Relativistic Quantum Becoming

Wayne C. Myrvold[†]
Department of Philosophy
University of Western Ontario
London, Ontario, Canada N6A 3K7
wmyrvold@uwo.ca

Forthcoming in *The British Journal for the Philosophy of Science*
© 2002 *The British Journal for the Philosophy of Science*

ABSTRACT

In a recent paper, David Albert has suggested that no quantum theory can yield a description of the world unfolding in Minkowski spacetime. This conclusion is premature; a natural extension of Stein's notion of becoming in Minkowski spacetime to accommodate the demands of quantum nonseparability yields such an account, an account that is in accord with a proposal which was made by Aharonov and Albert but which is dismissed by Albert as a 'mere trick'. The nature of such an account is clarified by an extension to a relativistic quantum context of David Lewis' picture of objective chances evolving in time.

1. Introduction
2. Classical relativistic becoming
3. Relativistic quantum becoming, without collapse
4. Relativistic quantum becoming, with collapse
5. Objective chance, conditional probability, and definite properties
6. The nature of the wave-function
7. Conclusion

[†] I would like to thank Amit Hagar for drawing my attention to David Albert's paper, Abner Shimony for helpful comments on an earlier draft, and Howard Stein for a meticulous reading of the paper, from which it has benefited greatly.

1. Introduction

In a recent paper, 'Special Relativity as an Open Question' (Albert [2000]), David Albert argues that it is unlikely that there can be any realistic theory that accommodates quantum phenomena while maintaining compatibility with special relativity. He urges that we not dismiss nonrelativistic collapse theories, such as the GRW theory (Ghirardi, Rimini, and Weber [1986]), and suggests that it would be worthwhile to do experiments to look for the explicit violations of Lorentz invariance that such a theory entails.

I am in full agreement with Albert that special relativity is not to be treated as sacrosanct. Moreover, if experiments of the type suggested by Albert can be performed, then it will be worthwhile to do so, as the results, however they turn out, will constrain the range of empirically viable collapse theories. I am less pessimistic than Albert, however, about the prospects for a collapse theory that is genuinely, metaphysically compatible with special relativity, in the sense of depicting the world as 'unfolding within a four-dimensional Minkowski space-time', to use Albert's picturesque phrase. The reason for this is that, unlike Albert, I do not regard the suggestion made by Aharonov and Albert ([1984]), described in section 4, below, as merely a formal 'trick' that entails relinquishing any notion of the world unfolding in Minkowski spacetime.[1] The move suggested by Aharonov and Albert, I claim, results from a natural extension of the ordinary notion of the occurrence of events to accommodate the demands of both relativity and quantum theory.

In accepting Albert's challenge, that quantum theory ought to be able to provide an account of the world unfolding in time, I am accepting, with Albert, that the quantum-mechanical state description is, or ought to be, a representation of the state of a system, and not merely a representation of someone's knowledge of the system. I will also take the 'collapse' of

the quantum state that occurs upon measurement—and, presumably, also in some situations that are not measurements—to be a real physical process, and not merely a change in someone's knowledge of the system, or a shift in perspective.

The issue is whether there can be any concept of the world evolving in time, of events occurring in succession, of things coming to be and passing away, that is consistent both with quantum theory and with Minkowski spacetime. The core of such a notion is this: at any time t, there are events that have already occurred as of t; at a later time t' there are further events, namely, those which occur between t and t', that have occurred as of t', while those events that had already occurred as of t retain that status at t'. It will be our task to give an account of the occurrence of quantum events that retains this core.[2]

This core notion could, if one wanted to do so, be expressed in a tenseless language, provided only that the language contains an expression for the 'earlier than' relation between events. There is a tradition, stemming ultimately from McTaggart ([1908]), that holds that whatever can be described in tenseless language is not genuine temporal becoming, that a description consisting of timelessly true propositions cannot be a description of events happening in time or of objects changing in time. This paper has nothing to say on this issue, because, in spite of persistent and sincere attempts, its author has been unable to understand what the problem is meant to be. This paper, therefore, commits itself neither to a 'tensed' nor a 'tenseless' theory of time, nor does it take either tensed or tenseless language as primary. Furthermore, those who hold that, independently of considerations of either relativity or quantum mechanics, the notion of temporal becoming is problematic, will find nothing to assuage their concerns in what follows, which addresses the claim that the conjunction of relativity and quantum theory poses a special problem for such a notion.

Nor does this paper take a stand on 'presentism', which asserts that only the present is real, or on related views concerning the ontological status of past, present and future. We have more than one usage of words such as 'real' or 'exists' (that is, more than one use of existential quantifiers). There is one usage according to which what exists is what exists now (so that we can say, for example, that the Library of Alexandria once existed but no longer exists), and there is another usage that distinguishes what has existed, does exist, or will exist from what never has and never will exist (in this sense, the Library of Alexandria is real, but Borges' Library of Babel is not). Both senses of such words have a legitimate place in our discourse, and it is difficult to make sense of the question as to which sense *really* captures reality. The distinction made by each of these two usages is a real distinction. What is essential is that we not conflate the two senses, as doing so leads to confusion.

2. Classical relativistic becoming

As part of a response to an argument by Nicholas Maxwell ([1985]), Howard Stein ([1991]) defined a two-place relation R, with 'Rab' to be read as 'the state at b is definite as of a', satisfying the conditions,

i) R is a relation between spacetime points.

ii) R is a transitive relation.

iii) R is reflexive, *i.e.*. the state at a is already definite as of a.

iv) For any point a, there are points b such that $\sim Rab$.

v) R is definable in terms of the geometry of Minkowski spacetime.

vi) Rab holds whenever b is in the causal past of a.

Stein showed that the only relation satisfying these criteria is the relation *Rxy* that holds if and only if either $x = y$, or y is in the causal past of x. He then argued that this is, indeed, an acceptable candidate for the notion 'as of x, y has already become'.

Condition (i) is the condition that we are seeking a relation between spacetime points. Sometimes, however, we wish to consider the states of extended objects. As is well-known, to speak of the instantaneous state of an extended object in a relativistic context requires picking out some spacelike slice of the object—or, if one prefers, an instantaneous stage, along some foliation, of the object's history.[3] Failure to remember this leads to paradoxical results, as is illustrated by the well-known pole and barn example. The state history of the pole and barn given with respect to the barn rest frame, and the associated foliation of spacelike hyperplanes orthogonal to the barn's worldline, includes a state in which the pole lies entirely inside the barn. The state history given with respect to the pole's rest frame, on which instantaneous states are taken along spacelike hyperplanes orthogonal to the pole's worldline, includes no states in which the pole lies entirely inside the barn. This sounds as if the two state histories conflict as to what actually happens, so that therefore at most one of them can be correct. That no such conflict exists is seen by remembering that the states of the extended system of which one account speaks are states along spacelike slices of the system different from those of which the other account speaks.

We want to extend the notion of becoming to spacelike slices of extended objects. This is easily done, and, in fact, one can extend the notion to arbitrary spacetime regions. Let α and β be any two sets of spacetime points. Define the causal past of α as the union of causal pasts of all points in α. We will then say that β is definite as of α if and only if every point of β not contained in α lies in the causal past of α. That is,

$D\alpha\beta$ ('β is definite as of α') $\equiv_{df.}$ for every $y \in \beta$ there exists $x \in \alpha$ such that Rxy.

Since R is invariant under all automorphisms of Minkowski spacetime that preserve time-orientation, so is D. Note that D, as so defined, has the property that, for any two parallel spacelike hyperplanes—or, more generally, any two spacelike hypersurfaces belonging to the same foliation of Minkowski spacetime—the earlier of the two is definite as of the later, which is of course what one should expect from a reasonable notion of becoming.

For any timelike curve, Stein's R will determine, for any two points on the curve, whether or not one is definite as of the other. Similarly, for any foliation of spacetime into spacelike hypersurfaces, the relation D will determine, for any hypersurface σ belonging to the foliation and any spacetime region α, whether or not α is definite as of σ. This can be used to yield a description of the entire world unfolding along some foliation. We have, therefore, (at least) two notions of becoming in a relativistic context. One is Stein's local notion, which is independent of a choice of foliation. There is also a global notion, relativized to a foliation.[4] Both have their uses—if we want to talk about the global unfolding of the *world* in Minkowski spacetime, in terms of successive states of the world, then it is necessary to choose a foliation along which it unfolds.

