We associate quantum states with preparation procedures that the system can be subjected to. A *quantum state* is an assignment ρ of numbers to operators, required to satisfy the conditions,

- *Positivity.* For any A, $\rho(A^\dagger A)$ is a non-negative real number.
- *Normalization.* $\rho(\mathbb{1}) = 1$.

- *Linearity.* For any complex numbers a, b,

$$\rho(a\,A + b\,B) = a\,\rho(A) + b\,\rho(B).$$

For self-adjoint A, the value of $\rho(A)$ is to be interpreted as the expectation value, in state ρ, of the outcome of an experiment with which is associated the operator A. The positivity condition ensures that self-adjoint operators are assigned real numbers. This condition, together with the normalization condition, ensures that, if the spectrum of a self-adjoint operator A is bounded, $\rho(A)$ is not above or below the bounds of the spectrum of A. This is required in order for these numbers to be interpreted as expectation values for the outcomes of experiments that yield results in the spectrum of A. The linearity condition is a non-trivial constraint, as it relates expectation values assigned to outcomes of incompatible experiments. It is a central principle of quantum mechanics, but is not something that is dictated by the operational significance of $\rho(A)$ as an expectation value of the outcomes of an experiment (one could imagine other, non-quantum theories that violate it).

For any state ρ, and any self-adjoint operator A, let a be the expectation value of A in state ρ, $\rho(A)$. We define the *variance* of A in state ρ as,

$$\mathrm{Var}_\rho(A) = \rho((A - a\mathbb{1})^2) = \rho(A^2) - \rho(A)^2. \tag{2}$$

This is one way to quantify the spread in the probability distribution of outcomes of an A-experiment. It is small if the distribution is tightly focussed near the expectation value $\rho(A)$. It is zero only when there is a single outcome that will be obtained with probability one. If this is the case—that is, if $\mathrm{Var}_\rho(A)$ is equal to zero—then ρ is said to be an *eigenstate* of A, with *eigenvalue* $\rho(A)$. For such a state, one in which there is a definite value of the observable corresponding to A that will with certainty be obtained as the outcome of an appropriate experiment, it is usual to ascribe the property of possessing this value to the system. This is known as the *eigenstate-eigenvalue link.* For example, a state that is an eigenstate of the operator corresponding to energy, with eigenvalue E, is taken to be a state in which the system has energy E. This has been, since the early days of quantum mechanics, a central interpretational principle of quantum theories.

Given any two states ρ_1, ρ_2, and any two positive numbers p_1 and p_2 that sum to one, we can always form the corresponding *mixture* of the states,

$$\bar{\rho} = p_1\,\rho_1 + p_2\,\rho_2. \tag{3}$$

An example of a preparation procedure with which a state like that would be associated is one that employs some randomizing device to choose between preparation of state ρ_1 and state ρ_2, with probabilities p_1 and p_2. Mixtures of more than two states are defined analogously. A state that is not a mixture of any two distinct states is called a *pure state*.

The content of the operational core of a quantum theory is completely encapsulated in the structure of the algebra $\mathcal{A}_Q$, the association of certain operators with experimental procedures, and of states with preparation procedures. However, for most purposes, this is not the most convenient formulation of the theory. It is often useful to construct a representation of the algebra as operators operating on vectors in a Hilbert space, in which any pure state ψ can be represented by a vector $|\psi\rangle$. If we're willing to countenance Hilbert spaces that are more capacious than they need to be, we can construct a representation in which *every* state, pure or mixed, is represented by state vector.

For any two distinct states ρ_1, ρ_2, represented by Hilbert-space vectors $|\psi_1\rangle$, $|\psi_2\rangle$, and any complex numbers a, b, there is another state that is represented by the vector

$$|\phi\rangle = a|\psi_1\rangle + b|\psi_2\rangle. \tag{4}$$

Vector addition is also referred to as *superposition*, and $|\phi\rangle$ is said to be a superposition of $|\psi_1\rangle$ and $|\psi_2\rangle$. Resist the temptation to talk as if some states are superpositions, and some are not. Any vector $|\phi\rangle$ is equal to infinitely many superpositions of other vectors.

If $|\psi_1\rangle$ and $|\psi_2\rangle$ are eigenvectors of some operator A, with eigenvalues a_1 and a_2, respectively, then the eigenstate-eigenvalue link tells us that in states represented by those vectors, the system has the corresponding properties. If a_1 and a_2 are distinct, then $|\phi\rangle$, as defined by (4), is *not* an eigenstate of A, and the eigenstate-eigenvalue link is silent on whether we are to ascribe any properties corresponding to A to the system in such a state. This is at the core of the so-called *measurement problem*, which will be presented in section 4.

In a quantum-mechanical theory of a system consisting of n particles without spin, a quantum-mechanical state can be represented by a square-integrable function $\psi(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n)$, called a *wave function*.

A wave function representing the state is not unique; any two functions that differ only on a set of measure zero represent the same state, and multiplying any wave function by any complex number yields another function that represents the same state. A wave function yields probabilities for outcomes of detection experiments as follows: the probability of finding particle 1 in a set Δ_1, particle 2 in Δ_2, *etc.*, is given by the integral of $|\psi|^2$ over the region of configuration space with x_1 in Δ_1 and x_2 in Δ_2, *etc.*, divided by the integral of $|\psi|^2$ over all of configuration space. For a system consisting of n particles with spin, the total spin state of the system can be represented by a vector in a finite-dimensional Hilbert space $\mathcal{H}_S$. A wave function for such a system is an assignment of a vector in $\mathcal{H}_S$ to each point in configuration space.

2.2 Entangled and unentangled states

Consider a quantum theory of two non-overlapping systems. The algebra of observables $\mathcal{A}_Q$ has two commuting subalgebras, $\mathcal{A}_A$ and $\mathcal{A}_B$, corresponding to the observables of the two subsystems. A state ρ is a *product state* if any only if

$$\rho(AB) = \rho(A)\rho(B) \tag{5}$$

for all A in $\mathcal{A}_A$ and B in $\mathcal{A}_B$. A state that is either a product state or mixture of product states is called a *separable state*. A state (pure or mixed) that is not a separable state is an *entangled state*.

We can also characterize pure entangled states more directly. A pure state ρ of $\mathcal{A}_Q$ is a product state if the restriction of ρ to $\mathcal{A}_A$ is a pure state of $\mathcal{A}_A$; it is an entangled state if the restriction of ρ to $\mathcal{A}_A$ is a mixed state of $\mathcal{A}_A$.

