Original Paper

In search of local beables

Peter J. Lewis

Department of Philosophy, University of Miami, P.O. Box 248054, Coral Gables FL 33124-4670, USA

E-mail: plewis@miami.edu

Abstract: The call to supplement the quantum wave function with local beables is almost as old as quantum mechanics. But what exactly is the problem with the wave function as the representation of a quantum system? I canvass three potential problems with the wave function: the well-known problems of incompleteness and dimensionality, and the lesser known problem of non-locality introduced recently by Myrvold. Building on Myrvold's insight, I show that the standard ways of introducing local beables into quantum mechanics are unsuccessful. I consider whether we really need local beables, and assess the prospects for a new theory of local beables.

Keywords: Wave function; locality; hidden variables; spontaneous collapse; many worlds; retrocausation; configuration space; Humean supervenience

1. The call for local beables

Forty years ago this year, J. S. Bell gave a talk called "The Theory of Local Beables".[1] In it, he introduces the term "beable" as a name for a putative element of reality in the quantum world, and suggests that the beables we should be particularly interested in are *local* in the sense that they can be assigned to some bounded space-time region (Bell 2004, 53). But even if the term originates with Bell, the call for local beables clearly echoes the EPR argument exactly forty years prior to that (Einstein, Podolsky and Rosen 1935).

---

[1] At the sixth GIFT seminar, Jaca, Spain, 2–7 June 1975. The paper is published in Bell (2004, 52–62).

The motivation is essentially the same: the wave function via which standard quantum mechanics represents physical systems is inadequate to the task, and hence needs to be supplemented with (or replaced by) something that genuinely represents the properties of the system.

The need for local beables remains controversial, and any particular account of them doubly so. My sense is that the debate itself is not terribly clearly defined. What, precisely, is the problem with the wave function that calls for the addition of local beables? What exactly would count as a local beable? And how far do the various accounts of local beables on offer succeed at solving the problems with the wave function representation? My present purpose is to try to make a little headway in answering these questions.

Let me start with the question of motivation. What's wrong with the wave function anyway? There are a number of concerns one might have. First, there's the EPR worry that the description of reality provided by the wave function is incomplete. The argument here is that since measuring one of a pair of entangled particles allows you to predict with certainty the outcome for the other particle, the latter particle must already have a property corresponding to the outcome of the measurement—and the wave function doesn't represent that property. Bell, of course, complicates this discussion by proving that any method of ascribing properties to entangled systems would have to violate some highly plausible physical assumption, e.g. causal locality. But his concern is essentially the same as Einstein's: we need local beables because the wave function is representationally incomplete.

The representational incompleteness of the wave function is not limited to entangled states, however. Schrödinger's cat thought experiment highlights a dilemma facing accounts of quantum mechanical measurement (Schrödinger 1935—also 80 years old this year!). Either you say that a measurement precipitates a collapse of the wave function—in which case you face the difficult task of defining just which physical processes constitute measurements—or else you don't—in which case there is nothing in the wave function representation corresponding to the unique outcome of the measurement. This is one way of expressing the measurement problem.

But representational incompleteness isn't the only concern you might have with the wave function. Another concern raised by Bell is that the wave function "propagates not in 3-space, but in $3N$-space" (2004, 128). That is, the wave function for an N-particle system is a function of $3N$ spatial coordinates: it inhabits a configuration space rather than ordinary three-dimensional space. Again, this problem takes a particularly stark form when the state of the system is entangled, since the wave function of the pair of particles in six-dimensional space can't be reduced to two wave functions in three-dimensional space without loss of information. But the problem is quite general: presumably the measurement outcomes we observe are localized in ordinary three-dimensional space, and the high-dimensional wave function doesn't directly represent such things. If we are to find

the familiar three-dimensional world represented in the wave function, we need to be told how.

Finally, as Myrvold (2014) has recently remarked, even in its own $3N$-dimensional space the wave function is an inherently non-local representation of a system, in the following sense: a non-zero wave function amplitude in some region (of $3N$-space) is incompatible with the squared wave function amplitude integrating to 1 over any disjoint region. Hence a non-zero amplitude in one region of $3N$-space carries implications for the amplitude right now in any region of $3N$-space, however distant. It seems reasonable to think that the properties of the system in my lab are local (in three-dimensional space) in the sense that they carry no implications for the properties of the far side of the moon right now. Of course, the properties of the system in my lab might carry implications for the properties of the far side of the moon *later*, via causal processes, but the non-locality at issue here is definitional rather than causal.[2] That is, the wave function cannot be defined over a bounded region even of $3N$-dimensional space without reference to the rest of the space, because the normalization of the wave function refers to *all* of the space.

For ease of reference, let us call these three problems with the wave function *incompleteness*, *dimensionality* and *non-locality*. Of course, not everyone is convinced that these are genuine problems that need to be solved. (Nor are they strictly independent problems: they interrelate in various ways to be considered below, and they all in some way implicate the representational completeness of the wave function.) But suppose you are impressed by some or all of these concerns. It seems that you must supplement (or even replace) the wave function as a representation of quantum systems; you must invent a new theory. And of course there is no shortage of contenders.

