

# On the Status of Primitive Ontology

Peter J. Lewis

February 2, 2017

## 1 Introduction

Spontaneous collapse theories provide a promising solution to the measurement problem. But they also introduce a number of problems of their own. First, the primary explanatory entity of a collapse theory—the wave function—inhabits a high-dimensional space, rather than the three-dimensional space of experience. Second, the continuity of the wave function introduces a new and potentially problematic form of vagueness when used to describe discrete physical systems such as particles or marbles. Third, the collapse of the wave function is hard to reconcile with special relativity.

In response to these problems, advocates of collapse theories have proposed various accounts of the *primitive ontology* of collapse theories—postulated underlying entities governed by the collapse theory and underwriting our observations. The most prominent of these are a mass density distribution over three-dimensional space, and a set of discrete “flash” events at space-time points. These ontologies provide strategies for solving some or all of the problems listed above.

My argument here is that, in the case of spontaneous collapse theories, primitive ontology is *redundant*, in the sense that the structures exhibited by the primitive ontologies that allow them to solve the problems facing spontaneous collapse theories are also present in the wave function. But redundancy is not *nonexistence*; indeed, the fact that the relevant structures are already there in the wave function shows that the mass density ontology and the flash ontology exist whether they are explicitly postulated or not. By the same token, there is no need to decide between a wave function ontology, a mass density ontology and a flash ontology.

In this regard, spontaneous collapse theories can be fruitfully contrasted with Bohm's theory. Bohmian particles are also sometimes described as primitive ontology, in response to the concern about the dimensionality of the wave function in particular. But in the Bohmian case, the primitive ontology is not redundant in the above sense: the particle structure is not already present in the wave function, but must be explicitly postulated. Hence there is a genuine question about whether Bohmian particles exist.

## 2 Collapse theories

Collapse theories were introduced to solve the measurement problem in quantum mechanics. The measurement problem arises because quantum states can be superposed, whereas measurement results (apparently) cannot. For example, suppose quantum state  $\psi_A$  represents a particle as located in some spatial region  $A$ , and quantum state  $\psi_B$  represents the particle as located in some disjoint spatial region  $B$ . Then according to quantum mechanics there is a third quantum state  $\psi_{AB}$  formed by superposing them:  $\psi_{AB} = 2^{-1/2}(\psi_A + \psi_B)$ . Such states are ubiquitous in quantum mechanics, and essential to its explanations.

Now consider a device that measures the location of the particle. When the state is  $\psi_A$  the device displays "A", and when the state is  $\psi_B$  the device displays "B". But when the state is  $\psi_{AB}$  the state of the device is not a superposition of displaying "A" and displaying "B"; rather, it displays "A" half the time and "B" half the time. Indeed, it is hard to see what a superposition of displaying "A" and displaying "B" could mean.

Von Neumann (1932, 186) dealt with this problem by proposing distinct dynamical laws for measurements and non-measurements. Between measurements, the quantum state evolves according to the continuous, linear Schrödinger dynamics—von Neumann's "process 2". But during a measurement, the quantum state jumps discontinuously into one of the eigenstates of the measured observable, with probabilities given by the square of the coefficient on that eigenstate in the pre-measurement state. This is von Neumann's "process 1", otherwise known as the *collapse postulate*. Applied to state  $\psi_{AB} = 2^{-1/2}(\psi_A + \psi_B)$ , the collapse postulate says that a measurement of the location of the particle causes the state to jump either to  $\psi_A$  or to  $\psi_B$ , with probabilities 1/2 each.

The trouble with von Neumann's proposal is that the two dynamical laws

are incompatible: neither can be reduced to the other. Furthermore, nothing physically distinctive takes place during a measurement: a measurement isn't a distinct kind of physical *process*, it is just one that has particular pragmatic importance *to us*. So there seems to be no *physical* distinction between measurements and non-measurements such that different fundamental laws could apply to measurements and non-measurements.

