Relativistic Constraints on Interpretations of Quantum Mechanics

Wayne C. Myrvold
Department of Philosophy
The University of Western Ontario

Contents

1 Introduction

Though it is fairly noncontroversial that an empirically adequate quantum theory must be a quantum field theory, and must be able to treat of the relativistic regime of energies, much of the literature on interpretations of quantum theory has been focussed on nonrelativistic quantum mechanics. There is some justification for this. Nonrelativistic quantum mechanics is conceptually and mathematically simpler, and does, after all, work very well in a low-energy regime. Any solution of the measurement problem, or any account of the ontology of quantum theory, must yield sensible results in this regime, and so nonrelativistic quantum mechanics can serve as a testing-ground for the viability of interpretations.

It should not be assumed, however, that an approach to the interpretation of quantum mechanics will carry over straightforwardly to the context of a relativistic quantum field theory. There are two sorts of challenges to be met. One is adaptation of the approach to a theory with relativistic causal structure. Another challenge arises from the fact that relativistic quantum theories are *quantum field theories*, that is, quantum theories of systems with infinitely many degrees of freedom.

As we shall see below, each of these challenges can pose a serious difficulty for an approach to the interpretation of quantum mechanics. We will focus on four main avenues of approach: (i) additional beables theories, which include "hidden-variables" theories and modal interpretations, (ii) dynamical collapse theories, (iii) Everettian, or "many worlds" interpretations, and (iv) non-realist interpretations, which deny that quantum states represent anything in physical reality independent of considerations of agents and their beliefs.

Summarizing briefly, the conclusions that will be drawn:

(i) Additional beables theories face serious difficulties with relativistic causal structure. This can be dealt with by introducing a dynamically distinguished (albeit unobservable) foliation of spacetime, or else by making the additional beables relational, so that some quantities, such as pointer positions, that we think of as local beables, are not local beables after all. Furthermore, modal interpretations that employ the spectral resolution of a system's quantum state to pick out the additional beables run into serious trouble when extended to quantum field theories, as they run the risk of picking out no non-trivial additional beables.

(ii) Dynamical collapse theories can be adapted, in a fairly straightforward way, to relativistic causal structure. However, a collapse theory that respects the full spacetime symmetry of Minkowski spacetime and its causal structure can do so only at the price of introducing additional, nonstandard quantum degrees of freedom that are unlike those that appear in familiar quantum field theories.

(iii) Everettian interpretations fare better; they extend to the relativistic context with little or no adjustment.

(iv) Some non-realist interpretations bypass difficulties with relativity by limiting the scope of application of quantum mechanics.

2   The Quantum Measurement Problem

If a quantum theory is capable, in principle, of affording a complete description of the physical goings-on in the world (or even just the laboratory), then it must be possible to model an experimental apparatus as a complicated array of quantum systems, and the interactions of the apparatus with the system of interest as a process conforming to the laws of quantum dynamical evolution. This, however, poses a problem, as typical interactions will result in a quantum state that is a superposition of terms corresponding to distinct experimental outcomes. The literature on what is (misleadingly) known as the interpretation of quantum mechanics deals with how to make sense of this.

Among approaches that attempt a realist construal of quantum theories, we can distinguish three main avenues of approach. One avenue accepts the usual, linear evolution of the quantum state, but does not saddle the quantum state with the burden of fully representing the physical world, but, rather, supplements it with additional structure. These are traditionally called "hidden-variables theories," but, as Bell (1981b, 1987a) has pointed out, this is misleading, since, according to a theory like this, it is via these additional variables that the world manifests itself. This is the reason we are calling theories of this sort *additional beables theories*. These include the de Broglie-Bohm pilot wave theory, and also modal interpretations.

Another avenue of attack modifies the dynamics of the quantum state so that, in appropriate circumstances, a process closely approximating the textbook collapse of the state vector occurs. This typically involves a stochastic modification of the dynamics, resulting in an indeterministic theory. The best-known theory of this type is the Ghirardi-Rimini-Weber (GRW) theory, on which collapses are instantaneous, discrete events; there is a continuous analog, known as *continuous spontaneous localization* (CSL).

These two approaches modify the quantum formalism, either by adding additional structure, or by modifying the dynamics. A third class of approaches retains the linear quantum dynamics, with its attendant superpositions of macroscopically distinct states, and attempts to construe a quantum state description as a complete description. These are *Everettian*, or *many-worlds* approaches to quantum theory.

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 Laudisa and Rovelli (2019), and references therein.

On any approach, empirical adequacy requires recovery of at least an approximation to the quantum probabilities for outcomes of experiments, insofar as these have been verified by experiment. This includes violations of Bell inequalities. It might seem that this, by itself, leads to an unavoidable conflict with relativistic causality. As we shall see, things are not so simple!

3   Relativistic Spacetimes

Compatibility with special relativity will be a concern in this chapter, and this amounts to the question of whether a theory can live happily in Minkowski spacetime, without adjoining any additional spacetime structure, such as a privileged relation of distant simultaneity. But, of course, special relativity is false, and the spacetime in which we live and move and have our being is not Minkowski spacetime, but, as general relativity teaches us, a spacetime of variable and dynamic curvature. Minkowski spacetime is of interest only as a local approximation, on scales at which curvature can be ignored. And general relativity itself is, almost certainly, an approximation to a deeper theory that incorporates quantum gravity.

For this reason, we will be interested, not only in Minkowski spacetime, but in a wider class of relativistic spacetimes. The radical departure of special relativity from previous conceptions of spacetime lies in its rejection of a privileged global temporal order. In Galilean spacetime, for any event p, any other event is either in the past of p, the future of p, or is simultaneous with p, and the relation of simultaneity is, as one would expect, an equivalence relation. In Minkowski spacetime, events that are spacelike separated are neither to the past nor to the future of each other, and, although there will always be some inertial reference frame whose accompanying Lorentz coordinates assign the two events the same t-coordinate, this is not simultaneity in any physically significant sense.

In Minkowski spacetime, define a relation $\prec$ that holds between spacetime points p, q, when there is a future-directed curve from p to q that is everywhere timelike or lightlike. This relation is irreflexive (no point stands in that relation to itself), antisymmetric (if $p \prec q$ then it is not the case that $q \prec p$), and transitive.

Now define the relation $\sim$ of temporal unconnectedness, that holds between two points p and q if neither $p \prec q$ nor $q \prec p$. Distinct points that are temporally unconnected are said to be *spacelike separated*. This relation is reflexive and symmetric, but not transitive. On the contrary, for any two points p, q such that $p \sim q$, one can always find a third, r, such that $p \sim r$ but $r \prec q$.

All of this remains true in the curved spacetimes of general relativity, provided that they are sufficiently well-behaved.[1] This is a striking difference between relativistic spacetimes and Galilean spacetime. In Galilean spacetime we have a relation $\prec$ of temporal precedence that is irreflexive and antisymmetric, but the corresponding relation of temporal unconnectedness is transitive, and hence spacetime can be uniquely be partitioned into equivalence classes under this simultaneity relation.

In this chapter, we will be concerned with spacetimes equipped with a relation $\prec$ of temporal precedence, from which we define a relation $\sim$ of temporal unconnectedness. We will assume that $\prec$ is irreflexive and antisymmetric (no temporal loops). We will say that a spacetime's temporal structure is *causally relativistic* iff for any two points p, q such that $p \sim q$, there exists a point r, such that $p \sim r$ and $r \prec q$. Note that this is a condition on causal structure and not a condition of symmetry under translations or boosts (which, of course, are additional conditions that one might impose).

---

[1]There is no need here for a precise characterization of "well-behaved," but global hyperbolicity suffices for our purposes.

On the usual notion of causation, a cause must temporally precede its effects, and so, if an event p is a cause of an event q, we must have $p \prec q$. It is this that motivates the claim that special relativity forbids action at a distance; without it, there is no reason to think that causal relations between spacelike separated events are incompatible with relativistic spacetime structure. Given a causally relativistic spacetime, we will say that a theory is *causally local* if all causal relations that the theory affords respect the condition that the cause be in the past of the effect and, hence, that there are no cause-effect relations between temporally unconnected events. There seems to be no harm in assuming that any spacetime point to the past of a point q is potentially a locus of an event that is a cause of q. If we make this assumption, then the relations of temporal precedence and potential causal influence coincide, and we may, as is usual in discussions of spacetime structure, treat temporal structure and causal structure interchangeably.

The condition of causal locality—that is, the condition that there be no cause-effect relations between spacelike separated points—is independent of symmetries such as boost symmetries. One can invent theories that are causally local but not invariant under boosts, and one can invent theories that violate the condition of causal locality without picking out a preferred reference frame.