A comment is in order to forestall a possible misinterpretation. Choosing a foliation is not the same as picking the perspective of some observer. The pedagogy of special relativity, with its tales of observers in trains and on spaceships, can lead the unwary to the impression that every observer is accompanied by a foliation consisting of hyperplanes orthogonal to the observer's world line, each of which constitutes that observer's infinitely extended present at some instant, and that accounts given with respect to different foliations correspond to the points of view of different observers. Even if such foliations existed,[5] an observer would have no knowledge of

events in this 'extended present' until such a time as the events are in the observer's causal past, and would have no privileged epistemic relation to such events. Nor is anyone under any obligation to refer events to his or her own rest frame, and we often do not—it is much more natural to say that one is travelling down the highway at 100 km/hr, than to say that the highway is passing beneath at 100 km/hr. When a speedometer in a car starts giving the speed of the car with respect to the speedometer's own rest frame, it's time to get the speedometer fixed.

Of our two notions of becoming, the first (local becoming) has the feature that there do not exist distinct spacetime points each of which is definite as of the other; on this notion, the *present* is the here-and-now. On the global, foliation-relative notion, we have, for each spacetime point P, a multitude of choices for the present of P, one for each spacelike hypersurface containing P. We do not have a notion of an instantaneous present that is both spatially extended and independent of an arbitrary choice of foliation. As Craig Callender ([2000]) has pointed out, both of these conditions have typically been implicit in philosophical discussions of becoming. Suppose, therefore, that a demand is made for a notion of an instantaneous present that includes more than the here-and-now, and which is *the* present, not merely one amongst a nondenumerable infinity of possible choices. To this demand it must be replied: neither our everyday dealings with the world, nor our physical theories, requires such a thing.

As Stein ([1991], § V) points out, endowing the present with spatial extent can be achieved, in a relativistic context, by endowing it with temporal extent. The event of my saying the word 'now' takes about one-third of a second, and hence fails to pick put a unique instant, but, in ordinary everyday contexts, no-one would object to this lack of precision. Following Stein, we will say that a point z is *contemporaneous* with a set η of spacetime points if and only if there exist x, y in η such that Rzx and Ryz. If η is any set of spacetime points containing my

utterance of the word 'now', the set of points contemporaneous with η has a spatial radius of at least 50,000 kilometres—large enough to encompass the entire Earth, and certainly large enough to include the people with whom I am speaking.

We therefore have a notion of extended present sufficient for everyday experience. It has been claimed (Maxwell [1985]; Dorato [1995], [1996]; see also the concluding paragraphs of Clifton and Hogarth [1995]) that Bell Inequality-violating quantum mechanical correlations between results of measurements performed at spacelike separation require reciprocal determinateness of spacelike separated events. It will be argued below that the foliation-relative notion of becoming suffices for this purpose.

3. Relativistic quantum becoming, without collapse

A feature of quantum mechanics, the feature that Schrödinger ([1935], 555) called 'not... *one* but rather *the* characteristic trait of quantum mechanics, the one that enforces its entire departure from classical lines of thought', is the fact that, for any two systems S_1 and S_2, among the possible states of the joint system $S_1 \oplus S_2$ are *entangled* states, in which neither of the composite systems possesses its own state vector, and in which the probabilities regarding the results of a pair of measurements, one on each system, do not factorize into independent probabilities regarding the individual measurements. This entanglement persists in principle, according to quantum mechanics, over an arbitrarily large spatial separation of the two systems, and has been observed in photons over distances greater than 10 kilometres (Tittel *et al.* [1998]). This means that the quantum-mechanical state description, in which the state of an individual system is represented by a ray in Hilbert space, or, equivalently, by a pure density operator, violates the condition that Don Howard ([1985], [1989]) has called 'separability', namely, that

the things dealt with in physics 'claim an existence independent of one another, insofar as they "lie in different parts of space"' (Einstein [1948], 321; translation by Howard, quoted in Howard [1985], 190). This is significant for our discussion, because Albert takes separability to be a necessary condition for what he calls 'metaphysical compatibility' with special relativity.

> What it is for a theory to be *metaphysically* compatible with special relativity (which is to say: what it is for a theory to be compatible with special relativity in the *highest* degree) is for it to depict the world as unfolding in a four-dimensional Minkowskian space-time. And what it means to speak of the world as unfolding within a four-dimensional Minkowskian space-time is (i) that everything there is to say about the world can straightforwardly be read off of a catalogue of the *local physical properties* at every one of the continuous infinity of *positions* in a space-time like that, and (ii) that whatever lawlike *relations* there may be between the values of those local properties can be *written down* entirely in the *language* of a space-time [like] that - that whatever lawlike relations there may be between the values of those local properties are *invariant* under *Lorentz-transformations* (Albert [2000], 3–4).

If this is right—if a failure of separability renders a theory metaphysically incompatible with relativity—then quantum mechanics is metaphysically incompatible with relativity. A specification of the values of all *local* properties that have definite values cannot possibly specify an entangled state vector. And if we can't specify the state vector (or, equivalently, the density operator), we cannot give a quantum-mechanical account of the state of the system.[6]

But this is too quick. As was pointed out above, any account, classical or otherwise, of the *world* unfolding in Minkowski spacetime, and not merely some pointlike bit of it, given in

terms of successive instantaneous states, requires a choice of foliation consisting of spacelike hypersurfaces along which the successive states of the world are defined. Albert's requirement (i) of separability is, therefore, too strong. An account of the world unfolding along some foliation may be reducible to an account of pointlike bits of the world unfolding along their worldlines, but it need not be, and the structure of Minkowski spacetime does not require that it be so reducible. To give an account of the world unfolding in Minkowski space-time, it suffices that, given any choice of foliation, (*a*) one can give an account of the world in terms of successive states along that foliation, in such a way that (*b*) no foliation is privileged as giving the uniquely correct state history, (*c*) the differences between state histories given with respect to different foliations can be ascribed entirely to the fact that the different foliations join up points of spacetime in different ways to form instantaneous hypersurfaces of simultaneity, on which global states are defined, and (*d*) a complete state history given with respect to one foliation uniquely determines the state history along any other foliation.

Condition (*c*) sounds as if it presupposes separability, but it does not. To satisfy this condition in a quantum context, what we will require is that differences in states on distinct hypersurfaces be attributable to *local operations* confined to the region between the hypersurfaces, in a sense that will now be explained. In algebraic quantum field theory, one associates with each open spacetime region O an algebra $R(O)$ of operators; the self-adjoint members of $R(O)$ correspond to observables measurable via operations confined to the spacetime region O. Given two disjoint bounded regions O_1, O_2, the state of affairs in the region $O_1 \cup O_2$ is not, in general, uniquely specified by specifying the definite values of all observables local to O_1 and O_2 (that is, observables represented by operators in $R(O_1)$ and $R(O_2)$)[7] that have definite values; one must specify also the definite values of those observables represented by operators in

$R(O_1 \cup O_2)$ that have definite values fixed by the state but not fixed, for that state, by definite values of observables local to O_1 and O_2. As a simple example, consider the singlet state of two spin-½ particles. This state is uniquely specified by the condition that the sum of spins of the two particles is zero, in every direction. The state cannot be specified by reference to spin states of the individual particles alone. Thus, the quantum state description violates Albert's requirement of separability; it is not true that everything there is to say about the world can be read off from a catalogue of local physical properties pertaining to spacetime points, or even from a catalogue of physical properties local to small bounded regions of spacetime.

Consider, now, two spacelike hypersurfaces, σ and σ', that coincide everywhere except for a small region in which σ' lies to the future of σ. That is: there is a bounded region Γ of spacetime lying in the future of σ and in the past of σ'; outside of Γ, what is in the past of σ is in the past of σ', and *vice versa* (See Figure 1). The quantum state of the world along σ and σ' can be represented by state vectors $|\psi(\sigma)\rangle$ and $|\psi(\sigma')\rangle$, respectively, or, equivalently, by pure density operators $\rho(\sigma)$, $\rho(\sigma')$.

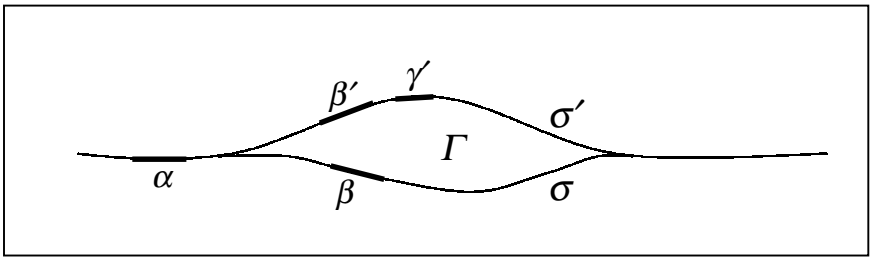

Figure 1.

