On closing the circle

Peter J. Lewis

Abstract:

Ghirardi sought to “close the circle”—to find a place for human experience of measurement outcomes within quantum mechanics. I argue that Ghirardi’s spontaneous collapse approach succeeds at this task, and in fact does so even without the postulation of a particular account of “primitive ontology”, such as a mass density distribution or a discrete “flashes”. Nevertheless, I suggest that there is a remaining ontological problem facing spontaneous collapse theories concerning the use of classical concepts like “particle” in quantum mechanical explanation at the micro-level. Neither the mass density nor the flash ontology is any help with this problem.

1. Introduction

I remember the first time I came across Gian Carlo Ghirardi’s work. I was a graduate student at U. C. Irvine, and my advisor, Jeff Barrett, sent me to read the original GRW paper (Ghirardi, Rimini and Weber 1986). When I was done, I thought to myself “Well, that’s it. The physicists have solved the measurement problem, and there’s nothing left for us philosophers to do”. Fortunately (for me), my initial thought was premature; there was plenty of work left, for both physicists and philosophers, both refining the various spontaneous collapse models, and clarifying the surrounding concepts. Indeed, Ghirardi himself was deeply involved in both sides of this work.

One of the conceptual projects involves what Shimony (1989) calls “closing the circle”. Physics begins with human experience: we postulate physical theories to explain what we observe. Physical theories, insofar as they are successful, tell us what the world is like. But that world, of course, includes human beings and their experiences. So, for consistency, we need to be able to locate human beings and human experiences within the world-view provided by our physical theories.

Ghirardi recognized the importance of closing the circle in physics, and proposed a particular strategy for doing so in the context of his spontaneous collapse approach to quantum mechanics (Ghirardi, Grassi and Benatti 1995). What I want to do in this paper is to locate Ghirardi’s proposal within a broader discussion of what it takes to close the circle. In particular, I will argue that, while some proposals in the foundations of quantum mechanics fail to adequately close the circle, Ghirardi’s proposal succeeds; however, it does so by making more physical commitments than are strictly necessary. I advocate a strategy for closing the circle very like Ghirardi’s but pared of excess physical structure. Finally, I argue that the more pressing problem for spontaneous collapse theories concerns explanation at the microscopic level, and that here Ghirardi’s proposal is little help.

2. How not to close the circle

Closing the circle might seem like a trivial exercise. After all, if our physical theories describe the behavior of matter at the smallest scales, and if human beings are just complicated chunks of matter, then our physical theories automatically describe the behavior of human beings, including their eyes and their brains. How, then could a physical theory *fail* to find a place for human experience?

Closing the circle does indeed look trivial from a classical perspective, but quantum mechanics challenges much that we thought we could take for granted. The basic difficulty is just the measurement problem. Take a spin-1/2 particle, and prepare it in a superposition of two *z*-spin eigenstates: $2^{-1/2}(|\uparrow\rangle_z + |\downarrow\rangle_z)$. Now measure the spin of the particle along the *z*-axis. In the spirit of closing the circle, take the measuring device to be described by quantum mechanics, where $|\text{up}\rangle_m$ is an eigenstate in which the measuring device reads "spin-up", and $|\text{down}\rangle_m$ is an eigenstate in which it reads "spin-down". Quantum mechanics entails that after applying the measuring device to the particle and allowing their states to become correlated, the final state of the particle plus the measuring device is $2^{-1/2}(|\uparrow\rangle_z|\text{up}\rangle_m + |\downarrow\rangle_z|\text{down}\rangle_m)$. It looks like there is nothing in this final state that represents the (unique) outcome of this measurement. And since the measuring device *m* could include a human observer, there is nothing in the final state that represents the (unique) experience of the observer. So it looks like there is nowhere to locate human experience within the quantum formalism, and closing the circle becomes a *problem*.

Closing the circle is related to von Neumann's *psychophysical parallelism* (1955, 419), and indeed discussions of closing the circle are often couched in terms of postulating a psychophysical parallelism (Bell 1987, 45; Ghirardi, Grassi and Benatti 1995, 33). But it important to distinguish psychophysical parallelism from closing the circle, in particular because von Neumann's account of psychophysical parallelism *fails* to close the circle.