For any state that is not a product state, the state of a composite system is not uniquely determined by the states of the components, even if the state of the composite is pure. This is a striking difference between quantum and classical theories. For a classical theory, the restriction of any pure state—that is, a maximally specific state description—of a composite to one of its components is a pure state of the component, and specification of the states of the components uniquely determines the state of the composite. Following Howard (1985), this feature of classical theories has come to be known as *separability*, and the fact that it is not satisfied by quantum theories, as *nonseparability*.

2.3 Temporal evolution: Schrödinger and Heisenberg pictures

Suppose that we have a system whose dynamical variables are $\{(q_i(t), p_i(t))\}$. To construct a quantum theory of the system, we require operators $\{(Q_i(t), P_i(t))\}$ to represent the dynamical variables.

The dynamical laws of our quantum theory specify how expectation values of variables at different times are related to each other. The basic equation of evolution is,

$$i\hbar \frac{d}{dt}\rho(A(t)) = \rho(A(t)H - HA(t)), \tag{6}$$

where H is the operator corresponding to the system's Hamiltonian H.

Suppose, now, we want to construct a Hilbert space representation of our theory. This means assigning, to each operator $A(t)$, a Hilbert space operator $\hat{A}(t)$, and choosing, for each time t, a density operator $\hat{\rho}(t)$ to represent the state, in such a way that

$$\rho(A(t)) = \mathrm{Tr}[\hat{\rho}(t)\hat{A}(t)]. \tag{7}$$

As we have to specify both $\hat{A}(t)$ and $\hat{\rho}(t)$, this gives us some lee-way. One way to do this is to choose the same Hilbert space operators $(\hat{Q}_i, \hat{P}_i)$ to represent $(q_i(t), p_i(t))$ at all times. Then the density operators $\hat{\rho}(t)$ will have to satisfy,

$$i\hbar \frac{d}{dt}\hat{\rho}(t) = \hat{H}\hat{\rho}(t) - \hat{\rho}(t)\hat{H}. \tag{8}$$

For a pure state represented by a state vector $|\psi(t)\rangle$, we have,

$$i\hbar \frac{d}{dt}|\psi(t)\rangle = \hat{H}|\psi(t)\rangle. \tag{9}$$

This is the *Schrödinger equation*, and the choice of Hilbert space representation on which the operators representing $(q_i(t), p_i(t))$ are time-independent, is called the *Schrödinger picture*.

Another choice is to choose a fixed density operator $\hat{\rho}$ to represent the state at any time. This requires the operators $\hat{A}(t)$ to satisfy

$$i\hbar \frac{d}{dt}\hat{A}(t) = \hat{A}(t)\hat{H} - \hat{H}\hat{A}(t). \tag{10}$$

This is the *Heisenberg equation of motion*, and the Hilbert space representation on which the density operator representing the state is time-independent, is called the *Heisenberg picture*.

It should be emphasized that these two choices are two Hilbert space representations of *one and the same quantum theory*; the physical content is the same. It would be a mistake to take the time-independence of the Heisenberg-picture density operator to suggest that what is being modelled is an unchanging physical situation. What is changed is the location of the time dependence in the mathematical apparatus used to model a changing physical situation.

In some cases a clear meaning can be given to the idea that quantities associated with a system at different times are values, at different times, of the "same" dynamical variable. For example, in a theory set in Galilean spacetime one can choose a reference frame, and consider the position-coordinate of a particle, at different times, as differing values of the same dynamical variable. This, of course, requires a notion of trans-temporal identity of particles.

In other cases it may be inconvenient to do so. Consider a classical field theory on Minkowski spacetime. A particular solution of the field equations will specify, for each spacetime point, a field value at that point. One can pick a set of timelike lines that jointly cover the spacetime (whether or not these are inertial trajectories), and consider how the field value changes with position on a line. If one is considering the response of some device that monitors the field, it makes sense to consider how the field variable changes along the worldline of the detector. But the structure of the theory may be more perspicuously represented without such considerations, and without identifying any two field values as values of the "same" dynamical quantity at different times. For this reason, the Heisenberg picture is usually regarded as more suitable for relativistic quantum field theories.

3 The collapse postulate: some history

Textbook formulations of quantum mechanics usually include an additional postulate about how to assign a state vector after an experiment, according to which one replaces the quantum state with an eigenstate of the "measured" observable, corresponding to the result obtained. Unlike the unitary evolution applied otherwise, this is a discontinuous change of the quantum state, sometimes referred to as *collapse of the state vector*, or *state vector reduction*. There are two interpretations of the postulate about collapse, corresponding to two different conceptions of quantum states. If a quantum state represents nothing more than our knowledge about the system, then the collapse of the state to one corresponding to the observed result can be thought of as representing nothing more than an updating of knowledge. If, however, quantum states represent physical reality, in such a way that distinct pure states always represent distinct physical states of affairs, then the collapse postulate entails an abrupt, perhaps discontinuous, change of the physical state of the system. Considerable confusion can arise if the two interpretations are conflated.

The collapse postulate is found already in Heisenberg's *The Physical Principles of the Quantum Theory*, based on lectures presented in 1929 (Heisenberg 1930a, 27; 1930b, 36). Von Neumann, in his reformulation of quantum theory a few years later, distinguished between two types of processes: Process ***1.***, which occurs upon performance of an experiment, and Process ***2.***, the unitary evolution that takes place as long as no measurement is made (von Neumann 1932; 1955, §V.I). He does not take this distinction to be a difference between two physically distinct processes. Rather, the invocation of one process or the other depends on a somewhat arbitrary division of the world into an observing part and an observed part (see von Neumann 1932, 224; 1955, 420).

There is a persistent misconception that, for von Neumann, collapse is to be invoked only when a conscious observer becomes aware of the result. This is the opposite of his attitude; for him it is essential that the location of the boundary between the observed part of the world and the observing part is somewhat arbitrary. It may be placed between the system under the study and the experimental apparatus. On the other hand, we could include the experimental apparatus in the quantum description, and place the cut at the moment when light indicating the result hits the observer's retina. Or we could go further, and include the retina and relevant parts of the observer's nervous system in the quantum system. That the cut may be pushed arbitrarily far into the perceptual apparatus of the observer is required, according to von Neumann, by the principle of psycho-physical parallelism.