Kraus ([1983]) argued that any 'physically realizable' operation can be modelled by a completely positive linear mapping φ of the set of trace-class operators into itself, and proved that any such function can be represented by a countable set of operators $\{K_i\}$ such that

$$\varphi(\rho) = \sum_i K_i \rho K_i^\dagger \text{ , with} \sum_i K_i^\dagger K_i \leq I.$$

An operation transforming pure states into pure states is called a *pure operation*, and can be represented via a single Kraus operator. Ordinary unitary evolution therefore has a Kraus representation consisting of a single unitary Kraus operator. A mapping φ that preserves the trace of ρ represents a *nonselective* operation; a mapping such that $\mathrm{Tr}(\varphi(\rho)) < \mathrm{Tr}(\rho)$ represents a *selective* operation. A selective operation may, for example, consist of a preparation of an ensemble followed by a measurement of some observable and a selection of those that yield a particular result. Collapse, regarded as a real physical process, can be thought of as a selective operation in which Nature does the selecting. Following Kraus, we will represent collapse as a pure Kraus operation that does not preserve the trace of the density operator; the new density operator in such a case is not $\varphi(\rho)$ but rather $\varphi(\rho)$ /Tr($\varphi(\rho)$). The advantage of choosing the function φ to be one that does not preserve norm is that the function can be a linear one (see Ghirardi *et al*. [1993]; Butterfield *et al* [1993] for discussion of this point). Alternatively, one can adopt the convention that quantum states are to be represented by trace-class operators whose traces are positive but need not be unity; if this convention is adopted, the expectation value of observable Ω in state ρ must be given by Tr($\rho\,\Omega$)/Tr(ρ).

Following Clifton and Halvorson ([2001]), let us define an operation *local* to a bounded spacetime region O to be an operation that has a Kraus representation consisting of operators belonging to $R(O)$. We will demand that state transitions satisfy the *local evolution condition*: that, if a hypersurface σ' lies nowhere to the past of the hypersurface σ, the state $\rho(\sigma')$ be obtainable from the state $\rho(\sigma)$ via an operation local to the region Γ between the two hypersurfaces σ and σ' .[8] The rationale for this requirement is that the two hypersurfaces of

simultaneity differ only on whether the events in the region between them are already past; any differences in the states on such hypersurfaces, therefore, should be attributable to events in Γ. If this is always the case, then condition (*c*) above is satisfied.

This requirement is satisfied by the usual versions of quantum field theory.[9] According to Albert, this is a step down from full metaphysical compatibility with relativity. Relativistic quantum field theories exhibit what Albert calls *dynamical compatibility* with special relativity, but not what Albert calls *metaphysical compatibility*, since the requirement of separability is not satisfied and cannot be satisfied by a relativistic quantum field theory. As I have said, regarding this as something less than full compatibility with relativity is unwarranted: at the risk of being repetitive, any foliation yields an account of the world unfolding in terms of successive states along that foliation, in such a way that no foliation is privileged as giving the uniquely correct state history, the differences between state histories given with respect to different foliations can be ascribed entirely to the fact that the different foliations join up points of spacetime in different ways to form an instantaneous hypersurface of simultaneity, on which a global state is defined, and a complete state history given with respect to any foliation uniquely determines the state history with respect to any other foliation.

To get a feel for how such accounts work, consider a system A located in a bounded subset α of the overlap of the two hypersurfaces σ and σ'. We can represent the state of A by a reduced density operator $\rho_A(\sigma)$ or $\rho_A(\sigma')$ obtained by tracing out the degrees of freedom of everything outside of α on σ or σ', respectively. If the transition from $\rho(\sigma)$ to $\rho(\sigma')$ is a pure unitary transformation—that is, representable by a single unitary Kraus operator—then we will have $\rho_A(\sigma) = \rho_A(\sigma')$.

Consider,[10] now, some system B, whose worldline (or world-tube) runs through Γ. Let β be the subset of σ occupied by B, and let β' be the subset of σ' occupied by B. Let C be some other system that on σ lies in β but in the interval between σ and σ' moves away, so that on σ' it lied entirely outside of β' in some region γ' disjoint from both α and β' (see Figure 1 again). Suppose that, on σ, the state of the combined system $A \oplus B \oplus C$ is a pure state representable by the state vector

$$|\psi(\sigma)\rangle_{ABC} = \left(c_1|+\rangle_A \otimes |+\rangle_B + c_2|-\rangle_A \otimes |-\rangle_B\right) \otimes |0\rangle_C ,$$

where $c_{1,2}$ are some nonzero constants, $|+/-\rangle_A$ are eigenstates of some local observable of A with eigenvalues ±1, and similarly for $|+/-\rangle_B$; for C, we want to consider an observable with three eigenvalues, 0, ±1.

The reduced density operator on σ for the system $A \oplus B$ is given by

$$\begin{aligned} \rho_{AB}(\sigma) = |c_1|^2 |+\rangle_A\langle+|_A \otimes |+\rangle_B\langle+|_B \quad &+ \quad c_1 c_2^* |+\rangle_A\langle-|_A \otimes |+\rangle_B\langle-|_B \\ + \quad c_1^* c_2 |-\rangle_A\langle+|_A \otimes |-\rangle_B\langle+|_B \quad &+ \quad |c_2|^2 |-\rangle_A\langle-|_A \otimes |-\rangle_B\langle-|_B . \end{aligned}$$

Define a *disentangled* state of A and B as a state whose density operator can be written as a convex sum of product states, and an *entangled* state as one that is not disentangled. The state $\rho_{AB}(\sigma)$ is entangled.

Now suppose that, before C moves away from B, the two systems interact in such a way that

$$|+\rangle_B \otimes |0\rangle_C \to |+\rangle_B \otimes |+\rangle_C , \qquad |-\rangle_B \otimes |0\rangle_C \to |-\rangle_B \otimes |-\rangle_C .$$

Then the state on σ' of the combined system $A \oplus B \oplus C$ is given by

$$|\psi(\sigma')\rangle_{ABC} = c_1|+\rangle_A \otimes |+\rangle_B \otimes |+\rangle_C + c_2|-\rangle_A \otimes |-\rangle_B \otimes |-\rangle_C ,$$

and the reduced density operator for $A \oplus B$ is the improper mixture,

$$\rho_{AB}(\sigma') = |c_1|^2 \, |+\rangle_A\langle+|_A \otimes |+\rangle_B\langle+|_B \quad + \quad |c_2|^2 \, |-\rangle_A\langle-|_A \otimes |-\rangle_B\langle-|_B .$$

This is a *disentangled* state, because it is a convex sum of two product states. Of course, what has happened is that we have exported the entanglement between A and B to the environment of $A \oplus B$; A remains entangled with the joint system $B \oplus C$. The point is that, as far as the relations between A and B considered in isolation from everything else are concerned, the operations between σ and σ' have destroyed the entanglement of A and B, and any broader entanglement the two participate in will not be revealed by measurements on A and B alone. On the hypersurface σ, A is entangled with B; on σ', it is not entangled with B (though it is entangled with $B \oplus C$). Therefore, in order to ask whether, at a certain stage of its history, the system A is entangled with a distant system B, one must specify *what stage* of B's history is being considered. One and the same stage of A's history may be entangled with one stage of B's history and not with another—a fact that we may call the *relativity of entanglement*.[11] This should perhaps not be surprising, as the question of whether or not A is entangled with B cannot be regarded as a matter of the intrinsic state of A alone, but is a matter of a relation between A and B.

4. Relativistic quantum becoming, with collapse

The usual linear, unitary dynamics of quantum theory do not yield a description of the sorts of events that we usually regard as occurring. These surely include measurements resulting in definite outcomes, and not merely in superpositions of the potential outcomes. To paraphrase Bell ([1987], 201) either the quantum state description is not everything, or the unitary quantum state evolution is not quite right.[12] Modal theories adopt the first alternative, supplementing the quantum state description with enough extra properties to account for definite outcomes of

experiments; collapse theories, the second, modifying the dynamics so as to yield definite outcomes (or at least close approximations to definite outcomes) for experiments.[13] The issue of the Lorentz invariance of modal theories has been discussed elsewhere (Dickson and Clifton [1998], Dieks [1998]). It is the purpose of this paper to evaluate Albert's remarks about the compatibility of collapse theories with relativity.