One response to this predicament is just to reject the collapse postulate, so that the Schrödinger dynamics applies to all systems at all times. This is the response embodied (in different ways) by Bohmian theories (Bohm 1952) and Everettian theories (Everett 1957). But a different response is also available—namely to precisify the collapse postulate so that it makes no essential appeal to measurement. This is the approach pursued by spontaneous collapse theories.

The original spontaneous collapse theory—the GRW theory of Ghirardi, Rimini and Weber (1986)—incorporates two incompatible dynamical processes, just like von Neumann's account. The first is the Schrödinger dynamics, which applies most of the time. The second is a collapse process, which localizes the quantum state in the coordinates of one particle. That is, for each particle in a system, there is a chance of  $10^{-16}$  per second that it will undergo a collapse, and if it does so, then the quantum state is multiplied by a Gaussian of width  $10^{-5}$  cm in the coordinates of that particle. The location of the center of the collapse is chosen at random so as to reproduce the statistical results of von Neumann's collapse postulate.

Unlike von Neumann's account, though, the GRW approach does not tie collapse to measurement: collapses just naturally occur at a certain rate, whether anyone is measuring a system or not. But the collapse process can nevertheless *explain* why measurements have unique outcomes. For microscopic systems, the collapse rate is so low that collapses are never observed. But for a macroscopic solid object containing of the order of  $10^{23}$  particles, the collapse rate is  $10^7$  per second, and since the positions of the particles in a solid object are highly correlated, a collapse for one particle suffices to localize the whole object.

Hence the GRW solution to the measurement problem: a single-particle state like  $\psi_{AB}$  is stable, but when the location of the particle is measured, its position is correlated with something we can see, such as the position of a pointer on a dial. The particle and pointer now form a composite macroscopic object, and the spontaneous collapse process very rapidly localizes the state of this object, either to a state in which the particle is in region *A* and

the pointer is pointing at “A”, or to a state in which the particle is in region  $B$  and the pointer is pointing at “B”. No special new process occurs during measurement; the measurement just amplifies the naturally-occurring spontaneous collapse process, due to the inevitable involvement of a large number of correlated particles.

### 3 Problems

As it stands, the GRW theory describes a world without particles. The representational machinery of the theory consists of the quantum state alone, and the quantum state, expressed as a wave function, is a spread-out entity rather than a spatially discrete particle. After a collapse the wave function “bunches up” in three coordinates, and under certain circumstances this “bunching” is relatively stable, so that we can use the language of particles and trajectories with some success to describe the way these “bunches” move around. But if the wave function describes all there is, then particles are not fundamental constituents of reality.

The view that some entity represented by the wave function is all there is (at the fundamental level) is called *wave function monism* (Ney 2013, 37). Wave function monism is a natural way of understanding the ontology associated with spontaneous collapse theories, but it raises some difficulties. The first is that the wave function is not defined over ordinary three-dimensional space, but over a  $3N$ -dimensional configuration space, where  $N$  is the number of particles in the system concerned (Bell 1987, 44). More precisely, since there are no particles at the fundamental level, the way to put it is that what we ordinarily think of as an  $N$ -particle system is described at the fundamental level by a  $3N$ -dimensional wave function. In particular, since we ordinarily think of the universe as a system containing a very large number of particles, then according to wave function monism it is better conceived as a single entity residing in a high-dimensional space.

The difficulty here concerns how a fundamentally high-dimensional world can nevertheless appear to us as three-dimensional. Albert (1996) maintains that the explanation is dynamical: the physical laws for macroscopic objects take a particularly simple form if the  $3N$  dimensions are grouped into threes, and so creatures like us should be expected to have an internal model of the world as three-dimensional. However, others are sceptical of this dynamical explanation. Monton (2006), for example, argues that if there is no three-

dimensional world, it is not clear that there *are* any “creatures like us” with internal models of the world.