The collapse postulate does not appear in the first edition (1930) of Dirac's *Principles of Quantum Mechanics*; it is introduced in the second edition (1935), which appeared subsequent to von Neumann's treatment. Dirac, in contrast to Heisenberg and von Neumann, appears to take the distinction between unitary and collapse evolution to be a distinction between two physical processes. Also, for Dirac it is an act of *measurement*, not *observation*, that causes a system to "jump" into an eigenstate of the observable being measured (Dirac, 1935, 26). According to Dirac, this jump is caused by the interaction of the system with the experimental apparatus.

A formulation of a version of the collapse postulate according to which a measurement is not completed until the result is observed is found in London and Bauer (1939). For them, as for Heisenberg, this is a matter of an increase of knowledge on the part of the observer.

Wigner (1961) combined elements of the two interpretations. Like those who take the collapse to be a matter of updating of belief in light of information newly acquired by an observer, he takes collapse to take place when a conscious observer becomes aware of an experimental result. However, like Dirac, he takes it to be a real physical process. His conclusion is that consciousness has an influence on the physical world not captured by the laws of quantum mechanics. This involves a rejection of von Neumann's principle of psycho-physical parallelism, according to which it must be possible to treat the process of subjective perception as if it were a physical process like any other.[2]

4 The so-called "measurement problem"

If it is possible for any quantum theory to be a comprehensive physical theory, it must be capable of treating of our experimental apparatus, and, indeed, everything else. Suppose, then, we analyze an experimental set-up quantum-mechanically.

Let S be the system to be experimented on, which we call the studied system. We suppose that it has at least two distinguishable states, $|+\rangle_S$ and $|-\rangle_S$, and that an apparatus, A, can be devised that distinguishes these states. This means that there are distinguishable sets A^+ and A^- of states of the apparatus, which we will call *indicator states*, such that the apparatus can be coupled to the system S in such a way that, if the apparatus is started out in a ready state and the

[2] Despite this, Wigner's proposal is sometimes wrongly attributed to von Neumann, and is sometimes called the "von Neumann-Wigner interpretation."

studied system in the state $|+\rangle_S$, the apparatus will evolve to a state in A^+, and, if the apparatus is started out in a ready state and the studied system in the state $|-\rangle_S$, the apparatus will evolve to a state in A^-. We do not assume that the apparatus is isolated from its environment, and so we include the relevant parts of the environment in our description.

The evolution of the composite system should satisfy,

$$\begin{aligned} |+\rangle_S |R\rangle_{AE} &\Rightarrow |`+'\rangle_{SAE}; \\ |-\rangle_S |R\rangle_{AE} &\Rightarrow |`-'\rangle_{SAE}; \end{aligned} \tag{11}$$

where $|R\rangle_{AE}$ is a state of the apparatus plus its environment in which the apparatus is ready to perform an experiment, and $|`+'\rangle_{SAE}$ and $|`-'\rangle_{SAE}$ are states of the composite system in which the apparatus is indicating the results $+$ and $-$, respectively.

Suppose, now, that the system A is started out in a state that is a nontrivial superposition of $|+\rangle_S$ and $|-\rangle_S$,

$$a|+\rangle_S + b|-\rangle_S,$$

with a and b both nonzero. If the evolution of the composite system is linear, as it must be if the Schrödinger equation applies, then we must have,

$$(a|+\rangle_S + b|-\rangle_S)|R\rangle_{AE} \Rightarrow a|`+'\rangle_{SAE} + b|`-'\rangle_{SAE}. \tag{12}$$

What to make of a state like this? It is not a state in which the apparatus is in either one of the indicator sets A^+ or A^-. It is not an eigenstate of the apparatus indicator variable, but a superposition of distinct eigenstates. The eigenstate-eigenvalue link, therefore, offers no guidance. The interaction of the apparatus with its environment may result in an entangled state of the apparatus and its environment that is such that the state of the apparatus is a mixture of distinct indicator states—a process known as *decoherence*—but this does not help with the interpretational issue, as the system containing the apparatus and a sufficiently wide portion of its environment is still a superposition of macroscopically distinct states. It is often said that a state like (12) conflicts with our experience, according to which experimental apparatus is, at the end of an experiment, always indicating some determinate result. This is highly misleading, because the real issue is whether we can make sense of (12) as a possible state of a system containing the experimental apparatus. If we don't know what it would be like to find the apparatus in such a state, then it makes no sense either to affirm or to deny that we have ever found the apparatus in such a state.

Though we have chosen an experimental set-up as an illustration, situations in which linear evolution would yield superpositions of macroscopically distinct states are ubiquitous. Nonetheless, the problem of what to make of this fact—that applying linear evolution to quantum states involving macroscopic objects will lead to superpositions of macroscopically distinct states—has come to be known as the *measurement problem.*

If there is a unique outcome of the experiment, and if (12) is the correct quantum state, then the outcome fails to be represented by the quantum state, which must be supplemented by something that *does* indicate the outcome. On the other hand, it might be that neither Schrödinger evolution nor any other linear evolution applies to situations like the one envisaged, and that the correct evolution leads to a state that we *can* take to be indicating a determinate outcome. These two options were summarized by J.S. Bell in his remark, "Either the wavefunction, as given by the Schrödinger equation, is not everything, or it is not right" (Bell 1987a, 41; 1987b and 2004, 201). This gives us a (*prima facie*) neat way of classifying approaches to the so-called "measurement problem."

- There are approaches that involve a denial that a quantum wave function (or any other way of representing a quantum state) yields a complete description of a physical system.
- There are approaches that involve modification of the dynamics to produce a collapse of the quantum state in appropriate circumstances.
- There are approaches that reject both horns of Bell's dilemma, and hold that quantum states undergo unitary evolution at all times and that there is no more to be said about the physical state of a system than can be represented by a quantum state.

We include in the first category approaches that deny that a quantum state should be thought of as representing anything in physical reality at all. If quantum states do not represent anything, and if there is something rather than nothing, then quantum states do not represent everything. In this category are Bohrian approaches, according to which there are principled reasons not to seek a complete description, and Einsteinian approaches, according to which seeking a theory that need not leave anything out in its descriptions is a project worthy of pursuit.

Also included in the first category are approaches that take quantum states to represent something, but not everything, in physical reality. These include "hidden-variables" theories, and modal interpretations (see Lombardi and Dieks 2017). The best-known and most thoroughly worked-out theory of this sort is the de Broglie-Bohm pilot wave theory, which takes particles with definite trajectories as the basic ontology. The role of the wave function is to provide dynamics for the particles. See Bacciagaluppi and Valentini (2009) for a historical introduction, and Dürr et al. (1992) and Pearle and Valentini (2006) for current perspectives.

