Collapse Theories

Peter J. Lewis

The collapse postulate in quantum mechanics is problematic. In standard presentations of the theory, the state of a system prior to a measurement is a sum of terms, with one term representing each possible outcome of the measurement. According to the collapse postulate, a measurement precipitates a discontinuous jump to one of these terms; the others disappear. The reason this is problematic is that there are good reasons to think that measurement *per se* cannot initiate a new kind of physical process. This is the *measurement problem*, discussed in section 1 below.

The problem here lies not with *collapse*, but with the appeal to measurement. That is, a theory that could underwrite the collapse process just described without ineliminable reference to measurement would constitute a *solution* to the measurement problem. This is the strategy pursued by *dynamical* (or *spontaneous*) collapse theories, which differ from standard presentations in that they replace the measurement-based collapse postulate with a dynamical mechanism formulated in terms of universal physical laws. Various dynamical collapse theories of quantum mechanics have been proposed; they are discussed in section 2.

But dynamical collapse theories face a number of challenges. First, they make different empirical predictions from standard quantum mechanics, and hence are potentially empirically refutable. Of course, testability is a virtue, but since we have no empirical reason to think that systems ever undergo collapse, dynamical collapse theories are inherently empirically risky. Second, there are difficulties reconciling the dynamical collapse mechanism with special relativity. Third, the post-collapse state is not the same as the post-measurement state of standard quantum mechanics, raising the possibility that dynamical collapse theories do not solve the measurement problem after all. These challenges are described in sections 3, 4 and 5, respectively.

Assuming that these challenges can be met, dynamical collapse theories can lay claim to being serious contenders for the correct description of the quantum world. And the description they provide has a number of interesting consequences. First, it makes indeterminism an irreducible fact about the physical world. It has been argued that this has important consequences for the foundations of statistical mechanics, for free will, and for consciousness. These claims are discussed in section 6.

Second, many dynamical collapse theories imply that the quantum wave function is fundamental, and that particles are just a temporary "bunching up" of the fundamental wave-like entity. This understanding of the ontology of the quantum world gives rise to a new kind of vagueness, since a wave can be fuzzy around the edges in a way that a particle cannot. Furthermore, the quantum wave function is defined over a high-dimensional space, not the three-dimensional space of experience, suggesting to some that three-dimensionality is an illusion. Dynamical collapse theories share these implications with other "wave function only" theories, such as Everettian approaches, and share some of them with wave function realist versions of Bohmian approaches. These implications of dynamical collapse theories are considered in section 7.

1. The Measurement Problem

Quantum mechanics represents the state of a system in various ways, but the representation that is most perspicuous for understanding collapse theories uses a wave function. For a single particle, the wave function is a complex-valued function of three spatial dimensions and time: $\psi(x, y, z, t)$. The wave function changes over time according to a linear differential equation, the Schrödinger equation.

To begin with, let us stipulate that a particle is located in a particular spatial region if and only if all the corresponding wave function amplitude is contained in that region. (We will have reason to relax this stipulation later.) So, for example, the wave function shown schematically (in

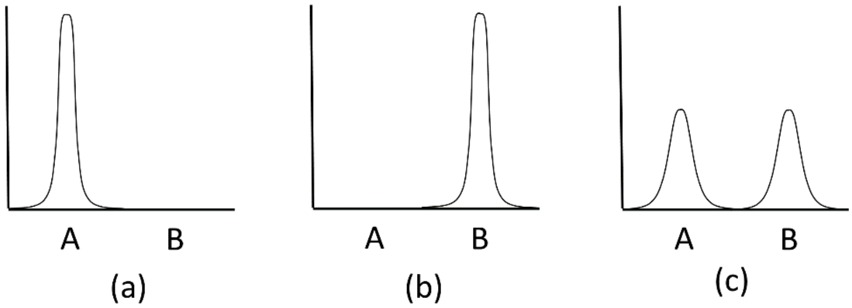

Figure 1: Three quantum states

one dimension) in figure 1(a) represents a particle in region A, and the one in figure 1(b) represents a particle in a distinct region B.