Von Neumann notes that there is at best a vague distinction between measured systems and measuring devices, and hence that it is arbitrary where physical analysis stops. When measuring temperature, for example, one can count the thermometer as external to the system, or as *part* of the system and hence subject to physical modelling. The same goes for human observers: one can treat the eye or the brain as external to the physical system, or as part of the physical system. Given this arbitrariness, it shouldn't matter to the predicted outcome of an experiment where we place this "cut". Von Neumann calls the principle that "the boundary between the observed system and the observer can be displaced arbitrarily" the "principle of the psycho-physical parallelism" (1955, 421)

This is a perfectly good methodological principle, and von Neumann's proof that his "collapse on measurement" formulation of quantum mechanics satisfies it is an important demonstration that it exhibits a particular kind of self-consistency. But it isn't the same as closing the circle. Note in particular his insistence that "we must always divide the world into two parts, the one being the observed system, the other being the observer" (1955, 420). The observed part is subject to physical modelling by quantum mechanics; for the observer, physical modelling "is meaningless" (1955, 420). This division crucial to his formulation: when the observed part does not interact with the observing part, we should model the observed part using the linear Schrödinger dynamics, but when the two portions interact, we should model the observed part using the non-linear collapse dynamics. In other words, the explanation for the collapse, and thereby for our experience of measurement outcomes, necessarily lies *outside* the system modelled by quantum mechanics. Von Neumann's approach blocks the possibility of closing the circle by design.

One might think that this barrier to closing the circle is inevitable. In introducing the concept, Shimony notes that “the greatest obstacle to “closing the circle” is the ancient one which haunted Descartes and Locke—the mind-body problem,” and conjectures that quantum mechanics may be “hospitable to a dualism of mind and body” (1989, 37). If the experiencing mind is non-physical, then of course it always lies beyond physical analysis, however far that analysis penetrates the workings of the brain. But positing dualism essentially just amounts to an admission that the circle can’t be closed. And it is a peculiarly unmotivated admission: our increasing knowledge of the brain strongly suggests that the varieties of human experience have a *physical* origin. The possibility of a “hard problem” of consciousness is irrelevant here: we know empirically that various experiential states are grounded in particular brain states, even if there is a residual explanatory gap concerning the phenomenal nature of those experiential states. The problem of “closing the circle”, then, is the problem of finding a way to explain those *brain states* quantum mechanically. It is here that von Neumann fails, since the quantum mechanical explanation of determinate post-measurement brain states appeals to the *deus ex machina* of interaction with an “observer”.

3. Closing the circle with quantum jumps

So Shimony’s concern about the mind-body problem is a red herring. Setting this concern aside, Shimony gives an indication of his preferred method for closing the circle, citing the GRW approach as the “most promising to date” (1989, 35). By adding a stochastic “collapse” term to the Schrödinger dynamics, the GRW approach can apparently explain measurement outcomes, including human brain states, without appeal to an extra-physical “observer”. The collapses ensure that for a system consisting of a large number of well-correlated particles, the wave function is, with high probability, well-localized around one point in configuration space. Since measuring devices and brains consist of large numbers of well-correlated particles, it looks like this is enough to close the circle.

However, there is a remaining gap in the circle, described here by Bell (1987, 44):

> There is nothing in this theory but the wavefunction. It is in the wavefunction that we must find an image of the physical world, and in particular of the arrangement of things in ordinary three-dimensional space. But the wavefunction as a whole lives in a much bigger space, of 3*N*-dimensions. It makes no sense to ask for the amplitude or phase or whatever of the wavefunction at a point in ordinary space. It has neither amplitude nor phase nor anything else until a multitude of points in ordinary three-space are specified.

The GRW dynamics governs the evolution of the wave function in 3*N*-dimensional space; on the face of it, it says *nothing* about a three-dimensional space inhabited by measuring instruments and brains. Certainly it can't be the *wave function* that is well-localized in three-space—that would be a mathematical category mistake. But if the GRW theory says nothing about the contents of three-space, obviously it can't close the circle.