The second category embraces the dynamical collapse theory programme, which seeks a modified dynamics that approximates unitary evolution in the domains in which we have good evidence for its correctness, and approximates collapse in other situations, including, but not limited to, experimental set-ups. The best-known version of this is the Ghirardi-Rimini-Weber (GRW) theory (Ghirardi, Rimini, and Weber 1986), referred to by its creators as *Quantum Mechanics with Spontaneous Localization* (QMSL). On this theory, Schrödinger evolution of the quantum state is punctuated by discontinuous jumps. The GRW theory has the defect that it does not respect the symmetrization/antisymmetrization requirements for states of a system containing identical particles. This is remedied in a successor theory, the *Continuous Spontaneous Localization* (CSL) theory (Pearle 1989; Ghirardi, Pearle, and Rimini 1990).

Approaches that reject both horns of Bell's dilemma are typified by Everettian, or "many-worlds" interpretations. The basic idea is denial that there is a unique experimental outcome; rather, there is a splitting, and different results obtain on different branches of the multiverse. These have their roots in the work of Hugh Everett III (see Barrett and Byrne 2012). See Saunders et al. (2010), Wallace (2012), and Carroll and Singh (2019) for some recent approaches along these lines.

An approach that does not fit neatly into these categories is the *relational interpretation* advocated by Carlo Rovelli. It is akin in some ways to Everett's original conception, which he called the *relative-state* interpretation. It differs from it in not taking quantum states to be representational. For more on this, see Rovelli's contribution to this volume, and also Laudisa and Rovelli (2019), and references therein.

5 Bell's theorem, nonlocality, and relativity

"Bell's theorem" is a term used for a family of theorems of the following form. From a condition on probabilities that is motivated, in part, by locality considerations, an inequality is derived constraining correlations between results of spatially separated experiments, which is violated by the predictions of quantum mechanics. See the *Stanford Encyclopedia of Philosophy* entry (Myrvold, Genovese, and Shimony 2019) for an overview more detailed than found here.

The distinctive condition needed to derive Bell-type inequalities is the condition that correlations between outcomes of spatially separated systems be *locally explicable*. This has two parts: that the correlations be explicable in a certain sense, and the explanation be local. The explicability condition was taken by Bell to be the condition that correlations between events that are not in a direct cause-effect relation with each other be attributable to a common cause (see Bell 1981, C2–55; 1987b and 2004, 152). This is a version of a principle that has been called by Reichenbach (1956, §56) the *Principle of the Common Cause*, and, though Reichenbach made no pretense to originality, has been called *Reichenbach's Common Cause Principle*. The Common Cause Principle, together with the assumption that experimental outcomes at spacelike separation are not in a direct cause-effect relation with each other, yields the condition that Shimony (1986; 1990) has called *outcome independence*. Causal locality requires that a choice of experiment made at one location does not affect the probabilities of outcomes of another experiment performed at spacelike separation. This condition is called *parameter independence*. In some of his writings Bell combined locality and causality considerations in a principle he called "The Principle of Local Causality" (Bell, 1976, 1990).

In addition to this condition of local explicability, there are supplementary assumptions of the sort taken for granted in all experimental science, such as the assumption that it is possible, via some randomizing procedure, to render one's choice of experiment statistically independent of the state of the system on which the experiment is done. This condition is referred to as a "no-conspiracies" assumption, or "measurement independence," or, in some of the recent literature, "statistical independence." Though the assumption has been denied by some, we will in what follows restrict our attention to views consistent with this assumption.

Violation of Bell inequalities has been abundantly confirmed by experiment. What does this tell us?

It is sometimes said that violation of Bell inequalities straightforwardly entails violation of relativistic causality. Things are not so simple, as there is no interpretation-independent answer to the question of compatibility with relativity.

The question is most straightforward in connection with hidden-variables theories such as the de Broglie-Bohm theory. Any deterministic theory that violates Bell inequalities must violate parameter independence, and thus must have cause-effect dependencies between spacelike separated events.

Because, in a multi-particle system, the velocity of each particle may depend on the positions of all the others, the de Broglie-Bohm theory requires a preferred relation of distant simultaneity for its formulation. There is a series of theorems that show that any theory of this sort, on which the quantum state is supplemented by extra variables that are required to have a probability distribution given by the Born rule, must employ a distinguished relation of distant simultaneity, as it is not possible to satisfy the postulate about probabilities on arbitrary spacelike hypersurfaces. See Berndl et al. (1996); Dickson and Clifton (1998); Arnztenius (1998); Myrvold (2002, 2009). This has the consequence that such theories require a dynamically distinguished relation of distant simultaneity; see Myrvold (2021, §5.5.1) for the argument.

Dynamical collapse theories, on the other hand, do not require a preferred relation of distant simultaneity for their formulation. There is an extension of the GRW theory to a relativistic context (Dove, 1996; Dove and Squires, 1996; Tumulka, 2006), which involves a fixed, finite number of noninteracting particles. There are also extensions of the CSL theory to the context of relativistic quantum field theories (Bedingham, 2011a,b; Pearle, 2015).

These theories involve probabilistic correlations between spacelike separated events that are *not* attributable to events in their common past. That is, they involve a rejection of the Reichenbach Common Cause Principle. The question of whether a theory such as this is in violation of any restriction on causal relations that is motivated by considerations of special relativity has been a hotly debated one. Several authors have argued over the years, in different ways, in favour of the compatibility of theories like that with the requirements of special relativity; these include Shimony (1978, 1984, 1986), Jarrett

(1984), Skyrms (1984), Redhead (1987), Ghirardi and Grassi (1996), and Ghirardi (2012). See Myrvold (2016) for a recent argument for compatibility of special relativity with violations of Bell inequalities.

6 Ontological questions concerning quantum states

6.1 The question of quantum state realism

We have introduced quantum states via their operational significance: they encode probabilities of outcomes of experiments. Should we think of them as representing some feature of the system to which they are ascribed?