The trouble with quantum mechanics starts from the *superposition principle*, which says that the sum of any two quantum states is also a quantum state. That is, if we take the function in figure 1(a) and add it to the function in figure 1(b), we obtain another possible state of the particle, shown in figure 1(c). (We also need to rescale the function, due to the connection between the area under the curve and probability, to be explained shortly.)

States like 1(c) are crucial to quantum mechanical explanations. But according to our earlier stipulation, the wave function in figure 1(c) is not a state in which the particle is in region A, and it is not a state in which the particle is in region B. Nevertheless, when the position of a particle in such a state is measured, it is always found in one location or the other. Why should a state in which the particle is not in A and not in B generate a measurement result in which the particle is found in one of these locations?

The founders of quantum mechanics were aware of this problem. Born (1926) noted that the wave function can be used to generate the *probabilities* of the two outcomes: the square of the area under the wave function in region A is the probability that the particle will be found in region A, and similarly for region B. This is the *Born rule*. But the Born rule just quantifies the problem: if the particle is not in region A, why is there a 50% chance of finding it there?

The standard response, codified by von Neumann (1932), was to postulate that there are *two* dynamical processes by which the wave function changes over time. Between measurements, the wave function evolves continuously according to the Schrödinger equation, but during a measurement, the wave function jumps discontinuously into a state corresponding to a determinate location, with probabilities given by the Born rule. The latter is the *collapse postulate*.

Applied to a particle in the state shown in figure 1(c), the collapse postulate says that if the position of the particle is measured, there is a 50% chance that it will collapse to the state shown in figure 1(a), and a 50% chance that it will collapse to the state shown in figure 1(b). Hence even though the pre-measurement state is not one in which the particle is in region A, and not one in which it is in region B, the collapse postulate together with the post-collapse state can explain our measurement results, as well as the sense in which these results reveal the actual (post-measurement) location of the particle.

As it stands, though, the collapse postulate is untenable. The continuous, linear Schrödinger evolution and the discontinuous, non-linear collapse process are incompatible: neither can be reduced to the other. So to be consistent, quantum mechanics must postulate a sharp division between those physical processes that count as measurement processes and those that do not. This seems like a tall order. What's more, since measuring devices are constructed out of particles that are not themselves being measured, the measuring device should operate according to the continuous Schrödinger evolution, and hence cannot instantiate the discontinuous collapse process. This is the much-discussed *measurement problem*.

It is worth noting, though, that the problem lies not with collapse *per se*, but with tying the collapse process to measurement. If a collapse process could be described that yields the same

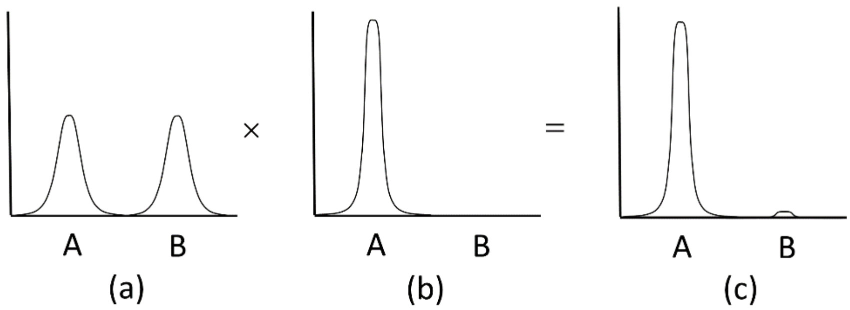

Figure 2: Collapse for a single particle

outcome without the appeal to measurement, it would not be subject to the same critique. This is the approach to the interpretation of quantum mechanics pursued by dynamical collapse theories.

2. Dynamical theories of collapse

The basic approach was first described by Pearle (1976): instead of one dynamical law describing measurements and another applying between measurements, a single dynamical law applies at all times, deviating just slightly from the Schrödinger equation. The first fully-developed theory along these lines was proposed by Ghirardi, Rimini and Weber (1986), and has become known as the GRW theory.

According to the GRW theory, the wave function for a single particle obeys a dynamical law that mostly coincides with the Schrödinger equation, except that there is a small chance per unit time that the wave function undergoes a spontaneous collapse process (or *hit*) in which it becomes highly localized around a point. More precisely, a hit multiplies the wave function by a narrow three-dimensional Gaussian (bell curve) in the coordinates of the particle concerned, where the width of the Gaussian is 10^{-5}cm.