A second difficulty with wave function monism stems from the fact that even after a collapse, the wave function in the coordinates of the collapsed particle is still non-zero everywhere. If we require that *all* the wavefunction amplitude associated with an object be contained in a particular region if the object is to count as being located in that region, then neither particles nor macroscopic objects ever have determinate locations according to spontaneous collapse theories, and we have no solution to the measurement problem after all (Albert and Loewer 1990). This is called the *tails problem*.

The obvious response here is to relax the requirement that all the wave function amplitude be contained in the relevant region: for an object to be located in a given region, it is sufficient that some large proportion of the wave function amplitude is contained in it (Albert and Loewer, 1996). This introduces a novel kind of vagueness, since presumably there is no precise matter of fact concerning the proportion that is required. And loosening the location requirement in this way produces other odd consequences: for example, each of a large collection of objects can individually count as being located in a region according to the requirement even when the collection as a whole does not count as being in the region (Lewis 1997). This is because when each of a collection of  $N$  objects has proportion  $p$  of its wave function amplitude in a region, where  $p$  is slightly less than 1, the proportion of the wave function of the whole system that is in the subspace of configuration space corresponding to all the objects being in the region is  $p^N$ , which can be considerably less than 1. This latter problem is called the *counting anomaly*.

A third difficulty with wave function monism concerns the reconciliation of the spread-out nature of the wave function with special relativity. A collapse in the coordinates of a particular particle affects the whole of space simultaneously, but according to special relativity there is no absolute standard of simultaneity, so the collapse seems ill-defined. Furthermore, for entangled particles, a collapse for any one of them instantaneously localizes all of them, no matter how far apart they are, which is in prima facie conflict with relativistic prohibitions on faster-than-light causation (Bell 1987).

None of these difficulties is unique to spontaneous collapse theories. The dimensionality problem and the counting anomaly affect any account of quantum mechanics that takes the wave function to represent a physical entity, including most ways of understanding Bohmian and Everettian theories (Lewis 2016, 97 and 151). The difficulty with relativity also affects Bohmian theo-

ries, as the Bohmian particle dynamics apparently requires an absolute standard of simultaneity (Bell 1981). Nevertheless, this particular combination of problems is most frequently addressed in the context of spontaneous collapse theories, and a particular kind of solution—adding ontology—is also most prevalent in the context of spontaneous collapse theories. It is this kind of ontological solution that I am interested in here.

## 4 Ontological solutions

The original proposal for adding ontology to a spontaneous collapse theory was made by Ghirardi, Grassi and Benatti (1995). They suggest adding a mass density distribution to the theory, defined in terms of the quantum state. Precisely, if  $M(\mathbf{r})$  is an operator for average mass density, then the mass density distribution  $\mathcal{M}(\mathbf{r}, t)$  can be defined in terms of the quantum state  $\psi(t)$  as  $\mathcal{M}(\mathbf{r}, t) = \langle \psi(t) | M(\mathbf{r}) | \psi(t) \rangle$ . It is the mass density distribution that directly accounts for the locations of macroscopic objects, and our perceptions of them.

Their motivations for adding a mass density distribution to the theory allude to the problems mentioned in the previous section. The mass density, unlike the quantum state, lives “in ordinary space” (Ghirardi, Grassi and Benatti 1995, 7). They regard this as an improvement over an ontology in  $3N$ -dimensional space, at least in the sense that it is “more traditional” (Ghirardi, Grassi and Benatti 1995, 37). Furthermore, they explain in detail how the mass density distribution avoids the tails problem: in essence, for a macroscopic object, the spontaneous collapse process localizes the wave function to such an extent that the mass density outside the immediate location of the object is undetectably small (Ghirardi, Grassi and Benatti 1995, 22–25).

A later article explains how the mass density ontology disposes of the counting anomaly (Bassi and Ghirardi 1999). If an object is a region of high mass density (surrounded by lower mass density), then when the whole of a collection of objects has a particular location precisely when each of the objects individually has that location. Since the location of objects according to the mass density ontology directly depends on the mass density in a particular region of three-dimensional space, and not directly on the proportion of wave function amplitude in a particular subspace of configuration space, the counting anomaly does not arise.