There is a tradition of invoking retrocausality—that is, backwards-in-time causation—in explaining quantum phenomena. This approach has its roots in a suggestion by O. Costa de Beauregard (1976) and has been championed by a number of authors; see Friederich and Evans (2019) for an overview. The idea is to make causal connectability time-symmetric by permitting cause-effect relations in both temporal directions. On this approach, causal influences can propagate along time-like or light-like lines in either temporal direction. Assuming that causal connectability is transitive, this turns the relation of causal connectability into the trivial relation that holds between any pair of spacetime points. Any approach of this sort is, therefore, straightforwardly causally nonlocal, in the sense we are using the phrase, as any two spacetime points, including those that outside of each other's lightcones, are causally connected. As Kastner (2013, Ch. 6) has argued, an approach of this sort is capable of respecting the symmetries of Minkowski spacetime.

4   Quantum theories in relativistic spacetimes

A quantum-mechanical theory is what results when one subjects a classical theory of a system with finitely many degrees of freedom (such as, for example, a finite number of particles) to quantization. A quantum field theory, on the other hand, results from quantization of a classical theory with infinitely many degrees of freedom. Introductory quantum field theory textbooks typically contain, in their early chapters, arguments that a relativistic quantum theory must be a quantum field theory. For a precise version of one such argument, see Malament (1996).

There are, therefore, two sorts of difficulties one might encounter when one seeks to extend to the relativistic domain an approach to the quantum measurement problem that has been formulated in terms of nonrelativistic quantum mechanics. One source of difficulties might lie in a adapting the theory to relativistic causal structure. Another source of difficulties might stem from differences between quantum mechanics and quantum field theories.

Typically, relativistic quantum field theories are expressed in the Heisenberg picture, rather than the Schrödinger picture. The Schrödinger picture, or an analogue of it, can, however, be formulated in a relativistic spacetime. On this picture, we track the rate of change of an observable via a changing state vector. Take a foliation of spacetime into spacelike hypersurfaces (most conveniently, hyperplanes, but we can also consider nonflat surfaces), taken to be smooth, so that there is, at any point of each of these hypersurfaces, a vector normal to the surface. Integral curves of this vector field form a timelike congruence. Pick one hypersurface σ_0, and consider a field operator $\hat{\phi}(x)$. For each x_0 on σ_0, there will be a family of operators $\hat{\phi}(x)$ corresponding to points on the integral curve containing x_0, and, for any quantum state, one can consider how the expectation value of $\hat{\phi}(x)$ changes as one follows these timelike curves. This can be expressed, either in the Heisenberg picture, using the same state vector and different operators $\hat{\phi}(x)$ for different points on an integral curve, or in the Schrödinger picture, using the same operator $\hat{\phi}(x_0)$ for every spacetime point on the integral curve containing x_0, and different state vectors for evaluating expectation values of observables on different elements of our foliation.

There is also an analog of the interaction picture, called the *Tomonaga-Schwinger picture*. On this picture, one writes the Lagrangian density as a sum of a free-field Lagrangian and a term incorporating interactions:

$$\mathcal{L}(x) = \mathcal{L}_0(x) + \mathcal{L}_I(x). \tag{1}$$

The operators employed are solutions of the free-field equations, and one associates, with each spacelike Cauchy surface σ (whether flat or not), a state vector $|\psi(\sigma)\rangle$. Evolution from a surface σ to another, σ', differing by an infinitesimal deformation about a point x, satisfies the *Tomonaga-Schwinger equation* (Tomonaga, 1946; Schwinger, 1948):

$$i\hbar c\,\frac{\delta|\psi(\sigma)\rangle}{\delta\sigma(x)} = \mathcal{H}_I(x)|\psi(\sigma)\rangle. \tag{2}$$

Integration of this equation yields, for any Cauchy surfaces σ, σ', a unitary mapping from $|\psi(\sigma)\rangle$ to $|\psi(\sigma')\rangle$. This approach is not commonly discussed in recent textbooks but can be found in some older texts; see, *e.g.*, Schweber (1961).

5 Additional-Beables Theories

5.1 De Broglie-Bohm theory

The best-known additional-beables theory is the de Broglie-Bohm pilot wave theory. On this theory, ordinary physical objects consist of point particles, or corpuscles, whose motion is guided by the quantum wave function. The dynamical laws of theory are the Schrödinger equation for the evolution of the quantum wave function, and the guidance equation, which fixes the velocities of Bohmian corpuscles.

$$m_i \mathbf{v}_i(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n, t) = \boldsymbol{\nabla}_i S(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n, t), \tag{3}$$

where $S/\hbar$ is the phase of the wave function:

$$S(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n, t) = \hbar \,\text{Im}\, \log \Psi(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n, t). \tag{4}$$

It is assumed that, at some time t_0, the probability density for configurations of the corpuscles is given by the square of the wave function amplitude. If this holds for some time t_0, the dynamics, consisting of the Schrödinger equation and the guidance equation, ensures that it holds for all time. This property is known as *equivariance* of the Born-rule distribution.

5.1.1 Causal structure

It can be seen from the guidance equation (3) that the velocity of any one of the corpuscles depends on an instantaneous configuration, and may, in principle, depend on the positions of arbitrarily many corpuscles, arbitrarily far away. Thus, the guidance equation requires a distinguished relation of distant simultaneity for its formulation.

As an illustration, consider a pair of spin-1/2 particles in the singlet state, spatially separated and sequentially subjected to Stern-Gerlach experiments for the same direction, say the z-direction, which we will regard as the vertical direction. The quantum state can be written,

$$|\Psi(t)\rangle = \psi_{+-}(\mathbf{x}_1, \mathbf{x}_2, t) |z^+\rangle_1 |z^-\rangle_2 + \psi_{-+}(\mathbf{x}_1, \mathbf{x}_2, t) |z^-\rangle_1 |z^+\rangle_2 \tag{5}$$

Before the experiment, there is no correlation between position and spin-state, and the two wavepackets overlap completely.

$$\psi_{+-}(\mathbf{x}_1, \mathbf{x}_2, t_0) = \psi_{-+}(\mathbf{x}_1, \mathbf{x}_2, t_0). \tag{6}$$

The outcome of both experiments is determined by the position of the particle on which the experiment is done first. The Stern-Gerlach experiment on particle 1 separates the wavepackets ψ_{+-} and ψ_{-+} along the z_1-direction. Since Bohmian trajectories cannot cross, there will be a plane or other two-dimensional surface in $\mathbf{x}_1$-space such that particle 1 will go in the $+$ direction if it is initially located above this plane, in the $-$ direction, if it is located below. After this experiment, there is no longer any overlap, in configuration space, between the two wavepackets, and the two particles are determinately located in only one of them. The second particle responds to its Stern-Gerlach apparatus accordingly, responding $+$ or $-$ depending on which wavepacket it is in.

This means that, if the experiment on particle 1 is done first, the outcome of both experiments is determined by the initial position of particle 1. By parity of reasoning, if the experiment on particle 2 is done first, the outcome of both experiments is determined by the initial position of particle 2. These outcomes could differ, meaning that, for the same initial wavefunction and initial configuration of the particles, the outcomes of the experiments could differ depending on which experiment is done first. Thus, the theory requires there to be a matter of fact about which experiment is done first, even if the experiments are done at spacelike separation.

We might ask: is this an artefact of the way this theory is formulated, or is it inherent in the very project of formulating a theory of this type, on which particles have definite trajectories and probabilities over configurations are to be given by the Born rule, with dynamics for the particles that ensures that this probability distribution is equivariant?

As Berndl et al. (1996) have shown, the answer to this question is that a theory of this type must, indeed, invoke a privileged foliation, as it is impossible for the distribution postulate to be satisfied on all spacelike hyperplanes.

Suppose that we have a pair of spin-1/2 particles, initially both located in the same small region. The particles are then separated a large distance, and each is passed through a Stern-Gerlach apparatus oriented in the x-direction, which separates the wavepacket (in configuration space) along the x-axis of the particle it's interacting with. The beams are then carefully recombined, with care taken not to let the particles interact with anything that might induce decoherence, and then are passed through Stern-Gerlach devices oriented in the z-direction. We suppose that all of this splitting, recombining, and resplitting of the two beams occurs at spacelike separation.

We consider four hypersurfaces (see Figure 1):

1. A hypersurface α, on which both particle 1 and particle 2 have their wavepackets split into $x+$ and $x-$ wavepackets.

2. A hypersurface β, on which both particle 1 and particle 2 have their wavepackets split into $z+$ and $z-$ wavepackets.

3. A hypersurface γ, on which the wavepackets are split along z_1 and x_2 axes, respectively.

4. A hypersurface δ, on which the wavepackets are split along x_1 and z_2 axes, respectively.
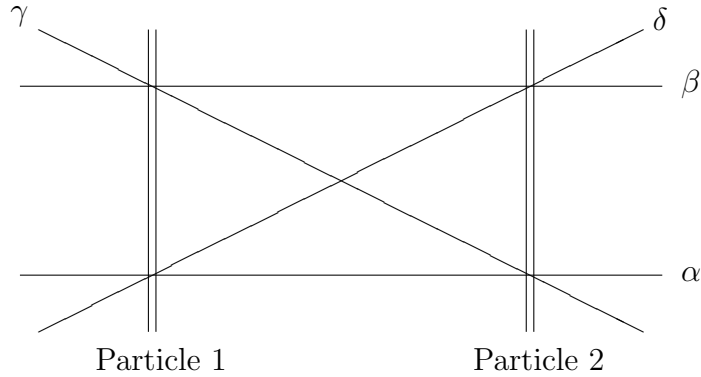

Figure 1: The hypersurfaces considered in the proof.