The collapse process is indeterministic in two senses. First, whether a given particle undergoes a collapse is a random matter: there is a chance of 10^{-16} that any given particle will undergo a collapse in any given second. Second, if a particle undergoes a collapse, the location of the hit is random, with probabilities chosen so as to recover the Born rule: the chance that the Gaussian is centred in a particular region is given by the integral over that region of the pre-collapse wave function multiplied by the Gaussian.

For a single particle in the superposition state of figure 1(c), the effect of a hit is shown schematically (in one dimension) in figure 2. In this state, there is a 50% chance of the hit being centred in region A, and a 50% chance of it being centred in region B. If it is centred in region A, then the effect is as shown: almost all the wave amplitude is now in region A, although a tiny amount (not shown to scale!) remains in region B, due to the fact that the "tails" of the Gaussian extend to infinity. It looks like this ought to be close enough for the particle to count as being in region A after the hit (although this question is addressed further in section 5).

For a single particle, a collapse is extremely rare—about one every 100 million years. This is just as well, because experimental physics shows that individual particles (and small collections of them) always obey the Schrödinger equation very closely. But for a macroscopic object, the collapse rate can be appreciable: for an object containing of the order of 10^{23} particles, there will be around 10 million collapse events per second. And if the positions of the particles are strongly correlated with each other, as is the case for a solid object, the collapse of a single particle is sufficient to localize all the particles.

This is shown schematically for two particles in figure 3. The wave function for N particles is defined over a *configuration space*—a 3N-dimensional space of *configurations* of particles. For two particles, then, the wave function is a function of six spatial coordinates, three for each particle, plus time: $\psi(x_1, y_1, z_1, x_2, y_2, z_2, t)$. Figure 3 represents two of those dimensions: the horizontal axis

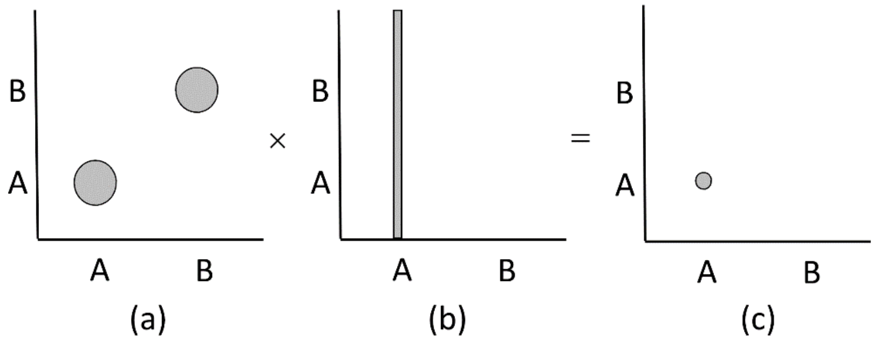

Figure 3: GRW collapse for two correlated particles

represents one of the spatial dimensions of particle 1, and the vertical axis represents one of the spatial dimensions of particle 2. The shaded areas represent regions where the wave function amplitude is high.

For two particles whose positions are strongly correlated, a typical superposition state is as shown in figure 3(a): all the wave amplitude is concentrated in areas representing the particles as occupying the same spatial region. If particle 1 undergoes a hit, the wave function is multiplied by a function that is a Gaussian in the coordinates of particle 1 and a constant in the coordinates of particle 2, so it is large only in the "stripe" shown in figure 3(b). The result is that the post-hit wave function is large only in region A for *both* particles. The same goes for a hit on particle 2. That is, if either particle undergoes a collapse, both particles acquire locations.

Of course, a collapse for two particles is still exceedingly rare. But for a macroscopic solid object, a hit on one of the particles making up the object will occur on average every tenth of a microsecond. Furthermore, because of the forces binding the particles in a solid object together, the wave function will only be large in regions of configuration space in which the particles are close together. Due to these strong correlations between the positions of the particles making up the object, a collapse for one particle is sufficient to localize the whole object.