Suppose that, for each foliation of space-time, we have an account of the unfolding of the states of the world along the successive hypersurfaces that make up the foliation. Can we accommodate collapse within such a picture without sacrificing the conditions (*b*) that no foliation is privileged as giving the uniquely correct state history, and (*c*) that the differences between state histories given with respect to different foliations can be ascribed entirely to the fact that the different foliations join up local states in different ways? The question reduces to whether or not a collapse theory can satisfy the local evolution condition while maintaining the usual quantum mechanical probabilities, which satisfy the no-signalling theorem (Eberhard [1978]; Ghirardi, Rimini and Weber [1980]) and do not pick out a preferred relation of distant simultaneity.

The local evolution condition may sound as if it presupposes determinism, but it does not. The condition says that from the state on σ we can obtain the state on σ' via an operation local to the region Γ between σ and σ'. *Which* local operation yields the state on σ' may be a matter of chance, undetermined by the state on σ and determined only by the actual course of events—which may include objective chance events— in Γ.

Suppose that an ideal measurement of an observable Ω with a nondegenerate, discrete spectrum is performed in the spacetime region Γ between σ and σ', and that all other changes of state in Γ are negligible. Suppose that the result of the measurement is a certain eigenvalue ω_k of

Ω. Then, if the state on σ is $\rho(\sigma)$, then the (unnormalized) state on σ' is given, according to the collapse postulate, by

$$\rho(\sigma') = P_k^{\Omega} \rho(\sigma) P_k^{\Omega},$$

where P_k^{Ω} is the projection onto the eigenspace of Ω corresponding to ω_k. Since Ω, the observable being measured, is local to the region in which the measurement occurs, so is P_k^{Ω}, and this state transition is a selective operation local to the region in which the measurement occurs. The local evolution condition is satisfied by the usual collapse postulate applied to an ideal measurement. All differences between the states on σ and σ' are due to occurrences within Γ—that is, they are attributable to the fact that Γ lies in the future of σ and the past of σ'. More general measurements will be representable as a combination of a local interaction between the system and the measurement apparatus with a collapse onto one of the pointer subspaces of the measurement apparatus, and these also will be selective operations local to the region in which the measurement occurs.

We will adopt, as a working hypothesis, the assumption that our collapse theory yields state transitions that admit of Kraus representations, and will assume that the theory satisfies the following conditions. For any two hypersurfaces σ and σ', with σ' nowhere to the past of σ, and any state $\rho(\sigma)$, the theory should specify a set of possible states on σ', each of which is obtainable from $\rho(\sigma)$ by a pure operation local to Γ, the region between σ and σ'. Moreover, the mixture consisting of the probabilistically weighted sum of all possible states on σ', given that the state on σ is $\rho(\sigma)$, should be representable as

$$\sum_i K_i \rho(\sigma) K_i^{\dagger},$$

where each K_i is local to Γ, and $\sum K_i^\dagger K_i = I$.[14] These conditions suffice to ensure the satisfaction of the local evolution condition and the no-signalling theorem.

We have obtained, not as an *ad hoc* assumption, but merely as a consequence of extending state transitions to collapse transitions induced by local operations, an account that mirrors a proposal made by Aharonov and Albert ([1984]). According to this proposal, for each foliation, the collapse occurs along that spacelike hypersurface that passes through the measurement event. This is precisely what must be the case if collapses can be represented as local operations. Assume that the collapse takes place within a bounded spacetime region Γ. If Γ lies in the future of a hypersurface σ, the state on σ will be the uncollapsed state, and if Γ lies in the past of σ, the state on σ will be the collapsed state; what the state is like on a hypersurface intersecting Γ will depend on the details of the collapse model. A similar idea, though restricted to foliations consisting of spacelike hyperplanes, has been defended by Gordon Fleming ([1986], [1989], [1996]), and a similar account of collapse has evolved naturally from the work of Ghirardi, Rimini, Weber, and Pearle as a result of the endeavour to extend the nonrelativistic continuous spontaneous localization model (CSL) to a relativistic context (see Pearle [1999], Ghirardi [2000], and references therein). Although he is one of the authors of this proposal, Albert disparages it, calling it a 'trick' by means of which the theory is made '*formally* compatible with special relativity' (Albert [2000], 6). Such merely formal compatibility, according to Albert, pays a high price.

> As things stand now we have let go not only of *Minkowski-space* as a realistic description of the stage on which the world is enacted, but (in so far as I can see) of *any* conception of that stage *whatever*. As things stand now (that is) we have let go of the idea of the world's having anything along the lines of a narratable story

> at *all*! And all this just so as to guarantee that the fundamental laws remain exactly invariant under a certain hollowed-out set of *mathematical transformations*, a set which is now of no particularly deep conceptual *interest*, a set which is now utterly disconnected from any idea of an *arena* in which the world occurs (Albert [2000], 6).

This conclusion is, I believe, premature. If one takes quantum nonseparability seriously, then the sort of picture that emerges from Aharonov and Albert's foliation-relative collapse postulate is just the sort of picture that one would expect to have of quantum states unfolding in time. We have, for each foliation, a narratable story of global states unfolding, in such a way that state transitions from one hypersurface to another are solely the result of local operations. This renders it possible to go from the narratable history along one foliation to the history along any other, in such a way that no foliation is uniquely picked out as giving the correct story. And all this takes place against the background of Minkowski spacetime.

The local evolution condition, together with the usual microcausality assumption that operators pertaining to spacelike separated observables commute, has some interesting consequences.

Theorem. Let σ and σ' be spacelike hypersurfaces such that no part of σ' lies in the past of σ. Let Λ be a region spacelike separated from the region Γ between σ and σ', such that the intersection of Λ with σ coincides with its intersection with σ'. If the state $|\psi(\sigma)\rangle$ is an eigenstate of an observable Ω local to Λ, with eigenvalue ω_k, then so is $|\psi(\sigma')\rangle$.

Proof. If $|\varphi\rangle$ is any eigenvector of an operator Ω, with eigenvalue ω_k, and K is any operator that commutes with Ω such that $K|\varphi\rangle$ is nonzero, then $K|\varphi\rangle$ is also an eigenvector of Ω with

eigenvalue ω_k. By the local evolution condition, the transition from $|\psi(\sigma)\rangle$ to $|\psi(\sigma')\rangle$ is induced by a pure operation local to the region Γ between σ and σ'. That is,

$$|\psi(\sigma')\rangle = K\,|\psi(\sigma)\rangle$$

for some K local to Γ. Since Γ is spacelike separated from Λ, the microcausality condition entails that K commutes with Ω; therefore, $K\,|\psi(\sigma)\rangle$ is an eigenvector of Ω with eigenvalue ω_k.

Because of the temporal asymmetry of the local evolution condition, we cannot remove the condition that σ' lies nowhere in the past of σ and conclude that, if the state $|\psi(\sigma)\rangle$ is an eigenstate of an observable Ω local to some region Λ, then $|\psi(\sigma')\rangle$ is also an eigenstate of Ω for *every* hypersurface σ' whose intersection with Λ coincides with that of σ. It does, however, follow that, *if* $|\psi(\sigma')\rangle$ is an eigenstate of Ω, it is an eigenstate with the same eigenvalue.[15]

Corollary. Let σ and σ' be spacelike hypersurfaces, and let Λ be a region spacelike separated from the region Γ between σ and σ', such that the intersection of Λ with σ coincides with its intersection with σ'. If $|\psi(\sigma)\rangle$ and $|\psi(\sigma')\rangle$ are both eigenstates of an observable Ω, local to Λ, with eigenvalues ω and ω', respectively, then $\omega = \omega'$.

Proof. Suppose that $|\psi(\sigma)\rangle$ and $|\psi(\sigma')\rangle$ are both eigenstates of an observable Ω, local to Λ, with eigenvalues ω and ω'. Consider a spacelike hypersurface σ'', which contains the common intersection of σ and σ' with Λ, but which lies nowhere to the past of either σ or σ'. Such a hypersurface can be constructed by joining the overlap of σ and σ' with the part of σ that lies to the future of σ' and the part of σ' that lies to the future of σ, and applying a smoothing procedure, if necessary, at the joints to yield a hypersurface that everywhere has a timelike normal. By the above theorem, since σ'' lies nowhere to the past of σ and $|\psi(\sigma)\rangle$ is an eigenstate of Ω with eigenvalue ω, so is $|\psi(\sigma'')\rangle$. Similarly, since $|\psi(\sigma')\rangle$ is an eigenstate of Ω with

eigenvalue ω', so is $|\psi(\sigma'')\rangle$. $|\psi(\sigma'')\rangle$ cannot be an eigenstate of Ω with two distinct eigenvalues, so $\omega = \omega'$.