It is not, in fact, always possible to satisfy the distribution postulate along all four of these hypersurfaces. Suppose that our pair of spin-1/2 particles is prepared in the Hardy state:

$$|\psi\rangle_H = \frac{1}{\sqrt{12}} \left(|x^+\rangle_1|x^+\rangle_2 - |x^+\rangle_1|x^-\rangle_2 - |x^-\rangle_1|x^+\rangle_2 - 3|x^-\rangle_1|x^-\rangle_2 \right) \tag{7}$$

$$= \frac{1}{\sqrt{3}} \left(|z^+\rangle_1|z^-\rangle_2 - |z^+\rangle_1|z^+\rangle_2 + |z^-\rangle_1|z^+\rangle_2 \right) \tag{8}$$

$$= \frac{1}{\sqrt{3}} \left(|z^-\rangle_1|z^+\rangle_2 - \sqrt{2}\,|z^+\rangle_1|x^-\rangle_2 \right) \tag{9}$$

$$= \frac{1}{\sqrt{3}} \left(|z^+\rangle_1|z^-\rangle_2 - \sqrt{2}\,|x^-\rangle_1|z^+\rangle_2 \right) \tag{10}$$

Let us write the states on these four hypersurfaces, with the position degrees of freedom included, writing $|+\rangle$ for the particle being in the $+$ beam, $|-\rangle$ for its being in the $-$ beam.

$$|\psi(\alpha)\rangle = \frac{1}{\sqrt{12}}\left(|x^+\rangle_1|+\rangle_1|x^+\rangle_2|+\rangle_2 - |x^+\rangle_1|+\rangle_1|x^-\rangle_2|-\rangle_2 \right.$$
$$\left. - |x^-\rangle_1|-\rangle_1|x^+\rangle_2|+\rangle_1 - 3|x^-\rangle_1|-\rangle_1|x^-\rangle_2|-\rangle_2\right) \quad (11)$$

$$|\psi(\beta)\rangle = \frac{1}{\sqrt{3}}\left(|z^+\rangle_1|+\rangle_1|z^-\rangle_2|-\rangle_1 - |z^+\rangle_1|+\rangle_1|z^+\rangle_2|+\rangle_2 + |z^-\rangle_1|-\rangle_1|z^+\rangle_2|+\rangle_1\right) \quad (12)$$

$$|\psi(\gamma)\rangle = \frac{1}{\sqrt{3}}\left(|z^-\rangle_1|-\rangle_1|z^+\rangle_2|+\rangle_2 - \sqrt{2}\,|z^+\rangle_1|+\rangle_1|x^-\rangle_2|-\rangle_2\right) \quad (13)$$

$$|\psi(\delta)\rangle = \frac{1}{\sqrt{3}}\left(|z^+\rangle_1|+\rangle_1|z^-\rangle_2|-\rangle_2 - \sqrt{2}\,|x^-\rangle_1|-\rangle_1|z^+\rangle_2|+\rangle_2\right) \quad (14)$$

If the distribution postulate is satisfied on α, there is a probability $1/12$ that both particles are in their respective x^+ beams. If the postulate is satisfied on γ, then in such cases particle 1 must be in its z^- beam on γ. If the postulate is satisfied on δ, then particle 2 must be in its z^- beam on δ. But this means that, on β, both particles will be in the z^- beams, which should have probability zero. Conclusion: the distribution postulate cannot be satisfied along all foliations, and the dynamics for the particles must be choosy about which, if any, foliation, it is with respect to which equivariance is to be maintained.

The argument does not depend on preparing the Hardy state exactly, or that our theory yield quantum probabilities exactly. On a theory of this sort, there must be a joint probability distribution over the positions of the particles at each point of intersection of the hyperplans considered that yields the correct correlations on each of the four hyperplanes. As demonstrated by Fine (1982a,b), violation of a Bell-type inequality for these variables is equivalent to impossibility of a joint distribution of that sort.

For example: if there is a joint distribution over these variables, we will have to have,

$$Pr(x_1^+(\alpha), x_2^+(\alpha)) \leq Pr(x_1^+(\delta), z_2^+(\delta)) + Pr(z_1^+(\gamma), x_2^+(\gamma)) + Pr(z_1^-(\beta), z_2^-(\beta)), \quad (15)$$

which is violated by the Hardy state.

A trajectory theory must choose a foliation along which the distribution postulate is to be satisfied. Moreover, it is fairly easy to see that, in order to maintain the distribution along this foliation, the dynamics must permit causal dependence of the behaviour of particles on choice of experiment at spacelike separation.

Consider three scenarios.

A. For particle 1, operations proceed as above, recombining the beams and resplitting them into z^+ and z^- wavepackets. Nothing is done to particle 2 between α and β, and there is a split into x^+ and x^- wavepackets for particle 2 on both α and β.

B. Nothing is done to particle 1 between α and β, and it remains split into x^+ and x^- wavepackets on both α and β. Recombining and resplitting proceed as above, for Particle 2.

C. Recombining and resplitting proceed as above, for both particles.

Suppose that we have dynamics according to which, if the probabilities for positions of particles match the Born-rule probabilities on α, they do so on β, for all three scenarios. Suppose, also, in scenarios A and B, the particle to which nothing is done does not hop from one of its wavepackets to another. That is, in scenario A, if particle 2 is in the x^+ beam on α, it is in the x^+ beam on β.

These conditions can be satisfied only if, for some initial configurations on α, transition probabilities—that is probabilities of a particle ending up in a certain position on β, given the configuration on α—for one particle depend on what is being done to the other particle.

To see this, suppose the contrary. Suppose that, for any configuration on α, the probability that particle 1 ends up in the z^+ beam on β is the same, whether scenario A or C is in effect, and that the probability that particle 2 ends up in the z^+ beam on β is the same, whether scenario B or C is in effect. Then, since, on Scenario A, particle 2 is in the same beam on α and β, if the Born rule is satisfied on β, it is also satisfied on γ. If transition probabilities for particle 1 are the same in scenarios A and C, it follows that, in scenario C, the Born rule is satisfied on γ.

Running the same argument with the roles of particles 1 and 2 reversed yields the conclusion that, if in scenario B the Born rule is satisfied on β, and if transition probabilities for particle 2 are the same in scenarios B and C, it follows that, in scenario C, the Born rule is satisfied on δ. But, as we have seen, in scenario C, the Born rule cannot be satisfied on all four hypersurfaces.

5.1.2 Bohmian quantum field theories

Even though there are good reasons to think that an additional-variables theory must employ a preferred foliation, the question still arises whether there could be a theory of this sort that recovers the empirical content of relativistic quantum field theories. This would, of course, entail that the preferred foliation is empirically inaccessible. Two avenues of approach suggest themselves for extending Bohmian theories to the context of quantum field theories. One is to retain a basic particle ontology, and permit particle creation and annihilation. The other is to take classical field configurations as the basic ontology, and to provide a non-classical dynamics for the evolution of these field configurations. For an overview of these options, see Struyve (2011). Struyve concludes that the field approach works best for bosonic fields, but is problematic when one attempts to extend it to fermionic fields. For fermionic fields, the particle approach seems more appropriate.

The task of extending additional-beables theories to the context of quantum field theories remains a work in progress. However, there do not seem to be in-principle obstacles to the creation of a theory of this sort that recovers the phenomena that are currently handled by quantum field theories.

5.2 Modal interpretations

Modal interpretations supplement the quantum state with additional beables. A variety of interpretations along these lines have been proposed; see Lombardi and Dieks (2017) for an overview. Many of these share the feature that the nature of the additional beables depends on the quantum state. Given a quantum system α, which may be a subsystem of a larger

system, one considers the reduced state of α, obtained by tracing out the degrees of freedom of the rest of the system. If the state of the larger system is a pure quantum state, the reduced state of α will be pure if α is not entangled with the rest of the system, mixed, if it is. One resolves the density operator representing the reduced state into its spectral components:

$$\rho_\alpha = \sum_i w_i \, P_i, \tag{16}$$

where $\{P_i\}$ are mutually orthogonally projections. In the non-degenerate case, in which all nonzero coefficients $\{w_i\}$ are distinct, this resolution is unique. The *basic modal rule* assigns definite values 1 or 0 to these projections; one of these projections represents a property definitely possessed by the system, the others, definitely absent. The probability that P_i is the possessed property is w_i.