This is the heart of the dynamical collapse solution to the measurement problem. Suppose we start with a single particle in the superposition state of figure 1(c) and we measure its position. Measuring its position requires us to correlate its position with something we can see, such as a pointer on a dial. But in doing that, we create a macroscopic object in a superposition of two distinct locations, pointing at "A" on the dial and pointing at "B". The GRW collapse process very rapidly reduces this superposition to one location or the other, and since the particle is correlated with the pointer, the particle too acquires a determinate location. Hence after the measurement, the particle is either in region A or in region B, and the pointer is pointing to the corresponding measurement result.

Note that there is no essential appeal to measurement in this account. When we correlate a superposition state of a microscopic system with the position of a macroscopic solid object, the GRW collapse process reduces the superposition state to one of its components, and it does so according to a single, precisely-specified dynamical law.

To accommodate symmetrisation requirements for identical particles, more recent dynamical collapse theories incorporate the collapse process as a non-linear correction to the Schrödinger equation, rather than as a distinct process, producing a collapse that is continuous rather than discrete (Ghirardi, Pearle and Rimini 1990). Diosi (1989) and Penrose (1996) further suggest that the collapse mechanism may be connected to the role of gravity. Penrose notes that the existence of a macroscopic object in a superposition of two distinct locations entails, via gravity, a superposition of two distinct space-time structures, and speculates that the latter superposition may be inherently unstable. At present, though, gravitational explanations of the collapse process have no empirical support. Indeed, there is no evidence for the existence of a collapse process at all—unless one takes the existence of definite experimental outcomes as evidence for collapse. One might take the lack of empirical support as an objection to the entire dynamical collapse project, so let us briefly consider the experimental situation.

3. Tests

Dynamical collapse theories make different empirical predictions from no-collapse interpretations of quantum mechanics, such as de Broglie-Bohm theory (Tumulka, this volume) and Everettian quantum mechanics (Saunders, this volume). So can't we just perform the relevant experiments to see which is right? Unfortunately, this is not at all straightforward.

There is no problem in principle. Dynamical collapse theories predict that a superposition of distinct locations for a macroscopic object is inherently unstable, and will rapidly evolve to some determinate location. No-collapse theories predict that such macroscopic superpositions are stable. And there are experiments that could in principle distinguish between an object in a superposition of locations and an object in one of those locations.

But in practice, these experiments are impossible to perform. Note first that a simple measurement of the position of the object won't suffice to distinguish the superposition from the determinate location. That is because both collapse and no-collapse interpretations have mechanisms to ensure that a measurement on a superposition state yields a determinate outcome, just as a measurement on a determinate location state does.

So we need something more subtle. What is required is an interference experiment, in which the two terms in the superposition are made to interact with each other, indicating that they are both present. Interference experiments are very sensitive to environmental effects. If a single outside particle becomes correlated with one term in the superposition but not the other, then the interference experiment fails. So to exhibit interference effects, a system has to be kept completely isolated from the environment. This is possible for microscopic systems, but is practically impossible for macroscopic systems.

Nevertheless, there has been considerable progress in demonstrating interference with larger and larger systems. Early attempts to detect quantum collapse centered on the behavior of superconducting quantum interference devices (SQUIDs). Using such devices, experimentalists are able to create and detect a superposition of a clockwise and a counter-clockwise electric current in a ring of approximately 1cm diameter. Such devices do not reveal any collapse effects. However, the superposition in a SQUID, though macroscopic, is of distinct currents, not of distinct locations. The wave function for the electrons involved is distributed over the whole ring whether the current is clockwise or counter-clockwise. So a GRW collapse to a precise location is not a collapse to one *current* over the other, and GRW collapses would not be expected to have a measurable effect on the current (Rae 1990).

Another important experimental development involves demonstrating interference effects for larger and larger molecules. Interference can now be demonstrated for relatively large organic molecules, such as $C_{48}H_{26}F_{24}N_8O_8$ (Juffmann et al. 2012). This shows that such molecules can exist in superpositions of distinct location states, at least for the short amount of time it takes to traverse the apparatus. But even such a large molecule involves fewer than 5000 fundamental particles, inducing a GRW collapse rate of 5000×10^{-16} per second, or about one every hundred thousand years—still far too rare to be detectable.