More recently, the mass density ontology has been taken up as an example of *primitive ontology* in quantum mechanics. Proponents of primitive ontology include Dürr, Goldstein, and Zanghi (1992), Allori (2013) and Maudlin (2013). The central constraints on primitive ontology are that its entities reside in three-dimensional space, and that the behaviour of these entities underwrites all the properties of the macroscopic objects we interact with (Allori 2013, 60). Clearly the mass density ontology fulfils these desiderata. The claim of the primitive ontologist is that any adequate quantum mechanical theory must be written in terms of the behaviour of some or other primitive ontology.

Allori and Maudlin both stress the importance of the three-dimensionality of the primitive ontology. Allori (2013, 62–63) echoes Ghirardi, Grassi and Benatti’s claim that the mass density ontology is “traditional”: “if you can account for everything you need to account for employing already successful and well-tested explanatory techniques, why not do so?” Wave function monism, on this view, is just too radical a break with the past. Maudlin goes further: the primitive ontology (or *primary ontology* in his terminology) *has* to be three-dimensional in order to determine our observations of three-dimensional objects (2013, 148). Every empirically adequate theory has to postulate a three-dimensional primitive ontology.

Primitive ontologies have also been suggested for other accounts of quantum mechanics. A mass density ontology has been developed for the branching worlds of Everettian quantum mechanics (Allori et al. 2011). Dürr, Goldstein, and Zanghi (1992) and Maudlin (2013) explore Bohmian quantum mechanics, and identify Bohmian particles as the primitive ontology in this case.

Conversely, other primitive ontologies have been proposed for spontaneous collapse theories. Most prominent is Tumulka’s *flash ontology* (2006), based on some comments by Bell (1987), according to which the primitive ontology consists of discrete events at space-time points corresponding to the centers of GRW collapse events. This ontology too, can straightforwardly address the issue of dimensionality, since the flashes are defined in three-dimensional space. It also bypasses the counting anomaly, since an object according to the flash ontology is just a swarm of flashes. Since the location of an object depends directly on the locations of flashes in three-dimensional space, and does not depend directly on the proportion of the wave function in a subspace of configuration space, the counting anomaly does not arise.

Additionally, the flash ontology suggests a way forward concerning the

conflict with special relativity: since the flashes are not spread out in space, a collapse event per se does not require instantaneous change everywhere. Indeed, the hope of a reconciliation with special relativity is the primary motivation behind Tumulka’s adoption of the flash ontology. Tumulka proposes that, 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 over simultaneity surfaces for which  $t$  is a constant. Hence for a single particle, a “flashy” version of the GRW theory can be made consistent with special relativity. But for two or more particles, there is still the worry that for entangled particles, a flash for one particle can instantly affect the future flash probabilities of the other, no matter how distant.

## 5 The primitive ontologist’s dilemma

The story so far is that wave function monism is problematic as an ontology for spontaneous collapse theories, and that either the mass density or the flash ontology can address some or all of these problems. This story can be resisted by arguing (as I have done elsewhere) that the dimensionality of the wave function and the counting anomaly are not really *problems* that need to be addressed (Lewis 2003; Lewis 2013). But perhaps some of them still look like problems—and certainly the problem of reconciling spontaneous collapse theories with special relativity remains. Here I take a different line: to the extent that dimensionality, the counting anomaly and consistency with special relativity are problems for spontaneous collapse theories, primitive ontology cannot provide distinctive solutions, despite initial appearances.

The argument is inspired by Hempel’s (1965) theoretician’s dilemma. Hempel points out that theoretical entities, if they do the job they are supposed to do, are redundant. And if they don’t do the job they are supposed to do, they are (trivially) redundant. So they are redundant. The job, as Hempel sees it, is to serve in explanations of observable phenomena. And the reason they are redundant if they do their job is that an explanation is a deductive argument from premises describing observable initial conditions to a conclusion describing an observable phenomenon, by way of suitable universal generalizations (laws). So an argument that uses laws mentioning theoretical entities can be replaced without loss by an argument whose laws



mention only observable things.