Positions that deny that quantum states represent features of physical reality have a history as old as quantum theory itself. This is one thing that Bohr and Einstein agreed upon. For Bohr, all description of physical reality must be couched in classical terms, and the limits of classical physics are the limits of physical description; quantum wave functions have only "symbolic" status (see Bohr 1934, 17). Einstein argued, in several places (see, *e.g.*, Einstein 1936), that quantum states should be regarded as akin to the probability distributions of classical statistical mechanics, that is, as representing incomplete knowledge of some deeper underlying physical state. The chief locus of difference between the two had to do with the propriety of seeking a deeper level of description.

The idea that quantum states are like that is an attractive one. It faces considerable obstacles, and it should be non-controversial that quantum states are not *just* like classical probability distributions.

A useful way of sharpening the question of realism about quantum states is afforded by the framework constructed by Harrigan and Spekkens (2010). This framework makes explicit some principles that are deeply embedded in our reasoning about the world.

Suppose that Alice has a choice of two or more preparation procedures that she can subject a system to. Having made the choice, she passes the system on to Bob, who can do an experiment on the system, and, from the outcome, reliably identify the procedure Alice has chosen. We would take this as an indication that distinct choices of preparation on Alice's part result in physical differences in the system being prepared, and that the outcome of Bob's experiment is informative about these differences.

Suppose, now, that we begin to consider what a theoretical model of a set-up like this would look like. This would involve some set Λ of possible physical states of the system (Harrigan and Spekkens call this the *ontic state space*). We associate with a preparation procedure ψ a probability distribution P_ψ over appropriate subsets of Λ. Suppose that Bob's experiment has potential outcomes $\{o_1, \ldots, o_k\}$. We associate with Bob's experiment *response probability functions* $f_k(\lambda)$. The function $f_k(\lambda)$ yields the probability that the kth outcome of the experiment is obtained, if the state of the system experimented on is λ.

We say that the preparations ψ and ϕ are *distinguishable* if there is an experiment whose outcome discriminates between them with certainty. If ψ and ϕ are distinguishable, then (as one would expect) there is no set of states that has nonzero probability of being realized on both preparations. Call a pair of preparations ψ, ϕ *ontically distinct* if there is no subset of the state space Λ that has nonzero probability of being realized on both preparations.

Preparations corresponding to distinguishable quantum states must be ontically distinct. The question to be addressed is whether, for all pairs of distinct pure quantum states, including those that are not distinguishable, the corresponding preparations are ontically distinct. Harrigan and Spekkens say that a theory is *ψ-ontic* if, according to the theory, preparations corresponding to distinct pure quantum states are always ontically distinct. They define *ψ-epistemic* as the negation of ψ-ontic. This is potentially misleading terminology. Consider, for example, a classical system, whose ontic state is represented by a point in its phase space. Suppose that one could learn either its position, or its momentum, but not both, though it always *has* determinate position and momentum. Any position is compatible with any momentum, and hence, for any position x and momentum p, the set of ontic states corresponding to position x overlaps with the set of states corresponding to momentum p. That doesn't mean that there is anything epistemic about position or momentum. Furthermore, to call a model "ψ-epistemic" if there are distinct pure quantum states whose associated probability distributions have *some* overlap, no matter how small, is potentially misleading, as it might suggest that the goal of constructing an interpretation on which quantum states are like classical probability distributions has been achieved. This, however, would require that the model be what has been called a *maximally*

ψ-epistemic model (Barrett et al., 2014). On such a model, the indistinguishability of quantum states is fully explained by overlap of the corresponding probability distributions on ontic state space.

There are a number of theorems concerning the viability of the programme of constructing a theory that is maximally ψ-epistemic, or, failing that, a theory that is not ψ-ontic. In particular, Barrett et al. (2014) show that no theory that reproduces quantum probabilities for outcomes of experiments and fits into the framework just sketched can be maximally ψ-epistemic, or even come close to being so. Pusey, Barrett, and Rudolph (PBR) show that, provided that the theory satisfies a postulate called the *Preparation Independence Postulate*, it must be ψ-ontic in order to reproduce quantum probabilities for outcomes of experiments (Pusey, Barrett, and Rudolph 2012).

The Preparation Independence Postulate is a postulate to the effect that it is possible to subject a pair of distinct systems A and B to preparation procedures that render their ontic states probabilistically independent of each other. This postulate involves an assumption, called the *Cartesian Product Assumption*, to the effect that, for a preparation of that sort, the state of the composite system AB can be fully represented by specifying a state of A and a state of B. This is a non-trivial restriction on the state spaces employed in the theory. A weaker assumption, called the *Preparation Uninformativeness Condition*, which makes no assumptions about the structure of the state spaces, was suggested by Myrvold (2018c, 2020). On the basis of this weaker assumption, a weaker conclusion can be derived. The conclusion that is derived from this condition is that, on any theory that satisfies it, pure quantum states $|\psi\rangle$ and $|\phi\rangle$ that are not too close to each other are ontically distinct. Here, the condition of not being too close is that the absolute value of their inner product be less than $1/\sqrt{2}$.

6.2 The ontological status of quantum states

Suppose that we are realists about quantum states. This means that distinct pure quantum states represent physically distinct states of affairs. This still leaves open the question of what sorts of physical reality these states represent. In this section we briefly discuss some options.

6.2.1 Quantum state monism

Could there be nothing more to the world than what is represented by a quantum state?

Recall that a quantum theory is not an uninterpreted formalism. A quantum theory involves an identification of physical quantities to be represented, and an association of operators with those quantities. The eigenstate-eigenvalue link yields property attributions in the special case of eigenstates. If we had a dynamical collapse theory that produced eigenstates of the right sorts of dynamical quantities—if, for example, it yielded definite mass or energy content for regions of space that are small on the macroscopic scale—then such a theory could, in a straightforward way, be a quantum state monist theory. Sometimes skepticism is expressed about this, but this skepticism seems aimed at a different project, a project that would involve starting with a mathematical formalism devoid of physical interpretation and attempting to interpret it physically.

Things are not so simple, because dynamical collapse theories do not produce eigenstates of appropriate physical quantities, and there are principled reasons for not expecting a dynamical collapse theory to do that. For this reason, Ghirardi and collaborators proposed a weakening of the eigenstate-eigenvalue link, according to which a system is to be ascribed a property if its quantum state is *sufficiently close* to being an eigenstate of the corresponding operator (Ghirardi, Grassi, and Pearle 1990, 1298; see also Ghirardi, Grassi, and Benatti 1995, 13). This modification has been dubbed, by Clifton and Monton (1999), the *fuzzy link*.[3] For a defense of quantum state monism along the lines proposed by the originators of the GRW and CSL theories, see Myrvold (2018a, 2019).