Further methods for detecting quantum collapses have been attempted (Leggett 2002; Bassi et al. 2013). So far we have no empirical evidence to suggest that quantum collapses occur, but neither are the particular models of the GRW theory and its continuous variants ruled out. So empirical tests are inconclusive. However, there is an indirect empirical argument against dynamical collapse theories, namely that they conflict with another well-confirmed theory: special relativity.

4. Relativity

The original GRW theory conflicts with special relativity in two distinct ways. First, when a particle undergoes a hit, the probability distribution for the centre of the collapse is based on the wave function distribution *at that time*. But according to special relativity, there is no absolute standard of simultaneity, so the state of a spread-out entity *at a time* is ill-defined. Second, the hit instantaneously multiplies the wave function over the whole of space by a Gaussian, and again this process is ill-defined according to special relativity. In particular, for correlated particles a hit on one particle has an instantaneous effect on the state of another particle, no matter how far apart, and instantaneous action at a distance is prima facie incompatible with special relativity.

However, progress has been made in constructing a dynamical collapse theory that is consistent with special relativity (see Myrvold, this volume, for further details). Tumulka (2006), following some remarks by Bell (1987, 205), suggests that the ontology of GRW-type theories should be understood as point-like: the spatio-temporal point at the centre of a hit event is a "flash" of reality, and the spread-out wave function can be interpreted instrumentally as governing the probability distribution for these flashes. So for a single particle system, what exists is a set of point-like events, about one per hundred million years. Since the wave function is treated instrumentally, there is no real collapse of the wave function that could conflict with relativity. Furthermore, given a flash event at a particular space-time point, the probability distribution for the location of the next flash event is defined over surfaces that are relativistically invariant—that is, over surfaces such that $x^2 + y^2 + z^2 - c^2t^2$ is a constant, rather than simultaneity surfaces for which t is a constant.

Alternatively, Ghirardi, Grassi, and Benatti (1995) suggest that the appropriate ontology for GRW-type theories is a mass density distribution defined over three-dimensional space. In relativistic versions, the mass density in a small region of space is determined by the state in the past light cone of that region (Bedingham et al. 2014). Hence a hit centred on a space-time point produces high mass density at that point, and a surrounding region of near-zero mass density that spreads outwards at the speed of light as the collapse event enters the past light cone of surrounding points (Myrvold 2016).

Hence for a single particle, versions of the GRW theory can be made consistent with special relativity. But for two or more particles, there is still the worry that for correlated particles, a hit on one particle can instantly affect the other, no matter how distant. Tumulka (2006, 350) suggests that this can be accommodated within special relativity by allowing that the direction of the influence is indeterminate. Consider two particles whose perfectly-correlated *z*-coordinates are measured at space-like separated *x*-coordinates. In some frames of reference, the measurement of particle 1 occurs first, and causes particle 2 to acquire a determinate *z*-coordinate; in other frames of reference, the measurement of particle 2 occurs first, and causes particle 1 to acquire a determinate *z*-coordinate. Neither causal story is to be preferred, so no absolute standard of simultaneity is required.

Even so, both causal stories require faster-than-light causation. However, this is not in direct conflict with special relativity; rather, it just requires that the probabilities ascribed to the measurement outcomes depend on the frame of reference. In the frame in which particle 1 is measured first, one particular measurement outcome for particle 2 has probability 1 prior to measurement, whereas in the frame in which particle 2 is measured first, no measurement outcome for particle 2 has probability 1 prior to measurement (Myrvold 2002, 461). Indeed, Myrvold (2016) prefers to say that there are non-local probabilistic correlations here, but no superluminal causation, on the grounds that the direction of causal relations cannot be indeterminate.

5. Tails

So consistency with special relativity is a problem for dynamical collapse theories, but perhaps not an insuperable one. However, solving this problem is moot if, as some have claimed, the dynamical collapse approach does not even minimally solve the measurement problem.

The source of this concern is that, as mentioned above, for a particle in the superposition state 2(a), a collapse does not ensure that *all* the post-collapse wave function amplitude is in region A (or in region B), just that *most* of it is. The same goes for a macroscopic object. If a human being, ends up in a superposition of occupying two distinct regions, for example as a result of correlating her location with the position of a particle in a superposition state, then collapses will very rapidly put her state *close* to one in which all her wave function amplitude is located in one of those regions. However, there remains a small but non-zero amplitude in the other location.