5. Objective Chance, Conditional Probability, and Definite Properties

The quantum state of a system is used to calculate the probabilities of the possible outcomes of measurements performed on a system. Unless we are to regard the quantum state as supervening on some more complete state, whose evolution is governed by a deterministic theory (and the prospects for a relativistic deterministic theory are not good), we should regard the quantum-theoretical probabilities as matters of objective chance.

David Lewis, in 'A Subjectivist's Guide to Objective Chance' (Lewis [1980]), presents a picture of objective chances as quantities that evolve in time as more and more events make the transition from the possible to the actual. The picture presented by Lewis is, I believe, both in accord with the way we usually think about chances and also appropriate for use in physical theory, once suitable adjustments are made to adapt it to a relativistic context. In his paper, Lewis provides the following example of chance evolving (in reading the example, think of a coin toss as an objective chance event).

> Suppose you enter a labyrinth at 11:00 A.M., planning to choose your turn whenever you come to a branch point by tossing a coin. When you enter at 11:00, you may have a 42% chance of reaching the center by noon. But in the first half hour you may stray into a region from which it is hard to reach the center, so that by 11:30 your chance of reaching the center by noon has fallen to 26%. But then you turn lucky; by 11:45 you are not far from the center and your chance of

reaching it by noon is 78%. At 11:49 you reach the center; then and forevermore your chance of reaching it by noon is 100% (271).

On Lewis's account, the way in which an objective chance function evolves is this: at each time t, there is a chance distribution Pr_t, that assigns a chance to each possible event, past, present, or future. The chance distribution $Pr_{t'}$ at a time t', later than t, is obtained from the distribution at time t by conditionalization on a complete history of events in the intervening time.[16]

In a Galilean spacetime, possessing a distinguished foliation, this prescription is unambiguous. In a relativistic context there is still more to be filled in. Given two points on an object's world-line, an earlier point P and a later point P', what do we include in the set of events on which we conditionalize in order to obtain the chances of events, as of P', and what do we exclude?

At a minimum, we should include the events in the timelike past of P', and exclude all events in the timelike future of P'. This would be enough to be unambiguous, if our chance function satisfied the Reichenbach Common Cause principle (Reichenbach [1956], §IV.19). According to this principle, if events A and B are correlated, such correlation is 'screened off' by events in the past[17] of A and B. That is, if $\Pr(A\&B) \neq \Pr(A)\Pr(B)$, then there is a partition $\{C_i\}$ of possible events in the past of A and B such that $\Pr(A\&B|\ C_i) = \Pr(A \mid C_i)\ \Pr(B \mid C_i)$ for all i. Although Reichenbach frames his Common Cause Principle in terms of a frequentist conception of probability (A and B are, therefore, not individual events but event-*types* in some ensemble of events), the principle can be recast in terms of objective chances, without doing violence to the underlying conception. To do so, we must regard the probability distribution involved as a chance distribution $\Pr_t$ as of some time t. Since past events are no longer chancy, and we want A

and B to be correlated on $\Pr_t$, t must be in the past of t_A and t_B. As was mentioned above, Lewis takes chances to be updated by conditionalization on a *complete* state history over the intervening time. Following Butterfield ([1989]), we will, therefore, take the common cause to be the complete state of the common past of A and B.

The Common Cause Principle, expressed in terms of objective chances, then, is as follows.

> (CC) Let $\Pr_t$ be the chance distribution as of some time t. Let A and B be events occurring at times t_A and t_B, respectively, in the future of t. Then, for each possible state C of the common past of A and B, $\Pr_t(A\&B|C) = \Pr_t(A|C)\ \Pr_t(B|C)$.

When the Common Cause Principle is satisfied, one can unambiguously update chances by conditionalization on the causal past of an event. Quantum probabilities involving spacelike entanglement do not, however, satisfy the Reichenbach principle, as van Fraassen ([1982]) has convincingly argued.[18] It matters, therefore, which events, spacelike separated from the event whose chance we wish to calculate, are included in our conditionalization. For each spacelike hypersurface σ, we will obtain a chance distribution Pr_σ that is conditionalized on all events in the past of σ.

This, of course, is not new, but has been implicit all along in our talk of relativizing the quantum state to a spacelike hypersurface. If we want to ask what the probability is of an outcome of a measurement on a system A in a spacetime region Γ, we must ask, relative to which spacelike hypersurface? Which is to say: on which events are the probabilities to be conditional?

Minimalists will want to restrict these to the causal past of Γ, on the grounds that these are the only events that are objectively in the past of Γ. This is all very well, but it should be

mentioned that, although we typically do regard chances as conditional on past events, it is not nonsensical to conditionalize on events not in the past. Suppose, for example, that a pair of spin-½ particles, S_1 and S_2, are prepared in the singlet state, and that measurements of spin-z are performed on them, not at spacelike separation, but consecutively, with the measurement on S_1 in the causal past of the measurement on S_2. Asked for the probability that the outcome of the measurement on S_1 will be +1, we will naturally answer that it is ½; the probability that the outcome of the measurement on S_2 will be +1 will then be either 0 or 1, depending on the outcome of the first measurement. Suppose, however, that one is asked for the probability that the outcome of the first measurement will be +1, conditional on the supposition that the outcome of the second measurement is –1. One way to make sense of this is to suppose that we are going to prepare a post-selected ensemble by preparing a large number of singlet pairs, performing spin-z measurements first on one member and then on the other member of the pairs, and then discarding those runs of the experiment for which the result of the second measurement is +1. The question then is: what fraction of the remaining subensemble will we expect to have +1 as the outcome of the first measurement? And the answer, of course, is 100%.

Both these probabilities, the one conditional only on events in the past of the measurement in question, the other including a future event in its domain of conditionalization, have equal right to be regarded as objective chances. The difference between them is not a subjective one; objectively different questions are being asked and answered in each case. Both probability assignments are susceptible of experimental confirmation or refutation, and in neither case do we regard them as having anything to do with incomplete knowledge of the objective states of affairs that give rise to them.

If it makes sense for chances to be conditional on future events, then it can hardly be nonsensical to have chances be conditional on spacelike separated events. This is precisely what quantum states on spacelike hypersurfaces do. Suppose, now, that the measurements on our singlet pair are performed at spacelike separation. Let σ be a spacelike hypersurface passing just to the past of both measurements, and let σ' be a spacelike hypersurface coinciding with σ in the region in which S_1 is located, but passing to the *future* of the measurement on S_2. Then the state on σ assigns a chance ½ to the proposition that the measurement on S_1 will yield +1, and, by the collapse postulate, the state on σ' either assigns a chance 0 or chance 1 to this same proposition. And this is as it should be, because the state on σ' yields chances *conditional on the entire past of* σ', and this includes the determinate outcome of the measurement performed on S_2.[19]

As the state of the world evolves, some local properties change their values, some local properties that previously did not have definite values acquire them, and some local properties that previously had definite values lose them. When do these events occur?

The usual eigenvalue-eigenstate link has it that an observable represented by an operator Ω has a definite value ω_k if and only if the quantum state $|\psi\rangle$ satisfies $\Omega|\psi\rangle = \omega_k |\psi\rangle$. Take, therefore, some spacelike slice α of an object A. Three possibilities initially present themselves as candidates for explication of the statement that, as of α, A has definite value ω_k of the observable represented by the operator Ω.

i) Relative to a hypersurface: A has definite value ω_k of the observable represented by Ω, relative to hypersurface σ containing α, if and only if $\Omega|\psi(\sigma)\rangle = \omega_k |\psi(\sigma)\rangle$.

ii) Absolute (that is, relative only to α). A has definite value ω_k of the observable represented by Ω, if and only if

a) $\Omega|\psi(\sigma)\rangle = \omega_k |\psi(\sigma)\rangle$ for some hypersurface σ containing α.

b) $\Omega|\psi(\sigma)\rangle = \omega_k\,|\psi(\sigma)\rangle$ for every hypersurface σ containing α.