Everettian theories seem to be best interpreted along these lines. Since such theories eschew collapse, on such a theory a quantum state will not typically be anywhere near an eigenstate of familiar macroscopic variables. However, the quantum state of any bounded region

[3] Peter Lewis (2016, 86–90) distinguishes between a *fuzzy link*, according to which there is some precise threshold p such that a system possesses the property $A = a$ if and only if the probability of finding some other value is less than p, and a *vague link* according to which possession of a definite property is a matter of degree. It is hard to imagine what arguments there could be (or even what it might mean) for there to be a precise threshold. Albert and Loewer (1996) argue, correctly in my opinion, that there could be no such precise threshold, and that the modified link must be somewhat vague. This is what I mean by a fuzzy link.

of spacetime will typically be a *mixture* of states in which macroscopic variables are near-definite, and the terms of these mixtures will evolve independently, and can be taken to represent quasi-classical domains.

6.2.2 The project known as "wave function realism"

Consider again a quantum theory of n distinguishable particles, and suppose that we choose to represent quantum states in this theory by wave functions on the configuration space of the particles (recall from section 2.1 that this is optional). The wave function representation is not unique; a wave function representing a quantum state ρ is represented by a class of functions, not a single one; two functions that differ only on a set of measure zero, or differ only by a multiplicative constant, represent the same quantum state.

When de Broglie introduced the precursors of our quantum-mechanical wave functions, the thought was that these would be akin to electromagnetic fields. A stumbling-block for an interpretation of this sort is that multi-particle wave functions are functions of n points in space, and have to be, to encode correlations between the positions of particles. That is, they are what Belot (2012) has called "multi-fields."

Suppose, however, one wanted—perhaps out of a commitment to separability—to construct an alternative to standard quantum theory according to which the basic ontology consisted of a specification of local conditions at every point of some fundamental space (see Albert 1996). This is a project that has come to be known as *wave function realism*—a misleading terminology, as every form of realism about quantum states will maintain that wave functions, as one way of representing quantum states, represent something physically real.

The basic idea is that quantum theories are to be embedded in a more encompassing framework that would include theories with no relation to quantum theories as we have characterized them. We are to imagine some class of dynamical laws for fields in this framework that are such that for some—but not all—choices of dynamics the evolution of the field on a $3n$-dimensional fundamental space will mimic the evolution of a wave function defined at n points in a 3-dimensional space; the theory restricted to evolutions of the sort will then be functionally equivalent to a quantum theory.

I describe this as a *project* because those engaged in it have not yet said in any detail what the field on the fundamental space is thought to be, except in the case of the nonrelativistic theory of n distinguishable particles. One can perhaps dimly see how the case of particles with spin is to be handled, and even more dimly, that of a quantum field theory (see Myrvold 2015 for some options). It is also unclear what sort of structure is presumed for the fundamental space: does it have some built-in metric or causal structure, or is this to be regarded as emergent also?

The project should be regarded as a work-in-progress. See Ney and Albert (2013) for a collection of essays connected with this project. Among the matters that require clarification are the goals and motivations of the project. For some discussion of this, see Ney (2019, 2021).

6.2.3 Primitive ontology, and the nomic view of quantum states

As mentioned above, the originators of dynamical collapse theories originally advocated a quantum state monist ontology, with a modified version of the eigenstate-eigenvalue link. In recent years this proposal has somewhat fallen out of favour, to be replaced with a "primitive ontology" approach on which a theory must posit some basic ontology, which is the stuff of which ordinary objects are made, and the role of quantum states is to provide dynamics for that ontology. On this view, what quantum states represent physically is more like a dynamical law than "stuff" as usually construed. See Allori et al. (2008, 2014) and Allori (2013) for exposition and defense of this view. This sort of conception applies most straightforwardly to the de Broglie-Bohm theory, on which the primitive ontology is particles with definite trajectories, and the wave function is a guiding wave.

One advantage of this view is that it makes transparent how quantum wave functions should transform under dynamical symmetries. If one starts with the primitive ontology, and takes the role of the wave function to provide dynamics for it, then the condition that the set of dynamically possible trajectories be invariant under a symmetry operation yields guidance as to how the mathematical representation of the wave function should change under symmetry transformations.

It should be emphasized that the nomic view of quantum states is a *realist* conception of quantum states, in the sense used in section 6.1. Preparation procedures have an influence on quantum states, and there is a matter of fact about *which* quantum state has been prepared. Where the nomic view differs from other conceptions of quantum states has to do with what sort of matter of fact it is.

6.2.4 A comment on "Spacetime State Realism"

Wallace and Timpson (2010) introduce the term "Spacetime State Realism," which they characterize as "a view which takes the states associated to spacetime regions as fundamental." This strikes me as misleading, as the terminology suggests that what is being proposed is some alternative to other sorts of realism about quantum states. But recall that a quantum state is an assignment of expectation values to operators representing physical quantities. If the basic quantities are taken to be ones that pertain to bounded spacetime regions (which are the ones that are relevant to outcomes of experiments, as these take place within bounded spacetime regions), then for any spacetime region there is an associated algebra, and hence an associated state. This is explicit in the algebraic formulation of quantum field theory, but any presentation of quantum theory will have to make sense of "observables" associated with bounded spacetime regions. Spacetime state realism, then, is simply realism about quantum states, and a view that takes states associated to spacetime regions as fundamental is simply a view that takes quantum states as fundamental.

7 Conclusion

This has by no means been an exhaustive overview of the philosophical discussions surrounding quantum mechanics. Among topics not touched upon are interpretational and conceptual issues peculiar to quantum field theories. We have said little about the arguments offered by those who reject realism about quantum states in favour of such a position. Also omitted is the project of reconstructing quantum theory on operational or information-theoretical principles, and any detailed discussion of the classical-quantum interface. The *Stanford Encyclopedia of Philosophy* article, "Philosophical Issues in Quantum Theory" (Myrvold, 2018b) provides some pointers to the relevant literature on these topics.

References

Albert, D. and B. Loewer (1996). Tails of Schrödinger's cat. In R. K. Clifton (Ed.), *Perspectives on Quantum Reality*, 81–92. Dordrecht: Kluwer Academic Publishers.

Albert, D. Z. (1996). Elementary quantum metaphysics. In Cushing, Fine, and Goldstein (Eds.), 277–284. Dordrecht: Kluwer.