Bell (1987, 45) proposes a way to plug the gap:

> However, the GRW jumps (which are part of the wavefunction, not something else) are well-localized in ordinary space. Indeed, each is centred on a particular spacetime point (**x**, *t*) … A piece of matter then is a galaxy of such events. As a schematic psychophysical parallelism we can suppose that our personal experience is more or less directly of events

> in particular pieces of matter, our brains, which events are in turn correlated with events in our bodies as a whole, and they in turn with events in the outer world.

The idea is that the GRW theory does, after all, describe the contents of three-space: it describes discrete, point-like events in three-space, with one such event corresponding to the center-point of each GRW collapse. The GRW collapse dynamics ensures that these point-like "flashes" will pick out a unique reading on the measurement apparatus or a unique brain state for an observer. Note that Bell's invocation of "psycho-physical parallelism" isn't a reference to von Neumann's principle about arbitrarily moving the observer-observed boundary. Bell simply means to point out that experience covaries with a person's brain state, so if the GRW dynamics can ensure determinate brain states, it can ensure determinate experience, and thus close the circle.

Ghirardi sees the same problem, but proposes a different solution. Rather than point-like events, Ghirardi proposes that the GRW theory describes a mass density distribution in three-dimensional space, defined in terms of the configuration-space wave function $\Psi(t)$ by

$$\mathcal{M}(\mathbf{r}, t) = \langle \Psi(t) | M(\mathbf{r}) | \Psi(t) \rangle,$$

where $M(\mathbf{r})$ is the mass density operator for location **r** in three-space (Ghirardi, Grassi and Benatti 1995, 16). This proposal has the advantage over Bell's that it is adaptable to continuous variants of the GRW theory (e.g. Ghirardi, Pearle and Rimini 1990), in which there are no discrete collapse events. As far as closing the circle goes, though, it proceeds very much like Bell's version: the GRW dynamics (discrete or continuous) makes sure that the mass density distribution picks out a unique reading of the measurement apparatus or a unique brain state of the observer (Ghirardi, Grassi and Benatti 1995, 36).

4. Primitive ontology

Hence we have two distinct proposals for closing the circle within the GRW model, one “flashy” and one “massy”. These have been described as distinct *primitive ontologies* for the GRW theory (Allori et al. 2008, 359). A primitive ontology plays a dual role: it is that which the theory is *about*, and it is that which explains the properties of everyday macroscopic objects (Allori 2013, 60). That is, the existence of a primitive ontology allows for the possibility of “closing the circle”, as it is in terms of the primitive ontology that the connection between the scientific image and the manifest image is spelled out. Indeed, Allori (2013, 66–69) takes the problematic nature of quantum mechanics to stem in part from the fact that it was developed *without* a primitive ontology, either because of the conviction that a realist understanding of the theory is impossible, or because of the conviction that it describes the *wave function*, which because of its high-dimensional nature is ill-suited to play the role of primitive ontology. Hence Allori (2013, 69–70) concludes that we should supply quantum mechanics with a primitive ontology after the fact—and in the case of the GRW theory, that means either flashes or a mass density distribution.

I do not dispute the need for a theory to “close the circle”, and if primitive ontology is ontology such that the circle can be closed, I do not dispute the need for that either. But I have some qualms about the particular proposals on offer. My initial worry is methodological: How do you find out what the primitive ontology of a theory is? Allori (2013, 63) is surely right that the primitive ontology of a theory can’t simply be read off its mathematical formulation. Since the dynamical law of quantum mechanics governs the evolution of the wave function over time, much as Newton's laws govern particle positions, one might think that the ontology of quantum mechanics directly corresponds to the wave function—that quantum mechanics is *about* a wave-like entity inhabiting a high-dimensional configuration space. Allori points out that this is not how we identify primitive ontology. In the case of Newtonian mechanics, an ontology of point masses is *presupposed* as the starting point for physical theory construction, not read off the theory after the fact. Indeed, given that many disparate physical systems can be modelled using the same mathematics, reading the ontology off the mathematics seems doomed to failure. Think of the variety of applications of the mathematics of the simple harmonic oscillator!