One form of the concern goes as follows. In the example above, the small term in the post-collapse state has exactly the same internal structure as the big term: it is the structure of a human being. Granted a rather plausible functionalism, it is the structure of a term, not its amplitude, determines what it represents (Wallace 2003). But then there is a human being in both locations, and dynamical collapse theories are ineffective at bringing about determinate measurement outcomes (Cordero 1999).

Since the source of this problem is the non-zero "tails" of the Gaussian collapse function stretching to infinity in every direction, the most obvious solution is to eliminate the tails (Wallace 2014). That is, is the collapse function were strictly *zero* at distances greater than 10^{-5}cm from the collapse centre, then after a collapse there would be a wave term with the structure of a human being at only one location.

But there is another form of the concern about tails that is resistant to this solution. Between collapses, the state of a particle evolves according to the Schrödinger equation, and this means that even if collapse makes the wave function strictly zero outside a given region at a time, it has tails extending to infinity an instant later. So if what it is for a particle to be located in a region is that *all* its wave amplitude is located in that region (as we assumed earlier), then after a collapse the particle is still not located in any finite region of space, and dynamical collapse theories do not solve the measurement problem after all (Albert and Loewer 1990).

One might take this as an additional motivation to adopt a "flashy" version of GRW, so that the ontology resides only at the precise instants of collapse. Albert and Loewer (1996) instead suggest solving this problem by relaxing the link between wave function amplitude and particle location: instead of demanding that *all* the wave amplitude be contained in a region for the particle to count as being located there, we only require that *almost* all of it is so located. This works, but at the cost of introducing a new kind of vagueness: there is presumably no fact of the matter about precisely how much of the amplitude needs to be in the region for the particle to count as located there. Whether this vagueness is problematic is considered in section 7.

6. Chance

Quantum mechanics is often taken to be an indeterministic theory. But this is not a straightforward consequence of the theory: the de Broglie-Bohm and Everettian versions are deterministic at the fundamental level. Nevertheless, dynamical collapse theories really are irreducibly indeterministic: they incorporate objective chances into fundamental physical law (see Emery, this volume, and Suarez, this volume, for more on chance and determinism).

What consequences does this have? Albert (2000) suggests that the role of chance in dynamical collapse theories can solve a problem in the foundations of statistical mechanics. It is well known that for a given macrostate, the microstates exhibiting normal thermodynamic behaviour vastly outnumber those exhibiting abnormal behaviour, but it is far from clear how this asymmetry can be used to *explain* thermodynamic behaviour (see Shahvisi, this volume, for further discussion). Albert notes that if there were some mechanism by which the states of systems were randomly and asymmetrically perturbed at the microscopic level, then normal thermodynamic behaviour would be straightforwardly explicable via this asymmetry. Dynamical collapse theories entail this random, asymmetric perturbation, but other interpretations of quantum mechanics, notably de Broglie-Bohm and Everett, do not. Hence dynamical collapse theories might gain indirect support from the foundations of statistical mechanics (Shenker, this volume, discusses some further ways in which quantum mechanics may bear on the foundations of statistical mechanics).

More controversially, some see the indeterminism of dynamical collapse theories as opening the door for a reconciliation of free will with physics. The trouble with such suggestions is that the collapse process is random, and randomness looks no more hospitable to free will than determinism. Still, Kane (1996) suggests that genuine indeterminacy, even if it is random, is essential to free will. The idea is that an agent must be ultimately responsible for their character if they are to be truly free, and a collapse in the brain at a suitable juncture, even if it is random, might be enough to secure ultimate responsibility.

Along similar lines, some have seen a connection between collapse and consciousness. Prior to the advent of dynamical collapse theories, Wigner (1961) famously speculated that consciousness might be required to explain the collapse postulate. More recently, Hameroff and Penrose (1996) suggest that dynamical collapses in the brain might explain consciousness, insofar as consciousness requires non-computability, and quantum collapse can introduce non-computability. It is worth noting, though, that these proposals concerning free will and consciousness add various highly contested philosophical claims to the already controversial status of collapse theories.