Allori, V. (2013). Primitive ontology and the structure of fundamental physical theories. In Ney and Albert (Eds.), 58–75.

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

Allori, V., S. Goldstein, R. Tumulka, and N. Zanghì (2014). Predictions and primitive ontology in quantum foundations: A study of examples. *The British Journal for the Philosophy of Science 65*, 323–352.

Arnztenius, F. (1998). Curiouser and curiouser: A personal evaluation of modal intepretations. In Dieks and Vermaas (Eds.), 337–377.

Bacciagaluppi, G. and A. Valentini (2009). *Quantum Theory at the Crossroads: Reconsidering the 1927 Solvay Conference*. Cambridge: Cambridge University Press.

Barrett, J., E. G. Cavalcanti, R. Lal, and O. J. E. Maroney (2014). No ψ-epistemic model can fully explain the indistinguishability of quantum states. *Physical Review Letters 112*, 250403.

Barrett, J. A. and P. Byrne (Eds.) (2012). *The Everett Interpretation of Quantum Mechanics: Collected Works 1955–1980 with Commentary*. Princeton: Princeton University Press.

Bedingham, D. (2011a). Relativistic state reduction model. *Journal of Physics: Conference series 306*, 012034.

Bedingham, D. (2011b). Relativistic state reduction dynamics. *Foundations of Physics 41*, 686–704.

Bell, J. S. (1976). The theory of local beables. *Epistemological Letters 9*, 11–24. Reprinted in *Dialectica 39* (1985), 86–96, and in Bell 1987b, 2004, 52–62.

Bell, J. S. (1981). Bertlmann's socks and the nature of reality. *Journal de Physique 42*, 611–37. Reprinted in Bell 1987b, 2004, 139–158.

Bell, J. S. (1987a). Are there quantum jumps? In C. Kilmister (Ed.), *Schrödinger: Centenary celebration of a polymath*, 41–52. Cambridge: Cambridge University Press. Reprinted in Bell 1987b, 2004, 201–212.

Bell, J. S. (1987b). *Speakable and Unspeakable in Quantum Mechanics*. Cambridge: Cambridge University Press.

Bell, J. S. (1990). La nouvelle cuisine. In A. Sarlemijn and P. Kroes (Eds.), *Between Science and Technology*. North-Holland. Reprinted in Bell 2004, 232-48.

Bell, J. S. (2004). *Speakable and Unspeakable in Quantum Mechanics* (2nd ed.). Cambridge: Cambridge University Press.

Belot, G. (2012). Quantum states for primitive ontologists: A case study. *European Journal for Philosophy of Science 2*, 67–83.

Berndl, K., D. Dürr, S. Goldstein, and N. Zanghì (1996). Nonlocality, Lorentz invariance, and Bohmian quantum theory. *Physical Review A 53*, 2062–2073.

Bohr, N. (1934). Introductory survey (1929). In *Atomic Theory and the Description of Nature*, 1–24. Cambridge: Cambridge University Press.

Carroll, S. M. and A. Singh (2019). Mad-dog Everettianism: Quantum mechanics at its most minimal. In A. Aguirre, B. Foster, and Z. Merali (Eds.), *What is Fundamental?*, 95–104. Cham, Switerland: Springer.

Clifton, R. and B. Monton (1999). Losing your marbles in wavefunction collapse theories. *The British Journal for the Philosophy of Science 50*, 697–717.

Cushing, J. T., A. Fine, and S. Goldstein (Eds.) (1996). *Bohmian Mechanics and Quantum Mechanics: An Appraisal.* Dordrecht: Kluwer.

Dickson, M. and R. Clifton (1998). Lorentz-invariance in modal interpretations. In Dieks and Vermaas (Eds.), 9–47.

Dieks, D. and P. E. Vermaas (Eds.) (1998). *The Modal Interpretation of Quantum Mechanics.* Dordrecht: Kluwer Academic Publishers.

Dirac, P. A. M. (1935). *Principles of Quantum Mechanics* (2nd ed.). Oxford: Oxford University Press.

Dove, C. and E. J. Squires (1996). A local model of explicit wavefunction collapse. `arXiv:quant-ph/9605047v1`.

Dove, C. J. (1996). *Explicit wavefunction collapse and quantum measurement.* Ph. D. thesis, Durham University.

Dürr, D., S. Goldstein, and N. Zanghì (1992). Quantum equilibrium and the origin of absolute uncertainty. *Journal of Statistical Physics 67*, 843–907.

Einstein, A. (1936). Physik und Realität. *Journal of the Franklin Institute 221*, 349–382. English translation in Einstein (1954).

Einstein, A. (1954). Physics and reality. In *Ideas and Opinions*, 290–323. New York: Crown Publishers, Inc. Translation of Einstein (1936).

Ghirardi, G. (2012). Does quantum nonlocality irremediably conflict with special relativity? *Foundations of Physics 40*, 1379–1395.

Ghirardi, G. and R. Grassi (1996). Bohm's theory versus dynamical reduction. In Cushing, Fine, and Goldstein (Eds.), 353–377.

Ghirardi, G., R. Grassi, and P. Pearle (1990). Relativistic dynamical reduction models: General framework and examples. *Foundations of Physics 20*, 1271–1316.

Ghirardi, G. C., R. Grassi, and F. Benatti (1995). Describing the macroscopic world: Closing the circle within the dynamical reduction program. *Foundations of Physics 25*, 5–38.

Ghirardi, G. C., P. Pearle, and A. Rimini (1990). Markov processes in Hilbert space and continuous spontaneous localization of systems of identical particles. *Physical Review A 42*, 78–89.

Ghirardi, G. C., A. Rimini, and T. Weber (1986). Unified dynamics for microscopic and macroscopic systems. *Physical Review D 34*, 470–491.

Harrigan, N. and R. W. Spekkens (2010). Einstein, incompleteness, and the epistemic view of quantum states. *Foundations of Physics 40*, 125–157.

Heisenberg, W. (1930a). *Die Physikalische Prinzipien der Quantentheorie*. Leipzig: Verlag von S. Hirzel. English translation in Heisenberg 1930b.

Heisenberg, W. (1930b). *The Physical Principles of the Quantum Theory*. Chicago: University of Chicago Press. Translation, by Carl Eckart and Frank C. Hoyt, of Heisenberg 1930a.

Howard, D. (1985). Einstein on locality and separability. *Studies in History and Philosophy of Science 16*, 171–201.