So if we can't read the ontology off the theory, how should we proceed? Allori (2013, 69) suggests that primitive ontologies are "proposals" about how to understand quantum mechanics. That is, we use our ingenuity to come up with an ontology that quantum mechanics *could* be about. This seems very much in the spirit of Ghirardi's proposal of a mass density ontology, and perhaps also in the spirit of Bell's proposal of a flash ontology—although note Bell's (1987, 45) insistence that the flashes are "part of the wavefunction, not something else". The basic idea is that the primitive ontology of quantum mechanics is a separate *hypothesis*, a hypothesis that is supported to the extent that it can explain the properties of macroscopic objects, including measuring devices and human brains.

Certainly one *can* postulate a primitive ontology for quantum mechanics. But I see danger in this approach. The most obvious danger is underdetermination: there are too many competing proposals, and no way to decide between them. In addition to a mass density ontology and a flash ontology, one might propose that the quantum state describes properties of spacetime regions (Wallace and Timpson 2010), or a collective property of a set of particles (Monton 2013). There are doubtless many other possibilities. How can we determine which is correct? Since each ontology is constructed to be fully consistent with the predictions of quantum mechanics, there is no possibility of an empirical answer. Perhaps extra-empirical virtues like explanatory power can come to the rescue here, but the historical track-record of this approach is debatable.

Furthermore, the explanatory power of the mass density and flash ontologies can be called into question. Consider, for example, a solid object whose quantum wave function is well localized in a particular region of configuration space, with “tails” extending elsewhere. The corresponding mass density distribution is large in the relevant region of 3-space, and small elsewhere. Now consider a second object passing through the region where the mass density of the first is small. How should we expect it to behave? The mass density picture suggests that the second object moves through a “sea” of rarefied matter, resulting in a small but constant force. But spontaneous collapse quantum mechanics tells us that there is no such force; instead, the small “tails” on the wave function tell us the probability of a collapse in which the first object moves discontinuously, say into the path of the second object. Rather than a small constant force, there is zero force, with a small probability of a large force.

The problem, then, is that the mass density ontology suggests a continuous effect where the reality is discontinuous. The flash ontology suffers from the opposite problem: it suggests discontinuity even in continuous cases. Consider, for example, a small object consisting of around 10^{19} particles. If each particle suffers a GRW collapse every 10^{16} seconds, there is a flash along the trajectory of the object roughly once every millisecond. Between these times, there is no ontology corresponding to the object whatsoever. Nevertheless, despite the discontinuity of the primitive ontology, the gravitational and electromagnetic forces exerted by the object on surrounding objects will be *continuous* in time.

One might object to both these examples that they ignore the full explanatory apparatus of quantum mechanics: quantum mechanics can explain both the discontinuous behavior in the first case and the continuous behavior in the second. That is correct, but the point is that the explanation in each case is given by the wave function and the Born rule, not the primitive ontology. The wave function may be explanatorily suspect because it inhabits a high-dimensional space (Allori 2013, 59), but if the idea is that the primitive ontology can "provide an explanatory scheme derived along the lines of the classical one" (Allori 2013, 70), neither the mass density ontology nor the flash ontology clearly meets the classical explanatory standard.

5. Wave function as structure

I have argued that simply *positing* a primitive ontology for quantum mechanics is a risky business, both because of the potential of radical ontological underdetermination, and because the proposed ontologies may fail to do the requisite explanatory work. How *should* we identify the appropriate primitive ontology, then? Allori (2013, 63) suggests a historical approach:

> The mathematical formalism of a theory has a *history* that constrains the interpretation of its formalism: the theory started with a metaphysical position and its appropriate mathematical representation, and it continued with the implementation of the suitable mathematical apparatus necessary to determine how the primitive ontology evolves.

So, for example, Newton *begins* with an ontological posit—that there are massive objects moving in a three-dimensional space—and *then* constructs the relevant mathematical tools to represent the ontology and its temporal evolution. Rather than trying to divine the appropriate ontology by gazing at the mathematics of the final theory, we should look to the interpretation intended by the developers of that theory.