5.2.1 Difficulties arising from relativistic causal structure

Can modal interpretations respect relativistic causal structure, or do they require a distinguished relation of distant simultaneity? Dickson and Clifton (1998) proved, for a broad class of modal interpretations, that is impossible to furnish the theory with transition probabilities in such a way as to yield Born-rule probabilities for possessed values on arbitrary hyperplanes. Dickson and Clifton's proof relies on an assumption concerning the transition probabilities for possessed values, the assumption they call "stability," but, as Arnztenius (1998) pointed out, the stability requirement is dispensable and the core of the proof concerns the nonexistence of certain joint distributions yielding the appropriate Born probabilities as marginals. Taking these proofs as a starting point, Myrvold (2002, 2009) adapted the argument of Berndl et al. (1996), outlined above, to show that any modal interpretation, so long as the possessed values of a system are capable of being regarded as local beables, cannot have Born-rule probabilities for these possessed values satisfied on arbitrary spacelike hyperplanes. The argument at the end section 5.1.1, above, can be straightforwardly adapted to show that these theories, also, must permit causal dependence of the behaviour of the additional beables on choice of experiment at spacelike separation.

Undeterred, Berkovitz and Hemmo(2005a,2005b,2006) have proposed a relational modal interpretation. The interpretation assigns properties to systems that are not thought to be intrinsic to that system, but are possessed by it only relative to some other system. In a relativistic context, these relational values may depend on what hypersurface is used to define a global state of the universe, and it can happen, in the presence of entanglement at spacelike separation, that some property of a system, at a single position on its worldline, may have different values relative to different hypersurfaces. Thus, the possessed values are not to be thought of as local beables, and, as the authors argue in some detail, the no-go theorems don't apply.

5.2.2 Difficulties arising from infinitely many degrees of freedom

In addition to the difficulties for modal interpretations arising from relativistic causal structure, there is another difficulty for attempts to extend modal interpretation to the context of relativistic quantum field theories, arising from infinitude of degrees of freedom.

There is, first, the issue of extending the basic modal rule to the context of a relativistic quantum field theory. We want to be able to apply the rule to local subsystems of the universe. If we consider a bounded region R of spacetime, and the algebra of operators $\mathcal{A}(R)$ formed from fields in R, there is a difficulty in applying the basic modal rule, since, as pointed out by Dieks (2000), these local algebras will not contain either a density operator representing the state, or projections onto finite-dimensional subspaces. Clifton (2000) provided a version of the modal rule that makes sense in the context of algebraic quantum field theory, and also showed that there exists a whole host of states for which the rule picks out no definite observables other than multiples of the identity (that is, observables that have the same value no matter what the state is). Earman and Ruetsche (2005) extend Clifton's result, and discuss the significance of these results for modal interpretations of quantum theory.

These trivialization results mean that a modal interpretation of this sort cannot respect relativistic causal structure and symmetries while accomplishing the goal of providing a solution to the measurement problem. The purpose of a modal interpretation is to have observables such as apparatus pointer observables take on definite values in physically realistic states. Obviously, this is not accomplished if no observables are definite other than multiples of the identity.

6 Collapse Theories

A dynamical collapse theory modifies the evolution of the quantum state. One approach that might be tried would be to replace the deterministic, unitary evolution of the quantum state by a deterministic and non-unitary evolution. There are arguments to the effect that, given plausible assumptions, this would lead to the possibility of superluminal signalling (Gisin, 1989; Simon et al., 2001).[2] For this reason, in this section we focus on stochastic collapse theories.

On a theory such as this, the dynamical laws do not specify a unique future state, given the state at a time t_0; rather, which state will be realized is, on such a theory, a matter of chance. The idea is to specify a probabilistic law that approximates the usual unitary, deterministic evolution within those domains where we have good evidence for it, but departs sufficiently from it in some situations—which should include experiments that lead to recorded outcomes—so as to closely approximate the textbook collapse postulate, according to which some definite outcome obtains at the end of an experiment, with probabilities as to which outcome is obtained given by the Born rule. One theory of this sort, with unitary evolution punctuated by discrete jumps, is the Ghirardi-Rimini-Weber (GRW) theory, referred to by the authors as *Quantum Mechanics with Spontaneous Localization* (QMSL) (Ghirardi et al., 1986). Another is the *Continuous Spontaneous Localization* (CSL) theory (Pearle 1989, Ghirardi, Pearle, and Rimini 1990).

Among the experiments to be accounted for are tests of Bell inequalities. Since our theory is required to produce correct probabilities for the outcomes of such experiments, we do not expect the probabilistic law to be one on which spacelike separated events are

---

[2]These assumptions may be questioned; see Kent 2005b.

probabilistically independent. It might seem obvious, at first, that this dooms compatibility with relativity from the start.

Things are not so simple, however. Bell himself said that, when he first encountered the GRW theory, "I thought I could blow it out of the water, by showing it was *grossly* in violation of Lorentz invariance" (Bell, 1989, p. 17). A careful examination led him to the opposite conclusion, namely, that, though couched in terms of nonrelativistic quantum mechanics, the theory exhibits "a residue, or at least an analogue, of Lorentz invariance" (Bell 1987a, p. 10 in Bell 1987b, pp. 206–207). The idea is this. On the usual, Schrödinger, dynamics, for two (or more) systems, widely separated from each other and not interacting with each other, evolution of each system is independent, and we can employ separate time variables for each system. The Schrödinger dynamics do not require any notion of synchronization of the time variables. This is a feature that Bell calls *relative time translation invariance.* The question he addressed is: for systems like that, does the GRW theory retain relative time translation invariance, or do its probabilistic laws require there to be a matter of fact about the temporal order of events at a distance? Bell's verdict: the probability of any sequence of GRW jumps is relative time translation invariant.

> I am particularly struck by the fact that the model is as Lorentz invariant as it could be in the nonrelativistic version. It takes away the ground of my fear that any exact formulation of quantum mechanics must conflict with fundamental Lorentz invariance (Bell 1987a, p. 14, in Bell 1987b, p. 209).

In his Trieste lecture, delivered in November 1989, Bell discussed the prospects for a genuinely relativistic version of a dynamical collapse theory, and concluded that the difficulties encountered by Ghirardi, Grassi, and Pearle in producing a genuinely relativistic version of CSL a theory that would be "Lorentz invariant, not just for all practical purposes but deeply, in the sense of Einstein, eliminating entirely any privileged reference system from the theory" (Bell, 2007, p. 2931), were "Second-Class Difficulties," technical difficulties, and not deep conceptual ones (see also Bell 1989, p. 25).

A relativistic collapse theory, if it is to yield even an approximation to the usual quantum probabilities, will include correlations between events at spacelike separation that violate Bell inequalities. The Bell inequalities can be derived from the condition that correlations 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 1981a, p. C2–55, in Bell 1987b, p. 152). This condition is violated in relativistic collapse theories, which involve probabilistic correlations between spacelike separated events that are *not* attributable in the usual way to events in their common past. Unlike cause-and-effect relations as usually conceived, the relation of probabilistic correlation between these distant events is symmetric, and does not require a temporal order between the events. For this reason, it is misleading to assimilate this relation to the causal relation and refer to it as nonlocal "influence" or "action at a distance." See Myrvold (2016) for further discussion.

As we shall see, there are difficulties associated with constructing a fully relativistic collapse theory. These difficulties arise, not because a collapse theory cannot be formulated in the absence of a distinguished relation of distant simultaneity, but rather from the problem

of constructing a relativistic collapse theory with a stable vacuum state. In order to see how all this works, we start with a few words about collapse theories in general.

6.1 Stochastic quantum state evolution

It is generally accepted that any physically realizable change of the state of a quantum system must be a *completely positive* mapping of the system's state space into itself. Any such mapping can be represented by a set $\{K_i\}$ of operators, such that the density operator ρ representing the state undergoes the change

$$\rho \to \sum_i K_i \, \rho \, K_i^\dagger, \tag{17}$$

where

$$\sum_i K_i^\dagger K_i \leq \mathbb{1}. \tag{18}$$

If $\sum_i K_i^\dagger K_i = \mathbb{1}$, the operation is called a *nonselective operation*; if $\sum_i K_i^\dagger K_i < \mathbb{1}$, it is *selective.* Such a representation of a completely positive mapping of the state space into itself is called a *Kraus representation* of the mapping; and the operators $\{K_i\}$, *Kraus operators.*

Selective operations reduce the trace-norm of the density operator, but this is not an issue, as normalization is only a convention. With an unnormalized density operator, we compute the expectation value of an observable represented by an operator A via

$$\langle A \rangle_\rho = \mathrm{Tr}(\rho A)/\mathrm{Tr}(\rho). \tag{19}$$

In the simplest case, the set of Kraus operators is a singleton, and the evolution is deterministic; this includes the case of unitary evolution. However, one can also consider stochastic processes. Suppose that we has a set of operators $\{K_i\}$, such that the state vector undergoes a stochastic transition of the following form. For some i,

$$|\psi\rangle \Rightarrow |\psi'\rangle = K_i|\psi\rangle/\|K_i|\psi\rangle\|, \tag{20}$$

with the probability for which transition it undergoes given by

$$p_i = \|K_i|\psi\rangle\|^2. \tag{21}$$

Since these probabilities must sum to unity for every vector $|\psi\rangle$, we must have $\sum_i K_i^\dagger K_i = \mathbb{1}$. For each i, we have a selective operation. One of these yields the actual state. The transition to the mixture of these candidate states, which corresponds to the proposition that some one of these transitions has occurred, without specification of which, is given by weighted mean of these.

$$\rho \to \bar{\rho}' = \sum_i p_i \frac{K_i \rho K_i^\dagger}{\mathrm{Tr}(K_i \rho K_i^\dagger)}. \tag{22}$$

In (20), we indexed the possible state transitions via a discrete index i. We can also consider stochastic processes of a more general sort. Let $\langle \Gamma, \mathcal{F}, \mu \rangle$ be a probability space.[3]

---

[3]That is, Γ is some set, $\mathcal{F}$ is a σ-algebra of subsets of Γ containing both the empty set $\emptyset$ and Γ itself, to be regarded as the measurable subsets of Γ, and μ is a non-negative countably additive set function on $\mathcal{F}$ with $\mu(\Gamma) = 1$.