The first of these two absolute notions has some initial plausibility—it satisfies a certain stability condition, namely, that if with respect to the state on some hypersurface, a definite prediction can be made regarding the outcome of a measurement of Ω within α, this is a fact about the state of affairs within α and does not depend on considerations of anything spacelike separated from α. The appeal of such a stability condition, however, stems from the assumption, perhaps implicit, of separability, an assumption which should be rejected. And as Clifton, Pagonis, and Pitowsky ([1992]) have shown, the condition (ii *a*) leads to absurd results when applied to the GHZ state, in the form of definite values being possessed by each of the three entangled systems in the GHZ state that are jointly *incompatible* with that state.

This leaves us with the hypersurface-relative notion (i) and the absolute notion (ii, *b*). Note that both are *invariant* notions, not frame- or coordinate- dependent (and certainly not observer-dependent). Since (i) requires picking a hypersurface along which the state is to be evaluated, definite values in the sense of (i) are not candidates for local, intrinsic properties of the systems located in the region α. Definite properties in sense (ii, *b*) can, on the other hand, be regarded as local, intrinsic properties. This is the sense favoured by Ghirardi and his collaborators (see Ghirardi and Pearle [1991]; Ghirardi and Grassi [1994]; Pearle [1997], [1999]; Ghirardi [2000]).[20] In a variant on their usual way of defining this notion, Ghirardi ([2000], 1364) proposes the following:

> A system at a spacetime point P possesses the objective property $\Omega = \omega_k$ if and only if the state on the past light-cone is an eigenstate of Ω belonging to ω_k.

As we can see from the theorem of §4, above, it is precisely the local evolution condition, together with the microcausality condition, that guarantees the equivalence of definite properties

so defined with the definition usually given by Ghirardi *et al*., which is our definition (ii, *b*). This notion may be regarded as a notion of a property definite, in Stein's sense, as of a point on the object's worldline. Sense (i), which is a notion of property definite as of a hypersurface σ, is the sense appropriate to the global notion of becoming identified in §2. Both concepts have their uses, and here also the question of which of the senses really captures reality does not seem to be well-posed. What is essential is that we not conflate the two senses, and in particular, that we remember that sense (i) makes reference to a spacelike hypersurface and does not ascribe local properties to systems.

6. The nature of the wave-function

In non-relativistic quantum mechanics, we are used to regarding the Schrödinger wave function $\Psi(\mathbf{x},t)$ as a field on space-time, such that the integral of $|\Psi(\mathbf{x},t)|^2$ over a volume V gives the probability that a position measurement at time t will find the particle in the volume V. This field, however, behaves oddly upon measurement; if the particle is found to be in V, the collapse postulate has it that $\Psi(\mathbf{x},t)$ goes instantaneously to zero everywhere outside of V. Even apart from relativistic considerations, this is unexpectedly nonlocal behaviour; it is not surprising that a realistic interpretation of the wave function was thought by some to be impossible, and that it was interpreted epistemically.

In a relativistic quantum theory, the concept of a localizable particle—and *ipso facto*, that of a wave function yielding the probability density for the location of such a particle—is problematic (see Malament [1996]; Halvorson and Clifton [2002]). One is tempted to simply say that there are no wave-functions in a relativistic quantum theory, and leave it at that. But this is too quick. Non-relativistic quantum mechanics ought to be obtainable by a limiting process from

relativistic quantum field theory; we ought, therefore, to construct, at least for certain special states, a quantity that acts at least approximately like the wave function.

Since we have relativized states to spacelike hypersurfaces, the wave function, or its surrogate, will not be a function on spacetime but rather a *functional* taking as arguments a pair consisting of a hypersurface and a point on that hypersurface. This may seem like an arbitrary move, a mere trick that ensures formal compatibility with relativity, but no more. Since only a single particle is involved, it may seem that entanglement will play no role, and that here violations of separability ought to be unimportant.

To see that this is not correct, let us consider, without going into too much detail, how something like a wave function can emerge from a relativistic quantum field theory. Take a spacelike hypersurface σ, and divide it up into disjoint regions $\{\Pi_i\}$. Let $|1_i\rangle$ be a state in which a particle detector located in Π_i is certain to detect one particle, and particle detectors located in all other regions are certain to detect none—or, at least, a state in which these conditions are *approximately* satisfied. It is probably too much too ask, in a relativistic context, that they be exactly satisfied. In a relativistic quantum field theory, *every* detector has a nonzero probability of going off in the vacuum state, or indeed in any state of bounded energy (Haag [1992]; 272). A more realistic assumption, therefore, would be that the probabilities mentioned are *close*, rather than equal, to one and zero, and that certainty is approached in a non-relativistic limit. This is possible if the regions $\{\Pi_i\}$ are large compared to the Compton wavelength of the particle.[21]

$|1_i\rangle$ is therefore (at least approximately) a single-particle state, one in which the position of the single particle is *located* in the region Π_i. It is not specifiable by giving the values of observables local to Π_i, however, as the condition that the state is a single-particle state is a

global condition. And if the state is not a single-particle state, there can be no sense attached to the notion of the position of *the* particle.

Consider, now, a state of the form

$$|\psi(\sigma)\rangle = \sum_i \psi_i |1_i\rangle .$$

This, too, is (at least approximately) a single-particle state. It is also a state that is entangled across the sets of observables belonging to the regions $\{\Pi_i\}$. The probability of finding a particle in one region is not independent of the probability of finding it in another region—if the particle is detected in any one of the regions $\{\Pi_i\}$, it is highly probable that *no* particle will detected in any other part of σ. Because of this entanglement, a collapse finding the particle in one of the regions, say, Π_k, can be realized via an operation local to the region Π_k. Let P_k be a projection onto the subspace in which a particle detector in Π_k is certain to yield a single particle. This projection can be local to Π_k because for a state to be an eigenstate of P_k entails, in general, no commitment as to whether detectors in other regions will fire. P_k serves as our selective Kraus operator for the collapse transition.

$$|\psi(\sigma)\rangle \quad \rightarrow \quad P_k|\psi(\sigma)\rangle = \psi_k |1_k\rangle .$$

If we mistakenly thought of the ψ_i's as local quantities, this transition would have a paradoxical air to it—the values of ψ_i are sent instantaneously to zero, for all $i \neq k$, by an operation local to Π_k. ψ_i is not a local quantity, though; it is the coefficient of expansion of our entangled single-particle state in terms of the globally specified state $|1_i\rangle$.

When one goes to the nonrelativistic limit, the Π_i's can be made arbitrarily small, and hence one can obtain a function $\Psi_\sigma(\mathbf{x})$ instead of the ψ_i's. Moreover, in a Galilean spacetime there is a uniquely distinguished foliation; one can, therefore, uniquely specify a hypersurface of

simultaneity by merely specifying the time t. We therefore obtain a function $\Psi(\mathbf{x},t)$. The values of this function, however, retain their character as coefficients of expansions of the state in terms of global states (and thus cannot be regarded as local quantities), and it remains true that any single particle state (or for that matter, any n-particle state for fixed n), is entangled across the set of observables belonging to any two regions in which the probability of finding the particle (or any one of the particles) is nonzero. This nonlocal character of the wave function is evinced in its behaviour in a collapse.

7. Conclusion

In his paper, Albert distinguishes five 'degrees of compatibility' with special relativity. Albert's 'metaphysical compatibility' includes separability as a requirement and is too strong; a theory that violates separability is not *ipso facto* incompatible with relativity. Violation of Albert's metaphysical compatibility is, therefore, not any sort of incompatibility with relativity at all. Albert's second category, which he calls 'dynamical compatibility', is satisfied by relativistic quantum field theories without collapse. Yet the very features that are responsible for this compatibility can be appropriated by a collapse theory—quantum theory, without or without collapse, can fall into this category, which does not fall short of full compatibility with relativity. With or without collapse, relativizing quantum states to spacelike hypersurfaces is not a merely 'formal trick', but a natural extension of the notion of becoming in a relativistic spacetime to accommodate quantum nonseparability; what is so yielded is not a hollowed-out 'formal compatibility'. Although a fully-worked out solution to the measurement problem is still required, there is ample room for hope for a solution that is fully compatible with relativity.

Notes

[1] Although I will speak of Minkowski spacetime throughout, nothing in this paper will depend on the precise structure of Minkowski spacetime, and the discussion will apply equally well to any globally hyperbolic spacetime—that is, any spacetime possessing a foliation of Cauchy surfaces.