Unfortunately, though, as Allori (2013, 67) is keenly aware, the development of quantum mechanics doesn't seem to fit this model. Although Schrödinger began with a particular interpretation of the wave function in mind—a three-dimensional field—he abandoned this interpretation when he realized that the wave function for multi-particle systems is defined on configuration space, not 3-space. With the blessing of Bohr and Heisenberg, quantum mechanics forged ahead *without* any conception of the ontology described by the mathematics.

Does this mean that we are forced to posit a primitive ontology for quantum mechanics after the fact, to remedy the oversight of its developers? Perhaps not: I think we can make some progress by considering other historical precedents. While quantum mechanics may be unique in being a theory developed in the *absence* of a primitive ontology, there are a number of historical examples of theories developed on the basis of a *mistaken* primitive ontology. Consider, for example, Fresnel's wave theory of light. Fresnel begins with an ontological posit—an all-pervasive elastic solid—and constructs a mathematical theory of transverse waves in this medium. We think Fresnel's mathematical theory was essentially correct, even though we now think there is no such elastic solid.

What should we make of cases like this? Worrall (1989) takes them as evidence for *structural realism*: scientific theories tell us about the structure of the world, but not in general about what instantiates that structure. This view has a good deal of plausibility. What do we know about the ontological nature of *mass* or *charge* or *spin*, over and above the mathematical structures of the physical theories containing those terms? “Nothing” seems like an appropriate answer.

Suppose, then, that we take a structural approach to the wave function. What would that mean? It would mean endorsing the claim that the wave function correctly describes the structure of physical systems in certain contexts, but without endorsing any particular account of the kind of thing that instantiates this structure. This is not the same as wave function realism—the position that the wave function describes a fundamental entity in a high-dimensional space. Rather, the wave function describes the structure instantiated by whatever fundamental entities there may be in ordinary three-dimensional space: particles, fields, flashes, mass density, or something else entirely. A structure is not in itself an *object*, but rather a way that objects relate to each other.

Of course, “structure” is a rather vague term, and it is reasonable to ask for more details about the sense of the term “structure” as it is used here. I wish I had more to say, but for now I only have a negative characterization to give. One kind of structure the world exhibits is *nomological* structure: events exhibit regularities, and those regularities are (or are produced by) *laws*. Dürr, Goldstein and Zanghi (1992) suggest that wave function structure is nomological structure. But the main motivation for taking wave function structure as nomological is that there is a sense in which it can be taken to govern the dynamical evolution of the primitive ontology. The wave function fixes the motion of the particles in Bohm's theory, the evolution of the mass density distribution in mass-density GRW, and the probability distribution of flashes in flashy GRW. If one withholds from endorsing any particular primitive ontology, then there is no particular reason to think that the relationship between the wave function and the ontology is best characterized as nomological.

Furthermore, even if we endorse one of the existing proposals concerning primitive ontology, there are well-known reasons to resist thinking of the wave function as nomological: the form of the wave function depends on the nature of the system under consideration, and it changes over time (Brown and Wallace 2005, 533). We don't usually conceive of laws this way. Of course, we can always extend our conception of law to include contingent, time-evolving laws (Callender 2015, 3157), but such a move threatens to elide an important distinction. Suppose we apply quantum mechanics to the motion of a set of charged particles. These particles exert forces on each other according to an inverse square law, and this law is reflected in the Hamiltonian term appearing in the Schrödinger equation. The inverse square law is neither contingent nor time-evolving. One could always propose that there are two basic kinds of law, but it does less damage to standard physical thought to conceive of the wave function as a summary of the relations between the three-dimensional entities involved (whatever they may be), rather than a law governing those entities.

6. Closing the circle—and opening another

My proposal, then, is that we should think of the wave function as a structure, but withhold commitment to any particular account of the ontology that instantiates this structure. And my claim is that this is enough to close the circle, at least for a spontaneous collapse theory. To appreciate what it takes to close the circle, consider again the case of Fresnel's wave theory of light. Poisson famously derived from this theory the existence of a bright spot in the center of a circular shadow. Mathematically speaking, what he actually derived was a region of high-amplitude wave structure surrounded by a region of low-amplitude wave structure. This was enough to underwrite the existence of the bright spot, even without any hypothesis concerning the nature of the ontology that instantiates the wave structure. That is, Fresnel's wave theory closes the circle: it enables us to locate observable experimental outcomes within the framework of the theory.