7. Ontology

If a dynamical collapse theory is true, what does this tell us about the furniture of the world? The clearest consequence is that particles are not fundamental. The fundamental law of a dynamical collapse theory governs the evolution of a wave function, and the wave function alone underlies all our empirical observations. After a collapse, the wave function becomes highly localized in three of its dimensions, and while that localization persists we can speak of a particle occupying a determinate position in three-dimensional space. But this "particle" is just a manner of speaking about the wave function.

However, there are several reasons to doubt that the ontology of dynamical collapse theories consists of the wave function and nothing but the wave function. First, as noted in section 4, some attempts to reconcile dynamical collapse theories with special relativity replace the wave function with discrete, point-like events or with a mass density distribution over three-dimensional space. However, special relativity arguably does not rule out the view that the wave function is fundamental, provided one is willing to ascribe a distinct wave function to every spacelike hypersurface (Myrvold 2002).

Second, the wave function for a system of N particles is defined over a 3N-dimensional configuration space. Hence it seems that if the wave function is fundamental, then the appearance of the world as three-dimensional is somehow illusory (Albert 1996). To avoid this conclusion, several commentators have again postulated that the fundamental ontology of dynamical collapse theories includes a mass density distribution over three-dimensional space whose value in a region can be derived from the wave function (Allori 2013). The issue of dimensionality can also be used to motivate a flash ontology, since flashes are defined at points of three-dimensional space. However, it is also possible to argue that the wave function indirectly describes objects in three-dimensional space, bypassing these concerns (Lewis 2013).

Third, some have argued that the wave function needs to be supplemented to counteract the effects of the novel quantum vagueness introduced by dynamical collapse theories (section 5). This vagueness allows a particle to count as occupying a region when almost all its wave amplitude is in that region. It follows from this and the high-dimensional nature of the wave function that for a set of particles, each can individually count as occupying a given region, even though the set taken as a whole does not count as occupying the region. The same goes for macroscopic objects (Lewis 1997). If objects are constituted by a mass density distribution rather than by a wave amplitude distribution, then this problem does not arise (Bassi and Ghirardi 1999). However, it is not clear that there is really a problem here that needs solution: the strange properties of compound objects can be regarded as an inevitable consequence of the mismatch between our classical concepts and the quantum world (Lewis 2003).

The ontological consequences of dynamical collapse theories remain an area of active debate.[1]

References

Albert, David Z. (1996), "Elementary quantum metaphysics," in J. Cushing, A. Fine and S. Goldstein (eds.), *Bohmian Mechanics and Quantum Theory: An Appraisal*. Dordrecht: Kluwer, 277-284.

Albert, David Z. (2000), *Time and Chance*. Cambridge: Harvard University Press.

Albert, David Z., and Barry Loewer (1990), "Wanted dead or alive: two attempts to solve Schrödinger's paradox," in A. Fine, M. Forbes and L. Wessels (eds.), *PSA 1990, Volume 1*. Chicago: University of Chicago Press, 277–285.

Albert, David Z., and Barry Loewer (1996), "Tails of Schrödinger's cat," in R. Clifton (ed.), *Perspectives on quantum reality*. Dordrecht: Kluwer, 81–92.

Allori, Valia (2013), "Primitive ontology and the structure of fundamental physical theories," in Ney and Albert (2013), 58–75.

Bassi, Angelo, and Gian Carlo Ghirardi (1999), "More about dynamical reduction and the enumeration principle," *British Journal for the Philosophy of Science* 50: 719–734.

Bassi, Angelo, Kinjalk Lochan, Seema Satin, Tejinder P. Singh, and Hendrik Ulbricht (2013), "Models of wave-function collapse, underlying theories, and experimental tests," *Reviews of Modern Physics* 85: 471-527.

Bedingham, Daniel, Detlef Dürr, Gian Carlo Ghirardi, Sheldon Goldstein, Roderich Tumulka, and Nino Zanghì (2014), "Matter density and relativistic models of wave function Collapse," *Journal of Statistical Physics* 154: 623-631.

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

Born, Max (1926), "Quantenmechanik der Stoßvorgänge," *Zeitschrift für Physik* 38: 803–827.

Cordero, Alberto (1999), "Are GRW tails as bad as they say?" *Philosophy of Science* 66: S59-S71.

Diosi, L. (1989), "Models for universal reduction of macroscopic quantum fluctuations," *Physical Review A* 40: 1165–1174.