Hempel's theoretician's dilemma is only a challenge to the extent that one buys into his deductive-nomological account of explanation. If explanations in physics are not typically deductive arguments, then the fact that one can replace a deductive argument mentioning theoretical entities with one that does not has no significant consequences. In particular, theoretical entities might well be essential elements in physical explanations.

But now consider the use of primitive ontology to solve the problems facing spontaneous collapse theories. These problems are not (directly) causal in nature, and the role of primitive ontology is not to act as a *causal* intermediary between the wave function and the phenomena. Rather, the relevant dependency is something like determination rather than causation. The wave function at a time determines a mass density distribution at that time over three-dimensional space, solving the dimensionality problem and the counting anomaly. Similarly, the evolution of the wave function over a period of time determines a distribution of flashes over the relevant chunk of four-dimensional space-time, solving the dimensionality problem and the counting anomaly, and hopefully getting us closer to a reconciliation of spontaneous collapse theories with special relativity.

Even if causal explanations are not typically deductive arguments, explanations involving determination relations quite plausibly do typically involve deductive arguments. In a classical world, a suitable arrangement of particles *entails* that there is a table here. Similarly in a spontaneous collapse world with primitive ontology: a suitable wave function *entails* a particular mass density distribution, which in turn *entails* that there is a table here. Alternatively, a suitable wave function evolution *entails* a particular distribution of flashes, which in turn *entails* that there is a table here.

In either case, the wave function entails the state of the primitive ontology, and the state of the primitive ontology entails the state of observable things. Hence Hempel's dilemma seems to apply here: if the primitive ontology succeeds in its role as an intermediate step in this deductive chain, then it is redundant. If the table supervenes on the flashes, and the flashes on the wave function, then the table supervenes directly on the wave function. If the primitive ontology doesn't succeed in this role, then it is trivially redundant. So it is redundant.

In particular, since the state of the primitive ontology is entailed by the wave function, whatever structure there is in the primitive ontology that allows it to address the dimensionality problem, the counting problem and

the conflict with relativity, is structure that is already present in the wave function. So if primitive ontology succeeds at solving these problems, it is redundant. And of course, if it *fails* to solve the problems, then it is redundant. The upshot is that primitive ontology plays no crucial role in solving the problems afflicting spontaneous collapse theories; the wave function by itself can perform all the relevant work. So why did we think we needed it in the first place?

The response may be: The primitive ontology is not supposed to be an additional layer of reality residing between the fundamental wave function and the observable world. Rather, the primitive ontology is an ontological *replacement* for the wave function. In particular, supporters of the primitive ontology approach often try to reconceive the wave function as a *law* governing the evolution of the primitive ontology (Dürr, Goldstein and Zanghi 1992; Allori 2013).

But this is a difficult road to follow. The wave function in nonrelativistic quantum mechanics changes over time, so in this context one would have to admit that the laws of nature change over time. Perhaps this is not impossible to contemplate (Callender 2015), but it is certainly at odds with the claim often made in defense of primitive ontology that it is traditional, whereas wave function monism is a radical break with the past. If the primitive ontology approach requires time-evolving laws, this is a quite radical break too—arguably more radical than admitting a fundamental entity that is most readily defined over a high-dimensional space.

Alternatively, one can hope that in the ultimate relativistic quantum theory, the quantum state will not change over time (Dürr, Goldstein and Zanghi 1992). Goldstein and Zanghi (2013) show how, in a Bohmian theory with a time-independent wave function, the particles can still move, and hence the wave function can act as a law directing the primitive ontology. But no such approach is applicable to a spontaneous collapse approach, in which the dynamical role of the wave function in undergoing collapses seems essential to the theory. Certainly the most straightforward approach to understanding spontaneous collapse theories is that the quantum state is indeed a *state*, and hence a representation of the ontology of the system.