Suppose that we have a family of operators $\{K_\gamma\}$, for $\gamma \in \Gamma$, such that

$$\int_\Gamma K_\gamma^\dagger K_\gamma \, d\mu(\gamma) = \mathbb{1}. \tag{23}$$

These can serve as the operators that induce our state transitions.

This gives us a rather general schema for a stochastic process in a Hilbert space. For the moment, we will presume that we have a unique global time function; extension to a relativistic spacetime will be considered in the next subsection.

Any theory satisfying the following conditions will give us stochastic evolution of the state vector.

1. We assume a measure space $\langle \Gamma, \mathcal{F}, \mu \rangle$. The elements of Γ are to be thought of as possible complete histories, specifying events for all times. For each time t there is a σ-algebra $\mathcal{F}_t$, whose elements are to be thought of as sets of possible histories up to time t. For $t < s$, we will have $\mathcal{F}_t \subseteq \mathcal{F}_s \subseteq \mathcal{F}$.

2. Given a history up until time t, and a state vector $|\psi(t)\rangle$, the state at a later time $t+s$ is a random variable, defined as follows.

    (a) We assume a family of operators $K_\gamma(s;t)$, which depend only on events up until time s (that, is, if two histories γ, γ' agree up to time s, then $K_\gamma(s;t) = K_{\gamma'}(s;t)$), such that

    $$\int_\Gamma K_\gamma^\dagger(s;t) K_\gamma(s;t) \, d\mu = \mathbb{1}.$$

    (b) For some γ,

    $$|\psi(t+s)\rangle = \frac{K_\gamma(s;t)|\psi(t)\rangle}{\|K_\gamma(s;t)|\psi(t)\rangle\|}.$$

    (c) The probability distribution for the realized outcome is given by

    $$\mathrm{Prob}(\gamma \in F) = \int_F \|K_\gamma(s;t)|()\psi(t)\|^2 \, d\mu.$$

    for any F in $\mathcal{F}_s$.

On a theory of this sort, given a state $|\psi(t_0)\rangle$, for any later time $t_1 = t_0 + s$ there will be an actual state vector $|\psi(t_1)\rangle$, and its corresponding pure-state density operator

$$\rho(t_1) = |\psi(t_1)\rangle\langle\psi(t_1)|. \tag{24}$$

There will also be what may be called an *ensemble*, or *mean* density operator, which is a weighted average of the various possibilities for $\rho(t_1)$, weighted by their respective probabilities.

$$\bar{\rho}(t_1;t_0) = \int_\Gamma K_\gamma(t_1;t_0) \, \rho(t_0) \, K_\gamma^\dagger(t_1;t_0) \, d\mu. \tag{25}$$

6.2 Collapse theories in Minkowski spacetime

6.2.1 Causality conditions

It is useful to work within what may be called the *stochastic Tomonaga-Schwinger picture*. We start with field operators that are solutions of the standard Heisenberg-picture equations (these may be free or interacting fields). We associate with each spacelike Cauchy surface σ a state vector $|\psi(\sigma)\rangle$. Our theory should deliver a stochastic evolution of $|\psi(\sigma)\rangle$ to $|\psi(\sigma')\rangle$ whenever σ' is nowhere to the past of σ. This will go much as it did in §6.1.

1. We assume a measure space $\langle \Gamma, \mathcal{F}, \mu \rangle$. The elements of Γ are to be thought of as possible complete histories, specifying events for all times. For each spacelike Cauchy surface σ there is a σ-algebra $\mathcal{F}_\sigma$, whose elements are to be thought of as sets of histories to the past of σ. For two Cauchy surfaces σ, τ, with $\tau \preceq \sigma$, we will have $\mathcal{F}_\tau \subseteq \mathcal{F}_\sigma \subseteq \mathcal{F}$.

2. Given two Cauchy surfaces σ, τ, with $\tau \preceq \sigma$, let δ be the spacetime region between them. Given a state vector $|\psi(\tau)\rangle$, the state vector $|\psi(\sigma)\rangle$ is a random variable, defined as follows.

   (a) We assume a family of operators $\{K_\gamma(\delta)\}$, which depend only on events in the past of σ (that, is, if two histories γ, γ' agree on events in the past of σ then $K_\gamma(\delta) = K_{\gamma'}(\delta)$), such that

   $$\int_\Gamma K_\gamma^\dagger(\delta) K_\gamma(\delta) \, d\mu = \mathbb{1}.$$

   We call these *evolution operators* for the region δ.

   (b) For some γ,
   $$|\psi(\sigma)\rangle = \frac{K_\gamma(\delta)|\psi(\tau)\rangle}{\|K_\gamma(\delta)|\psi(\tau)\rangle\|}.$$

   (c) The probability distribution for the realized outcome is given by

   $$\mathrm{Prob}(\gamma \in F) = \int_F \|K_\gamma(\delta)|\psi(\tau)\rangle\|^2 \, d\mu.$$

   for any F in $\mathcal{F}_\sigma$.

Here, again, we can define a mean density operator for the state on σ, given a state vector on τ, with $\tau \subseteq \sigma$.

$$\bar{\rho}(\sigma;\tau) = \int_\Gamma K_\gamma(\delta)\rho(\tau)K_\gamma^\dagger(\delta) \, d\mu. \tag{26}$$

Consider two Cauchy surfaces, τ, σ, which coincide everywhere except on the boundaries of two bounded regions δ_1, δ_2, where σ is to the future of τ (see Figure 2). We must have a unique law of stochastic evolution from τ to σ, through $\delta_1 \cup \delta_2$, and this should coincide with the composition of the evolution through δ_1 and the evolution through δ_2, in either order. The necessary and sufficient condition for this is,
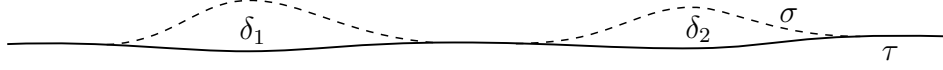
Figure 2: Spacelike separated regions between two spacelike hypersurfaces.

  3. Evolution operators corresponding to spacelike separated regions commute.

  One difference between a theories of this sort and a theory involving deterministic, unitary evolution, is that, for calculating probabilities of outcomes of experiments performed in some spacelike slice α that is common to two Cauchy surfaces τ, σ, it can make a difference whether the calculation is made using $|\psi(\tau)\rangle$ or $|\psi(\sigma)\rangle$. It is this feature that has given rise to concerns that a relativistic collapse theory must either be incoherent or exhibit radical contextualism about possessed properties. These concerns are misplaced; more on this, below, in subsection 6.2.2. However, it *would* be problematic if probabilities for events in α differed depending on whether one used $|\psi(\tau)\rangle$ or $\bar{\rho}(\sigma)$, which is the probabilistically weighted mixture of various possibilities for $|\psi(\sigma)\rangle$, given the state on τ. Moreover, if the evolution from τ to σ, through a region δ, spacelike separated from α, depended on a choice of experiments to be made, then, if different choices yielded different mixtures $\bar{\rho}(\sigma)$, the choice of experiment could be used for superluminal signalling. We, therefore, require $\bar{\rho}(\sigma)$ to yield the same probabilities for outcomes of experiments as $|\psi(\tau)\rangle$, for any allowable evolution through δ. The necessary and sufficient condition for this is (Arias et al., 2002),

  4. Evolution operators pertaining to a region δ commute with all operators representing observables at spacelike separation from δ.

The conditions 3, 4 ensure compatibility with relativistic causal structure.

6.2.2   Collapse theory ontology

If the region between two Cauchy surfaces τ, σ is a bounded spacetime region δ, then the transition from the state on τ to the state on σ can be attributed to events within δ. Quantum states, on the other hand, are in some sense "nonlocal"—a better term is *nonseparable*. They cannot be regarded as supervening on local matters of fact. A quantum state encodes, not merely local matters of fact, but also correlations between distant regions.