Emery, Nina (this volume).

Ghirardi, Gian Carlo, Renata Grassi, and Fabio Benatti (1995), "Describing the macroscopic world: Closing the circle within the dynamical reduction program," *Foundations of Physics* 25: 5-38.

Ghirardi, Gian Carlo, Philip Pearle, and Alberto Rimini (1990), "Markov processes in Hilbert space and continuous spontaneous localization of systems of identical particles," *Physical Review A* 42: 78–89.

Ghirardi, Gian Carlo, Alberto Rimini, and Tullio Weber (1986), "Unified dynamics for microscopic and macroscopic systems," *Physical Review D* 34: 470–491.

Hameroff, Stuart R., and Roger Penrose (1996), "Conscious events as orchestrated space-time selections," *Journal of consciousness studies* 3: 36–53.

Juffmann, Thomas, Adriana Milic, Michael Müllneritsch, Peter Asenbaum, Alexander Tsukernik, Jens Tüxen, Marcel Mayor, Ori Cheshnovsky and Markus Arndt (2012), "Real-time single-molecule imaging of quantum interference," *Nature Nanotechnology* 7: 297–300.

Kane, Robert (1996), *The Significance of Free Will*. Oxford: Oxford University Press.

Leggett, Anthony J. (2002), "Testing the limits of quantum mechanics: motivation, state of play, prospects," *Journal of Physics: Condensed Matter* 14: R415–R451.

Lewis, Peter J. (1997), "Quantum mechanics, orthogonality, and counting," *British journal for the philosophy of science* 48: 313–328.

[1] I would like to thank Wayne Myrvold and Alastair Wilson for very helpful comments on an earlier draft.

Lewis, Peter J. (2003), “Four strategies for dealing with the counting anomaly in spontaneous collapse theories of quantum mechanics,” *International Studies in the Philosophy of Science* 17: 137–142.
Lewis, Peter J. (2013), “Dimension and Illusion”, in Ney and Albert (2013), 110–125.
Myrvold, Wayne C. (2002), “On peaceful coexistence: is the collapse postulate incompatible with relativity?” *Studies in History and Philosophy of Modern Physics* 33: 435–466.
Myrvold, Wayne C. (2016), “Lessons of Bell's Theorem: Nonlocality, yes; Action at a distance, not Necessarily,” in Shan Gao and Mary Bell (eds.), *Quantum Nonlocality and Reality – 50 Years of Bell's Theorem*. Cambridge. Cambridge University Press: 238-260.
Myrvold, Wayne C. (this volume).
Ney, Alyssa and David Z. Albert (eds.) (2013), *The Wave Function*. Oxford: Oxford University Press
Pearle, Philip (1976), “Reduction of the state vector by a nonlinear Schrödinger equation,” *Physical Review D* 13: 857–868.
Penrose, Roger (1996), “On gravity's role in quantum state reduction,” *General Relativity and Gravitation* 28: 581-600.
Rae, A. I. M. (1990), “Can GRW theory be tested by experiments on SQUIDS?,” *Journal of Physics A* 23: L57-L60.
Saunders, Simon (this volume), “Everettian Quantum Mechanics”.
Shahvisi, Arianne (this volume).
Shenker, Orly (this volume).
Suarez, Mauricio (this volume)
Tumulka, Roderich (2006), “Collapse and relativity,” in A. Bassi, D. Dürr, T. Weber and N. Zanghì (eds), *Quantum Mechanics: Are there Quantum Jumps?* and *On the Present Status of Quantum Mechanics*. *AIP Conference Proceedings* 844: 340–352.
Tumulka, Roderich (this volume), “de Broglie-Bohm Theory”.
Wallace, David (2003), “Everett and structure,” *Studies in History and Philosophy of Modern Physics* 34: 87-105.
Wallace, David (2014), “Life and death in the tails of the GRW wave function,” *arXiv:1407.4746*.
Wigner, Eugene P. (1961), “Remarks on the Mind-Body Question,” in I. J. Good (ed.) *The Scientist Speculates*. London: Heineman: 284-302.
von Neumann (1932), *Mathematische Grundlagen der Quantenmechanik*. Berlin: Springer.