We can do the same with a spontaneous collapse theory. Consider a measurement of the spin of a spin-1/2 particle in which it is deflected by a magnetic field and then to one of a pair of suitably-positioned detectors, each of which responds to detection by raising a flag. If the particle is initially in a symmetric superposition of spins along the measurement direction, then the wave function of the particle plus detection apparatus evolves to a symmetric superposition in configuration space, but one term in this superposition is rapidly made many orders of magnitude larger than the other by the spontaneous collapse mechanism. That is, the vast majority of the wave function amplitude at the end of the measurement is concentrated around a particular small region of configuration space. Even in the absence of a preferred account of the underlying ontology, we know how to interpret this structure: it is the structure of things in three-space, whatever their underlying ontology may turn out to be. Ordinary objects like flags are made out of this three-dimensional ontology. That is, the high-amplitude region of configuration space tells us the locations of the two flags in three-space: it is either a structure in which the "up" flag is raised or a structure in which the "down" flag is raised.

Hence there is no need to endorse any particular account of primitive ontology so that a spontaneous collapse theory can close the circle. For macroscopic systems, spontaneous collapse picks out a particular small region of configuration space, and that region specifies the locations of ordinary objects in 3-space. Thus we can find the unique observed outcome of a measurement within the theory. More directly, since neuroscience suggests that our experience supervenes on the electrochemical configuration of our brains, the specification of a small region of configuration space also specifies human experiences.

Can we remain agnostic about primitive ontology then? I am not sure. I suspect that spontaneous collapse theories do face an explanatory problem, but that it concerns the micro-world rather than human experience. We have no serious difficulty locating measurement results and our experience of them within the structural framework of spontaneous collapse theories. The things that are more difficult to locate are the explanatory entities of classical physics—particles and fields. As Healey (2015) has forcefully argued, quantum explanation is parasitic on classical concepts. We measure the spin of a spin-1/2 *particle* by passing it through a magnetic *field*. The Hamiltonian term in the Schrödinger equation contains a term corresponding to the interaction of this point-particle with the field. But how do we find these particles and fields in the wave function structure?

One might think that it is here that the proposals for primitive ontology might do some work. The flash ontology, though, is clearly of no use: microscopic systems produce no flashes over reasonable time-scales, and hence correspond to no ontology whatsoever according to the flash proposal. The mass density ontology is more promising, the obvious approach being to treat a particle as a localized region of high mass density. Ghirardi cautions against such an interpretation, though. Essentially, the problem is the microscopic analog of the explanatory worry for the mass density ontology explained in section 4. When the mass density for a particle is spread out over a large region, a second particle passing through the region won't experience a continuous, small force, as the spread-out mass density would lead you to expect, but rather zero force, with a small probability of a large force ("collision"). Hence Ghirardi, Grassi and Benatti (1995, 18) construct a criterion for deciding when the mass density distribution is "objective", one that typically applies to macroscopic objects but not to single particles.

One could reject this latter move, and hold that microscopic systems have an associated mass density that is just as objective as that of macroscopic systems (Monton 2004). On this account, a single "particle" is really a localized region of high mass density, albeit one that can split in two or spread out. One might quite reasonably think that this is just how quantum "particles" behave: sometimes they spread out, and act more like waves. Concerning Ghirardi's worry about the behavior of regions of low mass density, Monton responds that the anomalous behavior is explained by the wave function alone, not the mass density. Monton is happy to concede that "mass density is epiphenomenal" (2004, 419). But while this may be acceptable for Monton's purposes, clearly mass density is not functioning as primitive ontology here.