  This is reflected in the way that a collapse affects the global state. Let τ and σ be two spacelike hypersurfaces, that coincide everywhere except on the boundaries of a bounded region δ, where σ lies to the future of τ. Let α be a bounded region that lies in both τ and σ, and hence is at spacelike separation from δ (see Figure 3). The two states can differ on probabilities assigned to observables pertaining to α. For example, consider the familiar EPR-Bohm scenario. Suppose the state on τ involves an entangled pair of spin-1/2 particles, one of which is located (to the extent that it can be located) within α, and suppose that in δ an experiment is performed on the other one, yielding, on σ, a near eigenstate of spin in some direction a that is correlated by the state on τ with spin of the other particle in direction b. Then the state on τ is a near-eigenstate of b-spin for the second particle. Thus, the two states $|\psi(\tau)\rangle$ and $|\psi(\sigma)\rangle$ differ on probabilities assigned to experiments performed in the forward domain of dependence of α. Relativistic causality precludes events occurring

17

Figure 3: A region common to two spacelike hypersurfaces.

in a region δ from affecting *local* beables pertaining to regions at spacelike separation from δ. If one thought that the differing probabilities assigned to observables pertaining to α by the states $|\psi(\tau)\rangle$ and $|\psi(\sigma)\rangle$ were local beables, this would give rise to a serious worry about the conceptual coherence of a relativistic collapse theory (see, *e.g.* Maudlin 1994, p. 209 for worries along these lines).

But the reduced states of α obtained from the global states $|\psi(\tau)\rangle$ and $|\psi(\sigma)\rangle$ are not obviously local matters of fact about α, and, indeed, what these reduced states are depends on events at spacelike separation from α. What *can*, and *should* be regarded as the intrinsic state of the spacetime region α, that is, its state stripped of implicit references to goings-on at a distance from α, is the past light cone state. This is the limiting state obtained from a sequence of spacelike hypersurfaces that converge on the past light cone of α. This state will reflect collapse events to the past of α, but not collapse events at spacelike separation.

This has implications for what the local beables of the theory could be. Ghirardi and Grassi (1994, p. 419) proposed a criterion for property attribution in the context of relativistic collapse theories. If A is a local observable with support in α, a system possesses a property $A = a$ when the expectation value of P_a, the projection onto the subspace $A = a$, is extremely close to 1, evaluated with respect to the past light cone state of α.

Ghirardi, Grassi, and Benatti (1995) proposed, as an appropriate local beable for collapse theories, a smeared mass density, with the smearing taken over lengths scales large compared to the atomic scale but small on a macroscopic scale. For a relativistic theory, the corresponding quantity would be a smeared energy-momentum density, whose 00 component is mass-energy density, which, in a nonrelativistic context, is dominated by the rest mass. Combining the choice of mass density as local beable with the past light cone criterion for property attribution, the mass density that in a relativistic context is taken to be a local beable in a region α is the mass density defined by the state on the past light cone of α (Tumulka, 2007; Bedingham et al., 2014). See Myrvold (2016, 2018, 2019) for further discussion of ontological issues in collapse theories.

6.2.3 A problem: stability of the vacuum

There is a difficulty with constructing a physically sensible collapse theory that respects the full symmetries of Minkowski spacetime. This has to do with the stability of the vacuum. A relativistic collapse theory must have a stable vacuum, which is not achievable via a straightforward extension of CSL to a relativistic quantum field theory.

If the theory were to produce excitations of the vacuum, it would have to furnish a probability distribution over possible excitations. The difficulty is the nonexistence of a bounded Lorentz-invariant measure over possible excitations.

This is readily seen in the simplest case, a free theory.[4] Let Δ be some subset of momentum space, and suppose that our collapse theory yields a certain probability $p(\Delta)$ per unit time of producing from the vacuum a particle with momentum in Δ. Lorentz invariance requires that there be the same probability for producing a particle with its momentum in any Lorentz boost of this volume.

This requirement cannot be satisfied. The problem is that there is no bounded invariant measure on the mass-shell consisting of all four-momenta of a given magnitude m. Any probability measure on the mass-shell, therefore, must break Lorentz symmetry. In the vacuum state, which is a Lorentz invariant state, there is nothing that could be used to define such a measure. To generalize beyond free theories, the issue is the absence of a bounded Lorentz-invariant measure over the space of possible excitations of the vacuum.

The requirement of a strictly stable vacuum applies only to theories that respect the full set of spacetime symmetries of Minkowski spacetime. One could, for example, have a theory on a lattice, with relativistic causal structure. The rate of energy increase could, for such a theory, be kept finite. Whether it could be kept below the threshold of current observations would depend on the lattice spacing and other parameters of the theory; see Ghirardi, Grassi and Pearle (1990, p. 1297) for discussion.

6.2.4 Impossibility of a relativistic collapse theory using only standard degrees of freedom

Though relativistic causality does not preclude the formulation of a relativistic collapse theory, there is a tension between the requirements of relativistic causality on a collapse theory and the need for a stable vacuum. It is possible to avoid an outright contradiction, but only at the cost of introducing nonstandard degrees of freedom, that do not obey anything like the usual laws of quantum evolution.

Here's the problem. Suppose we have a collapse theory with a stable vacuum state. Consider evolution from one surface τ to another, σ, through a bounded spacetime region δ. Stability of the vacuum means if the state on τ is the vacuum state, subsequent evolution is deterministic, in that there is only a single outcome that has nonzero probability, namely, maintenance of the vacuum. Deterministic evolution starting from vacuum means that, except perhaps for a set with zero probability of being realized, all collapse operators $\{K_\gamma(\delta)\}$ map the vacuum into multiples of the same vector. It can be shown (see Myrvold 2017 for details) that this will also hold for all standard states, where standard states are those that can be obtained by operating on the vacuum with operators of the sort that appear in standard quantum field theories. That is, for standard quantum field theories, the constraints imposed by relativistic causal structure, plus the constraint, imposed by Lorentz invariance, of a stable vacuum, are satisfied at the expense of rendering the theory deterministic, and thus incapable of suppressing the sorts of superpositions of macroscopically distinct states that collapse theories are meant to suppress.

This is not a no-go result for relativistic collapse theories *tout court.* Bedingham (2011a,b) and Pearle (2015) utilize a nonstandard field first introduced by Pearle (1993). This "index" or "pointer" field does not obey anything like an ordinary law of evolution; instead, its value

---

[4]This paragraph is based on the discussion in Pearle (2015, p. 2).

at any point is independent of its value at any other point. The vacuum state $|\Omega\rangle$ of this theory is the ground state of the standard fields *and* this pointer field. Standard operators operating on the vacuum state leave the pointer field in its ground state. On Bedingham's theory the pointer field couples to the standard fields, so that, if there are excitations of the standard fields, these produce excitations of the pointer fields. The theory tends to collapse towards eigenstates of smeared pointer-field operators, with accompanying collapse towards eigenstates of smeared matter fields.

One might also reject the requirement of having a theory that produces collapses in Minkowski spacetime. Okon and Sudarsky (2014, 2016, 2018) have proposed, in the context of semi-classical gravity, a modification of the CSL theory on which the collapse rate depends on local curvature. One might consider a modification of this proposal, according to which collapse rate goes to zero as curvature goes to zero (Bengochea et al., 2020). If there is a viable theory of this sort, it might satisfy the constraint of having a stable vacuum in Minkowski space, while producing collapse in the presence of matter.

6.3   Causal Quantum Theory

Standard quantum theory predicts correlations between outcomes of spacelike separated experiments that violate Bell inequalities. The crucial assumption of derivation of Bell inequalities is that the correlation be *locally explicable*, that is, that they can be explained by reference to conditions in the past lightcones of the experiments, such that a full specification of these conditions screens off the correlations. Experimental demonstrations of violations of Bell inequalities are usually taken to indicate correlations that are not locally explicable are a genuine physical phenomenon (see Myrvold, Genovese, and Shimony 2019 for an overview, and a discussion of the conditions required to reach this conclusion).

There is a loophole, however, identified by Kent (2005a), and called by him the *collapse locality loophole*. Tests of Bell inequalities that aim to establish nonlocality take care to ensure that the experiments involved take place at spacelike separation from each other, from choice of experimental setting to registration of result. A judgment of whether this condition is satisfied requires, therefore, a judgment of when an experiment is over. Typically it is assumed that this occurs once the result has been permanently recorded in some macroscopically readable device. However, as Kent observes, on some proposals, the quantum state will remain in a superposition corresponding to distinct outcomes beyond this point. One such proposal is the suggestion that state reduction takes place only when the uncollapsed state involves a superposition of sufficiently distinct gravitational fields (Diósi, 1987; Penrose, 1989, 1996). Another is Wigner's suggestion that conscious awareness of the result is required to induce collapse (Wigner, 1961).

Kent introduces a class of theories, which he calls *causal quantum theories*, on which collapses are localized events, collapse probabilities are conditioned only on events in the collapse's past light cone, and, as a consequence, collapses at spacelike separation from each other are probabilistically independent, conditional on events in their past lightcones. A theory of that sort differs in its predictions from standard quantum mechanics, but a decisive test requires experiments that are completed at spacelike separation from each other.