[2] In what follows, I will often adopt Albert's picturesque language, and speak of the world as 'unfolding' in time. This is not to be taken in a non-metaphorical sense, because it doesn't have a non-metaphorical sense, any more than 'time flows' does. What is meant by this picturesque language is: *events happen*, and, as time goes on, the stock of events that have happened increases.

[3] No effort will be made to be consistent in terminology as to whether 'objects' are temporally extended or not—that is, whether what exists on a hypersurface is the entire object at some stage of its history, or a three-dimensional slice of a four-dimensional object. This is a terminological matter of no physical or ontological significance.

[4] It might be thought that there is an intermediate notion, relative to a world-line. Clifton and Hogarth ([1995]) have disposed of such a notion.

[5] They don't, because they could only exist for observers undergoing inertial motion. No known observer satisfies this condition, and it is likely that no observer anywhere does.

[6] This isn't quite how Albert expresses the 'metaphysical incompatibility' of quantum mechanics with relativity. Instead, he points out that wave-functions are, in general, functions not on 4-dimensional spacetime but instead on a higher-dimensional configuration space. But it is entanglement that is responsible for the fact that such wave-functions are not reducible to catalogues of local properties in 4-D spacetime, and hence the diagnosis given here does not differ in substance from Albert's.

[7] 'Local' is here being used in its usual sense, in that a local observable is one that can be measured by local observations. A global property, such as the center of mass of a mass distribution on some spacelike hypersurface σ, can be *located* at a point on σ without being local to any bounded spacetime region containing that point. Fleming and Butterfield ([2000]) argue that Newton-Wigner position observables are to be construed in this way; they are *located* without being locally measurable.

[8] The condition that σ' not lie to the past of σ is included in order to leave room for collapse theories with irreversible dynamics. For such a theory, past states might not be obtainable from future states. It may seem that the local evolution condition imposes no constraints on the relations between the state on a hypersurface σ and another hypersurface σ', part of which is in the past and part of which is in the future of σ. This is not the case. If the system whose state we are interested in lies in a bounded region on both σ and σ', then the state of the system on both hypersurfaces must be obtainable from the state on some hypersurface in the common past of these bounded regions of σ and σ'. See theorem and corollary in §4, below.

[9] This is perhaps obscured in the usual presentations, which are often given in terms of the Heisenberg picture, and, frequently, in terms of nonlocal observables. Translated into the Schrödinger picture, one obtains states on hypersurfaces evolving according to the Tomonaga-Schwinger equation (Tomonaga [1946]; Schwinger [1948]; see Schweber 1961, 419–421, for a discussion), which makes the satisfaction of the local evolution condition manifest.

[10] The following example is indebted to an example given by Clifton and Halvorson ([2001]), 4–5.

[11] This term was introduced in Myrvold ([2002]). There, however, it is erroneously stated that the relativity of entanglement is a consequence of introducing collapse into a relativistic context. As this example makes clear, there is a sense in which it occurs even without collapse.

Clifton and Halvorson ([2001]) have pointed out that, since the axioms of relativistic algebraic quantum field theory entail that entanglement is generic, in such a context no physically reasonable operation can be *guaranteed* to disentangle a state. One should, therefore, not only consider not only whether or not a system is entangled with another, but also the degree of entanglement. The relativity of entanglement will hold for any reasonable measure of degree of entanglement: one and the same stage of *A*'s history may be strongly entangled with one stage of *B*'s history and only weakly entangled with another.

[12] In so saying, we are excluding the Everett interpretation and its kin, which take the quantum state description to be complete, and the unitary quantum evolution of the state to be exact. Of this tradition, Albert says that it '(alas, for a whole bunch of reasons) seems to me not to be a particular candidate either' and in a footnote adds, 'Foremost among these reasons is that the many-worlds interpretations seems [seem?] to me not to be able to account for the facts about chance. But that's a long story, and one that's been told often enough elsewhere' (Albert [2000], 3). Albert is right about this. See Barrett ([1999]) for a careful discussion of this issue.

[13] The classic example of a collapse theory, and the one discussed by Albert, is the GRW theory (Ghirardi, Rimini, and Weber [1986]). See Pearle ([1997], [1999]) for accounts of subsequent work, including progress towards a relativistic collapse theory.

[14] It will be shown in a separate paper that the relativistic generalization of the GRWP continuous spontaneous localization model (CSL) does indeed satisfy these conditions.

[15] See also Myrvold ([2002]).

[16] There is a potential problem here: typically the space of possible complete histories of this sort will be uncountable; hence, it is not possible (on standard probability measures) for every complete history to be assigned a nonzero probability, and this poses a problem for the usual definition of conditional probability. Lewis suggests that this is to be circumvented by having the probability function take on values in an appropriate extension of the reals that includes infinitesimals. We will assume, with Lewis, that the requisite conditional probabilities do, in fact exist. This seems to be easier to do in a quantum context than in a classical context.

[17] Reichenbach proposes to *define* the past direction of time as the direction in which common causes are to be found for every correlation. If so, then the common cause is in the past by definition. Even if this proposal is not adopted, however, the usual construal of the principle is that common causes lie in the past.

[18] This is independent of the issue of whether there is an appropriately weakened version of the common cause principle that quantum probabilities do satisfy. See Hofer-Szabó *et al*. ([1999]), Rédei and Summers ([2001]).

[19] This gives an answer to Tim Maudlin's question, 'Does it really make sense to say . . . that a detection event is stochastic when one assigns it to one hyperplane but deterministic when assigned to another?' (Maudlin [1994], 209) that is more complete than the answer given in Myrvold ([2002]). It does make sense, because one set of probabilities is conditional on the outcome of a spacelike separated measurement, and the other is not.

[20] Ghirardi and Pearle ([1991]) discuss both our (i) and (ii *b*), but reserve the term 'objective property' for the latter.

[21] More needs to be said about precisely how such an approximate localization is to be realized, in terms of configurations of quantum fields. This discussion is intended to be neutral with respect to the detailed workings-out of such a scheme.

References

Aharonov, Yakir, and David Albert [1984]: 'Is the Usual Notion of Time Evolution Adequate for Quantum-Mechanical Systems? II. Relativistic Considerations.' *Physical Review D* **29**, pp. 228–234.

Albert, David Z. [2000]: 'Special Relativity as an Open Question', in H.-P. Breuer and F. Petruccione (*eds*), *Relativistic Quantum Measurement and Decoherence: Lectures of a workshop, held at the Istituto Italiano per gli Studi Filosofici, Naples, April 9-10, 1999*, Berlin: Springer, pp. 1–13.

Barrett, Jeffrey A. [1999]: *The Quantum Mechanics of Minds and Worlds*. Oxford: Oxford University Press.

Bell, J. S. [1987]: 'Are There Quantum Jumps?', in *Speakable and Unspeakable in Quantum Mechanics*, Cambridge: Cambridge University Press, pp. 201–212.

Butterfield, Jeremy [1989]: 'A Space-Time Approach to the Bell Inequality', in Cushing and McMullin (*eds*), pp. 114–144.

Butterfield, J, G.N. Fleming, G.C. Ghirardi and R. Grassi [1993]: 'Parameter Dependence in Dynamical Models for Statevector Reduction', *International Journal of Theoretical Physics* **32**, pp. 2287–2304.

Callender, Craig [2000]: ‘Shedding Light on Time’. *Philosophy of Science* **67** (Proceedings), pp. S587–S599.

Clifton, Robert, and Hans Halvorson [2001]: ‘Entanglement and Open Systems in Algebraic Quantum Field Theory’. *Studies in the History and Philosophy of Modern Physics* **32**, pp. 1–31.

Clifton, Rob, and Mark Hogarth [1995]: ‘The Definability of Objective Becoming in Minkowski Spacetime’. *Synthese* **103**, pp. 355–387.

Clifton, Robert, Constantine Pagonis, and Itamar Pitowsky [1992]: ‘Relativity, Quantum Mechanics, and EPR’, in D. Hull, M. Forbes, and K. Okruhlik (*eds*), 1992, *PSA 1992*, Vol. 1, East Lansing, MI: Philosophy of Science Association, pp. 114-128.

Clifton, R. (*ed.*) [1996]: *Perspectives on Quantum Reality: Non-Relativistic, Relativistic, and Field-Theoretic*, Dordrecht: Kluwer Academic Publishers.

Cushing, James T., and Ernan McMullin (*eds*) [1989]: *Philosophical Consequences of Quantum Theory: Reflections on Bell's Theorem*, Notre Dame, IN: University of Notre Dame Press.