## 6 Primitive ontology made easy

The argument of the previous section is that primitive ontology is redundant. In causal contexts, redundancy arguments are typically followed by non-existence claims: if a postulated entity is explanatorily redundant, then we shouldn't say it exists. For example, if an absolute rest frame plays no explanatory role in mechanics, then we should say that there is no absolute rest frame. But that wouldn't be appropriate in the case of primitive ontology. The issue here isn't causal redundancy but what we might call *compositional* redundancy: if tables are nothing but distributions of mass density, and mass density distributions are nothing but wave function distributions, then tables are nothing but wave function distributions (and similarly for the flash ontology). Compositional redundancy doesn't lend itself to non-existence claims. I am nothing but an arrangement of cells (let's say), and my cells are nothing but an arrangement of molecules, but it doesn't follow that cells don't exist. Indeed, it trivially follows from the existence and arrangement of the molecules that my cells exist.

This is an instance the *easy* approach to ontology of Thomasson (2014). Tables exist, because all it takes for a table to exist is that the more fundamental ontology is arranged in a certain way. So, by the same token, if the wave function represents something physical, then the mass density ontology exists—and so does the flash ontology. The wave function distribution instantiates a particular mass density distribution over three-dimensional space. And the evolution of the wave function over time instantiates a set of flashes in space-time, since collapse events exist, and a flash is nothing but the center of a collapse event.

So the redundancy argument of the previous section should not be taken to be an argument against the *existence* of either the mass density ontology or the flash ontology. Rather, it should be taken as an argument that we don't have to explicitly *postulate* them, or to decide between them, since they both come along “for free” with wave function monism.

Indeed, both ontologies are *important*. The wave function is a highly structured entity. It contains the structures exemplified by the mass density distribution and by the flash distribution, in the sense that it entails both these structures. These structures do not *exhaust* the structure of the wave function, of course: the wave function entails the mass density distribution and the flash distribution, but not vice versa.

This is reflected in the fact that the quantum state cannot be represented

in four-dimensional space-time. (At least, it cannot *easily* be represented in a four-dimensional space-time: see Wallace and Timpson (2010) for a proposal about how to do so.) But nevertheless, the space-time structure is *there*, and the existence of three-dimensional primitive ontology *shows* that it is there. Similarly, the wave function contains structures we can use to solve the counting anomaly, and the use of the mass density language or the flash language to solve the counting anomaly shows this.

What about the conflict with special relativity? This is not so easily laid to rest. But the fact that the flash structure is present in the wave function gives some hope for the success of Tumulka's program. We do not need to say in addition that the wave function itself doesn't exist, even if there is a sense in which it is spread out in three-dimensional space (i.e. the sense expressed by the mass density distribution). Because of this spread-out nature, the wave function has different properties in different frames of reference, but this doesn't impugn it (Myrvold 2002). After all, cars and tunnels have different properties in different frames of reference: in one frame the car is longer than the tunnel, and in another the tunnel is longer than the car (Maudlin 1994, 54). And they certainly exist.

It is worth contrasting this with the case of Bohm's theory. Dürr, Goldstein and Zanghi (1992) and Maudlin (2013) count the particles of Bohm's theory as an example of primitive ontology, but it is important to note that, unlike the primitive ontologies proposed for spontaneous collapse theories, the existence and locations of the Bohmian particles are not entailed by the wave function. Rather, they have to be postulated separately, along with a new dynamical law to govern their motion. If you don't postulate particles in addition to the wave function in a no-collapse context, you end up with the ontology of a many-worlds theory. Bohm's theory and many-worlds theories are very different, with different strengths and weaknesses. So here there is a genuine, and difficult, ontological question. In the case of spontaneous collapse theories, though, questions about the existence of primitive ontology are easily answerable in the affirmative.