On a theory of the sort proposed by Diósi and Penrose, on which collapse occurs only when differing results correspond to sufficiently distinct gravitational fields, the condition

that the experiments be concluded at spacelike separation had not been enforced in any of the experimental tests of Bell inequalities as of the time of writing (2005). The experiment of Salart et al. (2008) closed the loophole for the particular proposals of Penrose and Diósi, though, as Kent (2018) pointed out, altering the Penrose-Diósi threshold by a few orders of magnitude would render these proposals compatible with the results of this experiment.

No experiment to date has addressed the collapse locality loophole if the collapse condition is taken to be awareness of the result by a conscious observer. See Kent (2018) for proposals of ways in which causal quantum theory could be subjected to more stringent tests.

7 Everettian Interpretations

Everettian interpretations make use of minimal machinery: they leave the usual, unitary dynamics untouched, and add nothing to the quantum state. The complications arise in connection with painting a sensible picture of the world with this minimal palette, and with making sense of probabilities or a working substitute for them.

Because of this minimality, worries about compatibility with relativity are also minimal. If the theory to be interpreted is a relativistic quantum field theory, then there is no question about the dynamics being compatible with relativistic causal structure; it is. As there are no additional beables, there is no question about whether their dynamics is compatible with relativity. Some proponents of Everettian interpretations have taken this to be one of the chief attractions of interpretations of this sort (see, *e.g.* Vaidman 1995, 2016, 2018).

Given a bounded subset α that is part of many Cauchy surfaces σ, σ', *etc.*, we can consider the reduced state of α,[5] evaluated with respect to various Tomonaga-Schwinger state vectors $|\psi(\sigma)\rangle$, $|\psi(\sigma')\rangle$. If the evolution from one Cauchy surface to another is unitary, and if the generators of evolutions through regions spacelike separated from α commute with operators representing observables pertaining to α, then the reduced state of α will be the same whether calculated from $|\psi(\sigma)\rangle$ or $|\psi(\sigma')\rangle$ or any Cauchy surface containing α. So, in that sense, on any quantum-state monist theory with deterministic, unitary dynamics, operations performed at a distance have no effect on states of affairs in α.

Is there nonlocality, in any sense, in an Everettian theory? It is not immediately clear how to pose the question. Unlike, say, the case of additional beables theories, we cannot ask (at least in any straightforward way) whether a choice of a parameter setting at one location affects the outcome of an experiment performed at another, because, on an Everettian theory, talk of *the outcome* of an experiment is unwarranted—experiments do not have unique outcomes. If the usual notions of causal dependence, such as act-outcome correlations, presuppose unique outcomes, then they cannot even be formulated in an Everettian context— and *ipso facto* they cannot be violated. Instead of saying that Everettian theories are causally local, perhaps it would be more apt to say: it is not the case that they are causally nonlocal.

Brown and Timpson (2016) have made a stronger claim about Everettian interpretations. Not only is there no dynamical nonlocality, they claim, there are no nonlocal correlations. The claim is that, in the Everettian context, talk of correlated outcomes of experiments makes

---

[5]Those who want to insist that states are to be assigned only to regions with nonempty interior may take this to mean: the state of the causal domain of dependence of α.

sense only when the outcomes of the distant experiments are compared, at which time both experiments are in the past. A similar argument has been made of which quantum theories are but a special case, by Brassard and Raymond-Robichaud (2013,2019). Analysis of this claim is beyond the scope of this chapter; we note only that compatibility of Everettian interpretations with relativity does not ride on whether a claim of this sort is accepted or rejected.

8   Non-Realist and Pragmatist Interpretations

There is a long tradition of denying that the ontology of the physical world contains anything corresponding to a quantum state. In recent years, views of this sort have been championed by advocates of a position called, by its proponents, *Quantum Bayesianism*, or *QBism* (Caves, Fuchs and Schack 2002; Fuchs, Mermin, and Schack 2014; Fuchs and Schack 2015). This is one variant of a class of views commonly referred to as *ψ-epistemic* views.[6] A somewhat related, but different, view is the pragmatist view of quantum mechanics defended by Richard Healey (2012, 2017a, 2020). See, also, Healey (2017b) for an overview of related views.

For Healey, quantum chances for events are assigned much as they have been in the discussion, above, of collapse theories. The chance of an event, say, an outcome of an experiment, is indexed by the spacetime event from which the chance is to be ascribed. The chance $Ch_x(e_A)$ of event e_A is, in effect, calculated using what in a dynamical collapse theory would be the past light cone state, which incorporates experimental outcomes that have occurred in the past light cone of x. On Healey's view, "Any agent who accepts quantum theory and is (momentarily) located at space-time point x should match credence in e_A to $Ch_x(e_A)$," where $Ch_x(e_A)$ is the chance evaluated conditional on events in the past light cone of the point x (Healey 2016, p. 179). Healey regards these chances as objective, in the sense that it is an objective fact that these chances are optimal credences for an agent that has epistemic access to all and only the facts in the past light cone of x. But they are not to be thought of as localized physical magnitudes.

> As we saw, the chance of outcome e_A does not attach to it in virtue solely of its physical description: the *chances* of e_A attach also in virtue of its space-time relations to different space-time locations. Each such location offers the epistemic perspective of a situated agent, even in a world with no such agents. The existence of these chances is independent of the existence of cognizers. But it is only because we are not merely cognizers but physically situated agents that we have needed to develop a concept of chance tailored to our needs as informationally deprived agents. Quantum chance admirably meets those needs: an omniscient God could describe and understand the physical world without it (Healey 2016, p. 183; see also Healey 2017c, pp. 175–176).

In Healey (2020), Healey argues that this view should be regarded as a form of scientific realism.

---

[6]See Harrigan and Spekkens (2010) for a discussion.

In their statement of QBism, Fuchs, Mermin, and Schack (2014) take a considerably more radical stance. First, they not only deny that quantum probabilities represent features of physical reality; for them, these probabilities are not objective in any sense—they are "personal judgments," not backed up by objective facts. Second, the scope of application of quantum mechanics is severely limited. In ordinary practice, physicists routinely use quantum mechanics to calculate probabilities about what will be observed by other people if they perform certain experiments, whether or not the agent carrying out the calculation will ever be aware of the results (which might, for example, not be realized within the agent's lifetime). In contrast, Fuchs *et al.* restrict an agent's assignment of probabilities to propositions concerning the agent's future experience.

The class of propositions to which the agent assigns probabilities all concern events along her own worldline, and therefore, contains no spacelike separated events.

> when any agent uses quantum mechanics to calculate "[cor]relations between the manifold aspects of [her] experience," those experiences cannot be space-like separated. Quantum correlations, by their very nature, refer only to time-like separated events: the acquisition of experiences by any single agent. Quantum mechanics, in the QBist interpretation, cannot assign correlations, spooky or otherwise, to space-like separated events, since they cannot be experienced by any single agent. Quantum mechanics is thus explicitly local in the QBist interpretation.

> And that's all there is to it (Fuchs, Mermin, and Schack, 2014, pp. 750–751).

There is no action at a distance, on this view, because, there are no distances between events that the agent assigns probabilities to. The spacetime of events to which the theory is applied is, in effect, a one-dimensional timelike worldline (or else a narrow world-tube). There is no superluminal signalling, because there is no signalling, of any kind (unless writing a note to your future self is said to be sending a signal to your future self).

Approaches that deny the reality of quantum states bypass some of the problems faced by realist approaches by limiting the scope of what is taken to be physical. In my opinion, however, such approaches must be regarded as promissory notes. If one denies that quantum states represent anything in physical reality, this raises the question of what *does* constitute physical reality. It does not seem, for example, that QBists wish to deny that distant places and other agents exists; they merely claim that such things are outside the scope of legitimate application of quantum mechanics. This raises the question of what physics *does* apply to them. If one says nothing at all, then ipso facto one says nothing that conflicts with relativity. The question still remains whether whatever it is that *does* exist can coexist peacefully with relativity.

9   Conclusion

There is a temptation to say that, because of quantum nonlocality, as revealed by Bell's theorem, that quantum theory is flatly incompatible with relativity. As we have seen, things are not so simple. It is only additional-beables theories, such as the de Broglie-Bohm theory and related theories, that must violate relativistic causal structure (at least, as long as the

additional beables are local beables). Collapse theories can respect relativistic causal structure; unlike the additional-beables theories, they need no distinguished relation of distant simultaneity. Full relativistic invariance, however, is possible only at the cost of introducing nonstandard degrees of freedom. For Everettian theories, relativity poses no special problem. For non-realist interpretations, it is difficult to render a verdict, as proponents of such views are usually less clear about what they do think *is* in the world than about what they think is *not*.

10 Acknowledgment

References

Arias, A., A. Gheondea, and S. Gudder (2002). Fixed points of quantum operations. *Journal of Mathematical Physics 43*(12), 5872–5881.