Dickson, Michael, and Rob Clifton [1998]: ‘Lorentz-Invariance in Modal Interpretations’, in Dieks and Vermaas (*eds*), 1989, pp. 9–47.

Dieks, Dennis [1998]: 'Locality and Lorentz-Covariance in the Modal Interpretation of Quantum Mechanics', in Dieks and Vermaas (*eds*), 1989, pp. 49–67.

Dieks, Dennis, and Pieter E. Vermaas (*eds*) [1998]: *The Modal Interpretation of Quantum Mechanics*. Dordrecht: Kluwer Academic Publishers.

Dorato, Mauro [1995]: *Time and Reality: Spacetime Physics and the Objectivity of Temporal Becoming*. Bologna: CLUEB.

—— [1996]: 'On Becoming, Relativity, and Nonseparability', *Philosophy of Science* **64**, pp. 585–604.

Eberhard, P.H. [1978]: 'Bell's Theorem and the Different Concepts of Locality'. *Il Nuovo Cimento* **46B**, pp. 392–419.

Einstein, Albert [1948]: 'Quanten-Mechanik und Wirklichkeit', *Dialectica* **2**, pp. 320–23.

Fleming, G. N. [1986]: 'On a Lorentz Invariant Quantum Theory of Measurement', in Daniel M. Greenberger (*ed*), 1986, *New Techniques and Ideas in Quantum Measurement Theory* (*Annals of the New York Academy of Sciences* **480**), pp. 574–575.

—— [1989]: ‘Lorentz Invariant State Reduction, and Localization’, in A. Fine and M. Forbes, (*eds*), *PSA 1988*, Vol. 2, East Lansing, MI: Philosophy of Science Association, pp. 112–126.

—— [1996]: ‘Just How Radical is Hyperplane Dependence?’, in R. Clifton (*ed*), 1996, pp. 11–28.

Fleming, G. N., and J. Butterfield [2000]: ‘Strange Positions’, in Jeremy Butterfield and Constantine Pagonis (*ed*), 2000, *From Physics to Philosophy*, Cambridge: Cambridge University Press, pp. 108–165.

Ghirardi, GianCarlo [2000]: ‘Local Measurements of Nonlocal Observables and the Relativistic Reduction Process’, *Foundations of Physics* **30**, pp. 1337–1385.

Ghirardi, GianCarlo, and Renata Grassi [1994]: ‘Outcome Predictions and Property Attribution: the EPR Argument Reconsidered’, *Studies in History and Philosophy of Science* **25**, pp. 397–423.

Ghirardi, G.C., R. Grassi, J. Butterfield, and G.N. Fleming [1993]: ‘Parameter Dependence and Outcome Dependence in Dynamical Models for State Vector Reduction’, *Foundations of Physics* **23**, pp. 341–364.

Ghirardi, GianCarlo, and Philip Pearle [1991]: 'Elements of Physical Reality, Nonlocality and Stochasticity in Relativistic Dynamical Reduction Models', in Arthur Fine, Micky Forbes, and Linda Wessels (*eds*), 1991, *PSA 1990*, Vol. 2, East Lansing, MI: Philosophy of Science Association, pp. 35–47.

Ghirardi, G. C., A. Rimini, and T. Weber [1980]: 'A General Argument against Superluminal Transmission through the Quantum Mechanical Measurement Process', *Lettere al Nuovo Cimento* **27**, pp. 294–98.

——— [1986]: 'Unified Dynamics of Microscopic and Macroscopic Systems', *Physical Review D* **34**, pp. 470–91.

Haag, Rudolf [1992]: *Local Particle Physics: Fields, Particles, Algebras*, Berlin: Springer-Verlag.

Halvorson, Hans, and Rob Clifton [2002]: 'No Place for Particles in Relativistic Quantum Theory?', *Philosophy of Science* **69**.

Hofer-Szabó, Gábor, Miklós Rédei, and László E. Szabó [1999]: 'On Reichenbach's Common Cause Principle and Reichenbach's Notion of Common Cause', *British Journal for the Philosophy of Science* **50**, pp. 377–399.

Howard, Don [1985]: 'Einstein on Locality and Separability', *Studies in History and Philosophy of Science* **16**, pp. 171–201.

—— [1989]: 'Holism, Separability, and the Metaphysical Implications of the Bell Experiments', in Cushing and McMullin (*eds*), 1989, pp. 224–253.

Kraus, Karl [1983]: *States, Effects, and Operators: Fundamental Notions of Quantum Theory*. (*Lecture Notes in Physics* **190**), Berlin: Springer-Verlag.

Lewis, David [1980]: 'A Subjectivist's Guide to Objective Chance', in Richard C. Jeffrey (*ed.*), 1980, *Studies in Inductive Logic and Probability, Vol. II*, Los Angeles: University of California Press, pp. 263–293.

Malament, David B. [1996]: 'In Defense Of Dogma: Why There Cannot Be A Relativistic Quantum Mechanics Of (Localizable) Particles', in Clifton (*ed*), 1996, pp. 1–10.

Maudlin, Tim [1994]: *Quantum Non-Locality and Relativity: Metaphysical Intimations of Modern Physics*, Cambridge, MA: Blackwell Publishers.

Maxwell, Nicholas [1985]: 'Are Probabilism and Special Relativity Incompatible?', *Philosophy of Science* **52**, pp. 23–43.

McTaggart, J.M.E. [1908]: 'The Unreality of Time', *Mind* **68**, pp. 457–474.

Myrvold, Wayne C. [2002]: 'On Peaceful Coexistence: Is the Collapse Postulate Incompatible with Relativity?', forthcoming in *Studies in History and Philosophy of Modern Physics*.

Pearle, Philip [1997]: 'Tales and Tails and Stuff and Nonsense', in Robert S. Cohen, Michael Horne, and John Stachel (*eds*), 1997, *Potentiality, Entanglement, and Passion-at-a-Distance: Quantum Mechanical Studies for Abner Shimony, Volume One*, Dordrecht: Kluwer Academic Publishers, pp. 143–156.

—— [1999]: 'Collapse Models', in Heinz-Peter Breuer and Francesco Petruccione (*eds*), *Open Systems and Measurement in Relativistic Quantum Theory: Proceedings of the workshop held at the Istituto Italiano per gli Studi Filosofici Naples, April 3–4, 1998*, Berlin: Springer.

Rédei, Miklós, and Stephen J. Summers [2002]: 'Local Primitive Causality and the Common Cause Principle in Quantum Field Theory', forthcoming in *Foundations of Physics*. Available at PhilSci Archive (http://philsci-archive.pitt.edu/), PITT-PHIL-SCI00000352.

Reichenbach, Hans [1956]: *The Direction of Time*, Maria Reichenbach (*ed*), Berkeley and Los Angeles: University of California Press.

Schrödinger, E. [1935]: 'Discussion of Probability Relations Between Spatially Separated Systems', *Proceedings of the Cambridge Philosophical Society* **31**, pp. 555–563.

Schweber, Silvan S. [1961]: *An Introduction to Relativistic Quantum Field Theory*, Evanston, IL: Row, Peterson and Company.

Schwinger, Julian [1948]: 'Quantum Electrodynamics I. A Covariant Formulation', *Physical Review* **74**, 1439–1461.

Stein, Howard [1968]: 'On Einstein-Minkowski Space-Time', *The Journal of Philosophy* **65**, pp. 5–23.

—— [1970]: 'A Note on Time and Relativity Theory', *The Journal of Philosophy* **67**, pp. 289–294.

—— [1991]: 'On Relativity Theory and Openness of the Future', *Philosophy of Science* **58**, pp. 147–167.

Tittel, W., J. Brendel, H. Zbinden, and N. Gisin [1998]: 'Violation of Bell Inequalities by Photons More Than 10 km Apart', *Physical Review Letters* **81**, pp. 3563–3566.

Tomonaga, S. [1946]: 'On a Relativistically Invariant Formulation of the Quantum Theory of Wave Fields', *Progress of Theoretical Physics* **1**, pp. 1–13. Reprinted in J. Schwinger, (*ed*), 1958, *Selected Papers on Quantum Electrodynamics*, New York: Dover Publications, pp. 156–68.

van Fraassen, Bas C. [1982]: ‘The Charybdis of Realism: Epistemological Implications of Bell’s Inequality’, *Synthese* **52**, pp. 25–38. Reprinted, with additions, in Cushing and McMullin (*eds*), 1989, pp. 97–113.