## 7 Conclusion

Often the ontology of spontaneous collapse theories is presented as a choice: either one adopts wave function monism, or one adopts a mass density ontology, or one adopts a flash ontology (e.g. Ney 2013, 40). My argument here

is that there is no pressing need to make such a choice. The wave function ontology entails the mass density ontology and the flash ontology, and with it the solutions to the apparent problems facing spontaneous collapse theories. Indeed, using the language of a mass density distribution or a flash distribution may be useful in explaining exactly how the problems can be resolved. But using these languages doesn't undermine the idea that ultimately, for a spontaneous collapse theory, all there is is the wave function.

## 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. 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: Proceedings of the 1990 Biennial Meeting of the Philosophy of Science Association*, vol. 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.
- Allori, Valia, Sheldon Goldstein, Roderich Tumulka, and Nino Zanghì (2011). "Many worlds and Schrödinger's first quantum theory," *British Journal for the Philosophy of Science* 62: 1–28.
- Bassi, Angelo, and GianCarlo Ghirardi (1999), "More about dynamical reduction and the enumeration principle" *British Journal for the Philosophy of Science* 50: 719–734.
- Bell, John S. (1981), "Quantum mechanics for cosmologists," in C. Isham, R. Penrose and D. Sciama (eds.), *Quantum Gravity 2*. Oxford: Clarendon Press, 611–637.

- Bell, John S. (1987), “Are there quantum jumps?” in C. W. Kilmister (ed.), *Schrödinger: Centenary Celebration of a Polymath*. Cambridge: Cambridge University Press, 41–52.
- Bohm, David (1952), “A suggested interpretation of the quantum theory in terms of “hidden” variables, parts I and II,” *Physical Review* 85: 166–193.
- Callender, Craig (2015), “One world, one beable,” *Synthese* 192: 3153–3177.
- Dürr, Detlef, Sheldon Goldstein, and Nino Zanghi (1992), “Quantum equilibrium and the origin of absolute uncertainty,” *Journal of Statistical Physics* 67: 843–907.
- Everett III, Hugh (1957), ““Relative state” formulation of quantum mechanics,” *Reviews of Modern Physics* 29: 454–462.
- 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, Alberto Rimini, and Tullio Weber (1986), “Unified dynamics for microscopic and macroscopic systems,” *Physical Review D* 34: 470–491.
- Goldstein, Sheldon, and Nino Zanghi (2013), “Reality and the role of the wave function in quantum theory,” in Ney and Albert (2013), 91–109.
- Hempel, Carl (1965), “The theoretician’s dilemma: A study in the logic of theory construction” in *Aspects of Scientific Explanation*. New York: Free Press, 173–226.
- Lewis, Peter J. (1997), “Quantum mechanics, orthogonality, and counting,” *British Journal for the Philosophy of Science* 48: 313–328.
- 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.

- Lewis, Peter J. (2016), *Quantum Ontology: A Guide to the Metaphysics of Quantum Mechanics*. Oxford: Oxford University Press.
- Maudlin, Tim (1994), *Quantum Non-Locality and Relativity*. Oxford: Blackwell.
- Maudlin, Tim (2013), “The nature of the quantum state,” in Ney and Albert (2013), 126–153.
- Monton, Bradley (2006), “Quantum mechanics and  $3N$ -dimensional space,” *Philosophy of Science* 73: 778–789.
- 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.
- Ney, Alyssa (2013), “Introduction,” in Ney and Albert (2013), 1–51.
- Ney, Alyssa and David Z. Albert (eds.) (2013), *The Wave Function*. Oxford: Oxford University Press.
- Thomasson, Amie L. (2014), *Ontology Made Easy*. Oxford: Oxford University Press.
- 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.
- von Neumann, Johann (1932), *Mathematische Grundlagen der Quantenmechanik*. Berlin: Springer.
- Wallace, David, and Christopher G. Timpson (2010), “Quantum mechanics on spacetime I: Spacetime state realism,” *British Journal for the Philosophy of Science* 61: 697–727.