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

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.

Bedingham, D., D. Dürr, G. Ghirardi, S. Goldstein, R. Tumulka, and N. Zanghì (2014). Matter density and relativistic models of wave function collapse. *Journal of Statistical Physics 154*, 623–631.

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

Bell, J. S. (1981b). Quantum mechanics for cosmologists. In C. J. Isham, R. Penrose, and D. W. Sciama (Eds.), *Quantum Gravity 2: A Second Oxford Symposium*, pp. 611–37. Oxford University Press. Reprinted in Bell (1987b, 117–138).

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

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

Bell, J. S. (1989). Towards an exact quantum mechanics. In S. Deser and R. J. Finkelstein (Eds.), *Themes in Contemporary Physics II: Essays in Honor of Julian Schwinger's 70th Birthday*, pp. 1–26. Singapore: World Scientific.

Bell, J. S. (2007). The Trieste lecture of John Stewart Bell. *Journal of Physics A: Mathematical and Theoretical 40*, 2919–2933.

Bell, M. and S. Gao (Eds.) (2016). *Quantum Nonlocality and Reality: 50 Years of Bell's Theorem*. Cambridge: Cambridge University Press.

Bengochea, G. R., G. León, P. Pearle, and D. Sudarsky (2020). Discussions about the landscape of possibilities for treatments of cosmic inflation involving continuous spontaneous localization models. *European Physical Journal C 80*, 1021.

Berkovitz, J. and M. Hemmo (2005a). Can modal intepretations of quantum mechanics be reconciled with relativity? *Philosophy of Science 72*, 789–801.

Berkovitz, J. and M. Hemmo (2005b). Modal interpretations of quantum mechanics and relativity: A reconsideration. *Foundations of Physics 35*, 373–397.

Berkovitz, J. and M. Hemmo (2006). A new modal interpretation in terms of relational properties. In W. Demopoulos and I. Pitowsky (Eds.), *Physical Theory and its Interpretation: Essays in honor of Jeffrey Bub*, pp. 1–28. Dordrecht: Springer.

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

Brassard, G. and P. Raymond-Robichaud (2013). Can free will emerge from determinism in quantum theory? In A. Suarez and P. Adams (Eds.), *Is Science Compatible with Free Will?*, pp. 41–61. Dordrecht: Springer.

Brassard, G. and P. Raymond-Robichaud (2019). Parallel lives: A local-realistic interpretation of "nonlocal" boxes. *Entropy 27*, 87.

Brown, H. R. and C. G. Timpson (2016). Bell on Bell's theorem: The changing face of nonlocality. In Bell and Gao (2016), pp. 91–123.

Caves, C. M., C. A. Fuchs, and R. Schack (2002). Quantum probabilities as Bayesian probabilities. *Physical Review A 65*, 022305.

Clifton, R. (2000). The modal interpretation of algebraic quantum field theory. *Physics Letters A 271*, 167–171.

Costa de Beauregard, O. (1976). Time symmetry and interpreation of quantum mechanics. *Foundation of Physics 6*, 539–550.

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

Dieks, D. (2000). Consistent histories and relativistic invariance in the modal interpretation of quantum mechanics. *Physics Letters A 265*, 317–325.

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

Diósi, L. (1987). A universal master equation for the gravitational violation of quantum mechanics. *Physics Letters A 120*, 337–381.

Earman, J. and L. Ruetsche (2005). Relativistic invariance and modal interpretations. *Philosophy of Science 72*, 557–583.

Fine, A. (1982a). Hidden variables, joint probability, and the Bell inequalities. *Physical Review Letters 48*, 291–295.

Fine, A. (1982b). Joint distributions, quantum correlations, and commuting observables. *Journal of Mathematical Physics 23*, 1306–1310.

Friederich, S. and P. W. Evans (2019). Retrocausality in quantum mechanics. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2019 ed.). Metaphysics Research Lab, Stanford University.

Fuchs, C. A., N. D. Mermin, and R. Schack (2014). An introduction to QBism with an application to the locality of quantum mechanics. *American Journal of Physics 82*, 749–754.

Fuchs, C. A. and R. Schack (2015). QBism and the Greeks: why a quantum state does not represent an element of physical reality. *Physica Scripta 90*, 015104.

Ghirardi, G. and R. Grassi (1994). Outcome predictions and property attribution: the EPR argument reconsidered. *Studies in History and Philosophy of Science 25*, 397–423.

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.

Gisin, N. (1989). Stochastic quantum dynamics and relativity. *Helvetica Physica 62*, 363–371.

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

Healey, R. (2012). Quantum theory: a pragmatist approach. *The British Journal for the Philosophy of Science 63*, 729–771.

Healey, R. (2016). Local causality, probability, and explanation. In Bell and Gao (2016), pp. 172–194.

Healey, R. (2017a). Quantum states as objective informational bridges. *Foundations of Physics 47*, 161–173.

Healey, R. (2017b). Quantum-Bayesian and pragmatist views of quantum theory. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2017 ed.). Metaphysics Research Lab, Stanford University.

Healey, R. (2017c). *The Quantum Revolution in Philosophy*. Oxford: Oxford University Press.

Healey, R. (2020). Pragmatist quantum realism. In S. French and J. Saatsi (Eds.), *Scientific Realism and the Quantum*, pp. 123–146. Oxford: Oxford University Press.

Kastner, R. (2013). *The Transactional Interpretation of Quantum Mechanics: The Reality of Possibiliy*. Cambridge: Cambridge University Press.

Kent, A. (2005a). Causal quantum theory and the collapse locality loophole. *Physical Review A 72*, 012107.

Kent, A. (2005b). Nonlinearity without superluminality. *Physical Review A 72*, 012108.

Kent, A. (2018). Testing causal quantum theory. *Proceedings of the Royal Society A 474*, 20180501.

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.

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.

Malament, D. B. (1996). In defense of dogma: Why there cannot be a relativistic quantum mechanics of (localizable) particles. In R. Clifton (Ed.), *Perspectives on Quantum Reality: Non-Relativistic, Relativistic, and Field-Theoretic*, pp. 1–10. Dordrecht: Kluwer Academic Publishers.

Maudlin, T. (1994). *Quantum Non-Locality and Relativity: Metaphysical Intimations of Modern Physics*. Cambridge: Cambridge University Press.

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. (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*, pp. 238–260. Cambridge: Cambridge University Press.

Myrvold, W. C. (2017). Relativistic Markovian dynamical collapse theories must employ nonstandard degrees of freedom. *Physical Review A 96*, 062116.

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

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

Okon, E. and D. Sudarsky (2014). Benefits of objective collapse models for cosmology and quantum gravity. *Foundations of Physics 44*, 114–143.

Okon, E. and D. Sudarsky (2016). A (not so?) novel explanation for the very special initial state of the universe. *Classical and Quantum Gravity 33*, 225015.

Okon, E. and D. Sudarsky (2018). The weight of collapse: dynamical reduction models in general relativistic contexts. pp. 312–345. Cambridge: Cambridge University Press.

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

Pearle, P. (1993). Ways to describe dynamical state-vector reduction. *Physical Review A 48*, 913–923.

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

Penrose, R. (1989). *The Emperor's New Mind*. Oxford: Oxford University Press.

Penrose, R. (1996). On gravity's role in quantum state reduction. *General Relativity and Gravitation 28*, 581–600.

Salart, D., A. Baas, J. A. van Houwelingen, N. Gisin, and H. Zbinden (2008). Spacelike separation in a Bell test assuming gravitationally induced collapses. *Physical Review Letters 100*, 220404.

Schweber, S. S. (1961). *An Introduction to Relativistic Quantum Field Theory*. New York: Harper & Row.

Schwinger, J. (1948). Quantum electrodynamics. I. a covariant formulation. *Physical Review 74*, 1439–1461.

Simon, C., V. Bužek, and N. Gisin (2001). No-signalling condition and quantum dynamics. *Physical Review Letters 87*, 170405.

Struyve, W. (2011). Pilot-wave approaches to quantum field theory. *Journal of Physics: Conference Series 306*, 012047.

Tomonaga, S. (1946). On a relativistically invariant formulation of the quantum theory of wave fields. *Progress of Theoretical Physics 1*, 27–41.

Tumulka, R. (2007). The 'unromantic pictures' of quantum theory. *Journal of Physics A: Mathematical and Theoretical 40*, 3245–3273.

Vaidman, L. (1995). On the paradoxical aspects of new quantum experiments. In D. Hull, M. Forbes, and R. M. Burian (Eds.), *PSA 1994: Proceedings of the Biennial Meeting of the Philosophy of Science Association, Volume One: Contributed Papers*, pp. 211–217. East Lansing, MI: Philosophy of Science Association.

Vaidman, L. (2016). The Bell inequality and the many-worlds interpretation. In Bell and Gao (2016), pp. 195–203.

Vaidman, L. (2018). Many-worlds interpretation of quantum mechanics. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2018 ed.). Metaphysics Research Lab, Stanford University.

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