Jarrett, J. (1984). On the physical significance of the locality conditions in the Bell arguments. *Noûs 18*, 569–589.

Laudisa, F. and C. Rovelli (2019). Relational quantum mechanics. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2019 ed.). Metaphysics Research Lab, Stanford University.

Lewis, P. J. (2016). *Quantum Ontology: A Guide to the Metaphysics of Quantum Mechanics*. Oxford: Oxford University Press.

Lombardi, O. and D. Dieks (2017). Modal interpretations of quantum mechanics. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2017 ed.). Metaphysics Research Lab, Stanford University.

Lombardi, O., S. Fortin, C. López, and F. Holik (Eds.) (2019). *Quantum Worlds: Perspectives on the Ontology of Quantum Mechanics*. Cambridge: Cambridge University Press.

London, F. and E. Bauer (1939). *La Théorie de l'Observation en Méchanique Quantique*. Paris: Hermann & Co. English translation in London and Bauer 1983.

London, F. and E. Bauer (1983). The theory of observation in quantum mechanics. In J. A. Wheeler and W. H. Zurek (Eds.), *Quantum Theory and Measurement*, 217–259. Princeton: Princeton University Press. Translation of London and Bauer 1983.

Myrvold, W., M. Genovese, and A. Shimony (2019). Bell's theorem. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2019 ed.). Metaphysics Research Lab, Stanford University.

Myrvold, W. C. (2002). Modal intepretations and relativity. *Foundations of Physics 32*, 1773–1784.

Myrvold, W. C. (2009). Chasing chimeras. *The British Journal for the Philosophy of Science 60*, 635–646.

Myrvold, W. C. (2015). What is a wave function? *Synthese 292*, 3247–3274.

Myrvold, W. C. (2016). Lessons of Bell's theorem: Nonlocality, yes; action at a distance, not necessarily. In M. Bell and S. Gao (Eds.), *Quantum Nonlocality and Reality: 50 Years of Bell's Theorem*, 238–260. Cambridge: Cambridge University Press.

Myrvold, W. C. (2018a). Ontology for collapse theories. In S. Gao (Ed.), *Collapse of the Wave Function: Models, Ontology, Origin, and Implications*, 97–123. Cambridge: Cambridge University Press.

Myrvold, W. C. (2018b). Philosophical issues in quantum theory. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2018 ed.). Metaphysics Research Lab, Stanford University.

Myrvold, W. C. (2018c). ψ-ontology result without the Cartesian product assumption. *Physical Review A 97*, 052109.

Myrvold, W. C. (2019). Ontology for relativistic collapse theories. In Lombardi et al. (Eds.), 9–31.

Myrvold, W. C. (2020). On the status of quantum state realism. In S. French and J. Saatsi (Eds.), *Scientific Realism and the Quantum*, 219–251. Oxford: Oxford University Press.

Myrvold, W. C. (2021). Relativistic constraints on interpretations of quantum mechanics. In E. Knox and A. Wilson (Eds.), *The Routledge Companion to Philosophy of Physics*. Routledge.

Ney, A. (2019). Locality and wave function realism. In Lombardi et al. (Eds.), 164–182.

Ney, A. (2021). *The World in the Wave Function: A Metaphysics for Quantum Physics*. Oxford University Press.

Ney, A. and D. Z. Albert (Eds.) (2013). *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*. Oxford University Press.

Pearle, P. (1989). Combining stochastic dynamical state-vector reduction with spontaneous localization. *Physical Review A 39*, 913–923.

Pearle, P. (2015). Relativistic dynamical collapse model. *Physical Review D 91*, 105012.

Pearle, P. and A. Valentini (2006). Generalizations of quantum mechanics. In J.-P. Françoise, G. L. Naber, and T. S. Tsun (Eds.), *Encyclopaedia of Mathematical Physics*, 265–276. Elsevier.

Pusey, M. A., J. Barrett, and T. Rudolph (2012). On the reality of the quantum state. *Nature Physics 8*, 475–478.

Redhead, M. (1987). *Incompleteness, Nonlocality, and Realism*. Oxford: Oxford University Press.

Reichenbach, H. (1956). *The Direction of Time*. Berkeley: University of Los Angeles Press.

Saunders, S., J. Barrett, A. Kent, and D. Wallace (Eds.) (2010). *Many Worlds? Everett, Quantum Theory, and Reality*. Oxford: Oxford University Press.

Shimony, A. (1978). Metaphysical problems in the foundations of quantum mechanics. *International Philosophical Quarterly 18*, 3–17.

Shimony, A. (1984). Controllable and uncontrollable non-locality. In S. Kamefuchi (Ed.), *Foundations of Quantum Mechanics in Light of New Technology*, 225–30. Tokyo: The Physical Society of Japan. Reprinted in Shimony 1993, 130–138.

Shimony, A. (1986). Events and processes in the quantum world. In R. Penrose and C. J. Isham (Eds.), *Quantum Concepts in Space and Time*, 182–203. Oxford University Press. Reprinted in Shimony 1993, 140–162.

Shimony, A. (1990). An exposition of Bell's theorem. In A. I. Miller (Ed.), *Sixty-Two Years of Uncertainty*, 33–43. New York: Plenum. Reprinted in Shimony 1993, 90–103.

Shimony, A. (1993). *Search for a Naturalistic Worldview, Volume II: Natural Science and Metaphysics*. Cambridge University Press.

Skyrms, B. (1984). EPR: Lessons for metaphysics. *Midwest Studies in Philosophy 9*, 245–255.

Tumulka, R. (2006). A relativistic version of the Ghirardi-Rimini-Weber model. *Journal of Statistical Physics 125*, 825–844.

von Neumann, J. (1932). *Mathematische Grundlagen der Quantenmechanik*. Berlin: Verlag von Julius Springer. English translation in von Neumann 1955.

von Neumann, J. (1955). *Mathematical Foundations of Quantum Mechanics*. Princeton: Princeton University Press. Translation, by Robert T. Beyer, of von Neumann 1932.

Wallace, D. (2012). *The Emergent Multiverse*. Oxford University Press.

Wallace, D. and C. G. Timpson (2010). Quantum mechanics on spacetime I: Spacetime state realism. *The British Journal for the Philosophy of Science 61*, 697–727.

Wigner, E. (1961). Remarks on the mind-body question. In I. J. Good (Ed.), *The Scientist Speculates*, 284–302. London: Heinemann.