So neither the mass density ontology nor the flash ontology adequately explains the role of particles in quantum mechanics. One possible move at this point would be to insist that the primitive ontology *is* just particles in three-dimensional space. The primitive ontology approach is flexible: in addition to the obvious particle-based quantum theory (Bohm's theory), versions of the GRW theory can be constructed that have a primitive ontology of particles (Allori et al. 2014; Allori 2019). The challenges facing this approach are well known—most notably that it is hard to square the law governing the evolution of the particles with relativity. But this looks to me like the direction to take. That is, I submit that the real ontological puzzle of the quantum world doesn't concern human experience, but rather concerns how our physical theories—quantum, relativistic, and classical—hang together.

7. Conclusion

Ghirardi's physical and philosophical insight ran deep. He realized the need for an account of the quantum world in which it is possible to locate our experience of measurement outcomes, and unlike von Neumann, he succeeded at producing one. In fact, matters are simpler than he realized: there is no need for a mass density distribution, or any other particular account of primitive ontology, in order to close the circle, since the wave function, understood as a structure of three-dimensional things, can do the job by itself.

However, this doesn't mean that there is no further work to be done. Quantum mechanics does not only have to be hospitable to human experience; it also needs to be hospitable to classical explanation, since quantum explanation is parasitic on classical. Ghirardi's mass density ontology is of little help here. Unless, with Healey (2015, 11), we give up thinking of quantum mechanics as *descriptive* at all, we need to look elsewhere for an understanding of the ontology of the quantum world.

References

Allori, Valia (2013), “Primitive ontology and the structure of fundamental physical theories,” in A. Ney and D. Z. Albert (eds.), *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*. New York: Oxford University Press, 58–75.

Allori, Valia (2019), “Scientific Realism without the wave-function: an example of naturalized quantum metaphysics,” In: J. Saatsi and S. French (eds.), *Scientific Realism and the Quantum.* Oxford: Oxford University Press.

Allori, Valia, Sheldon Goldstein, Roderich Tumulka, and Nino Zanghí (2008), “On the common structure of Bohmian mechanics and the Ghirardi-Rimini-Weber theory,” *British Journal for the Philosophy of Science* 59: 353–389.

Allori, Valia, Sheldon Goldstein, Roderich Tumulka, and Nino Zanghí (2014) “Predictions and primitive ontology in quantum foundations: a study of examples,” *British Journal for the Philosophy of Science* 65 (2): 323–352.

Bell, J. S. (1987), “Are there quantum jumps?” in C. W. Kilmister (ed.), *Schrödinger: Centenary of a Polymath*, Cambridge University Press, 41–52. Reprinted in Bell, J. S. (2004), *Speakable and Unspeakable in Quantum Mechanics*, Second Edition. Cambridge: Cambridge University Press, 201–212.

Bell, J. S. (2004), *Speakable and Unspeakable in Quantum Mechanics*, Second Edition, Cambridge University Press.

Brown, Harvey R., and David Wallace (2005), “Solving the Measurement Problem: De Broglie--Bohm Loses Out to Everett,” *Foundations of Physics* 35: 517–540.

Callender, Craig (2015), “One world, one beable,” *Synthese* 192: 3153–3177.

Dürr, Detlef, Sheldon Goldstein, and Nino Zanghì (1992), "Quantum equilibrium and the origin of absolute uncertainty," *Journal of Statistical Physics* 67: 843–907.

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

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

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

Healey, Richard A. (2015) "How quantum theory helps us explain," *British Journal for the Philosophy of Science* 66: 1–43.

Monton, Bradley (2004), "The problem of ontology for spontaneous collapse theories," *Studies in History and Philosophy of Modern Physics* 35: 407–421.

Monton, Bradley (2013), "Against 3N-dimensional space," in A. Ney and D. Z. Albert (eds.), *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*. New York: Oxford University Press, 154–167.

Shimony, A. (1989), "Search for a World View that will Accommodate our Knowledge of Microphysics," in J. Cushing and E. McMullin (eds.), *Philosophical Consequences of Quantum Theory: Reflections on Bell's Theorem*, pp. 25–37. Notre Dame: Notre Dame University Press.

Von Neumann, J. (1955), *Mathematical Foundations of Quantum Mechanics*, Princeton University Press. Originally published (1932) as *Mathematische Grundlagen der Quantenmechanik*, 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.

Worrall, John (1989), “Structural realism: The best of both worlds?” *Dialectica* 43: 99–124.