2. Incompleteness

Several attempts have been made to add local beables to quantum mechanics, most notably hidden variable theories in the style of Bohm (1952), and spontaneous collapse theories in the style of Ghirardi, Rimini and Weber (1986). Bohm's theory supplements the wave function with particles that are "pushed around" by the wave function. These particles are local in Bell's ontological sense: the beables are the particle positions (at a time), which can be assigned to space-time points.[3] The GRW theory supplements standard quantum

---

2   There are three senses of "locality" to keep straight here. There is Bell's sense: a beable is local in some space iff it can be assigned to a bounded region of that space. Let us call this *ontological* locality. There is Myrvold's sense: a beable is local in some space iff it can be defined over a bounded region of that space, without reference to the rest of the space. Let us call this *definitional* locality. And finally there is *causal* locality: a causal process is local iff it can be ascribed to a time-like trajectory. I will endeavor to specify which I mean.

3   I take the beables to be local properties rather than local objects. But for the purposes of this paper I don't

mechanics with a spontaneous collapse mechanism that postulates a small chance per unit time per particle that the wave function will become localized in the coordinates of that particle. The local beables are a little harder to discern here: the localization is imperfect, in the sense that a wave function that is non-zero everywhere before collapse will still be non-zero everywhere after collapse. But as Bell (2004, 205) notes, the center-point of a collapse event can be regarded as a local beable in his ontological sense: the collapse event can be regarded as assigning a property to a space-time point. This is the "flashy" ontology for GRW endorsed by Tumulka (2006). Alternatively, one can supplement the wave function of GRW with a mass density distribution over three-dimensional space, defined so that the mass density is large where the squared wave function amplitude is large (Ghirardi, Grassi and Benatti 1995; Allori et al. 2008). In this "massy" ontology for GRW, the beables are mass density assignments to space-time points.

Both of these theories are most directly aimed at the incompleteness problem. In Bohm's theory, the positions of the particles represent the outcomes of our measurements, without the need for collapse. So even though the wave function does not represent the unique outcome of a measurement, the particle positions perform this job. In the GRW theory there are collapses, but there is no need for a problematic collapse-on-measurement postulate, because measurements automatically precipitate collapse: when a quantum system is correlated with a macroscopic pointer, the sheer number of particles involved means that a spontaneous collapse is almost certain in a very small period of time. In the case of the flashy ontology, the location of the initial collapse event picks out the unique outcome of the measurement, and in the case of the massy ontology, the collapse event means that the mass density is concentrated on one outcome.[4]

Similar stories are available for entangled states, although the solution to the incompleteness problem for such states is less than satisfactory. In Bohm's theory, a measurement on one particle can determine the properties of both particles. So, for example, although Bohmian particles always have determinate position properties, they do not always have determinate spin properties, and for an entangled pair in the singlet state $2^{-1/2} \left(|{\uparrow_z}\rangle_1 |{\downarrow_z}\rangle_2 - |{\downarrow_z}\rangle_1 |{\uparrow_z}\rangle_2 \right)$, a z-spin measurement on either particle causes the Bohmian particles to move within the wave function so as to fix the spin properties of *both* particles. The particle positions succeed at representing the outcomes of our measurements, but in this case a causally non-local influence between the two particles is required as part of the explanation. In the GRW theory, a measurement on either particle in the singlet

---

want to make anything of the property/object distinction: hopefully it doesn't matter.

[4]   Worries can be raised about the adequacy of both Bohm's theory and the GRW theory in explaining measurement results (e.g. Brown and Wallace 2005; Cordero 1999). But these worries can, I think, be addressed (Lewis 2007).

state precipitates a wave function collapse centered on one term or the other—either on $|\uparrow_z\rangle_1 |\downarrow_z\rangle_2$ or on $|\downarrow_z\rangle_1 |\uparrow_z\rangle_2$—and hence both particles acquire determinate spin values. The evolution of the wave function explains the outcomes we observe, but again a causally non-local influence is involved in the explanation of the outcomes. The beables introduced by Bohm's theory and the GRW theory are ontologically local—they reside at well-defined points—but the laws that direct them are causally non-local. Hence these theories succeed at solving the incompleteness problem for entangled states, but at the cost of introducing causal non-locality.

3. Dimensionality

Bohm's theory and the GRW theory were devised to solve the incompleteness problem. But what about the other problems with the wave function? Prima facie, both Bohm's theory and the GRW theory provide a solution to the dimensionality problem. Bohmian particles always have determinate locations in ordinary three-dimensional space, even when the wave function requires a higher-dimensional representation. Similarly, the flashes of flashy GRW always have determinate locations in three-dimensional space, and the mass density of massy GRW is distributed in a determinate way over three-dimensional space, even when the wave function cannot be represented in such a space. Hence the ontology in all three cases is undeniably three-dimensional.

But there still is a remaining dimensionality problem. In Bohm's theory, the dynamical law for the particles is such that their motion depends on the wave function: this is the sense in which the wave function "pushes around" the particles. But if the wave function inhabits a high-dimensional configuration space, and the particles inhabit a separate 3-space, it is hard to see how the wave function can push around the particles. Similarly in the GRW theory, the flashes and the mass density distribution have a law-like connection to the wave function, and hence the evolution of flashes or mass density over time depends on the evolution of the wave function over time. But since the wave function inhabits one space and the flashes or mass density inhabit another, it is hard to see how the dependence of one on the other could be mediated. This is a separate issue from the causal non-locality just mentioned; it is not that the causation involved here operates instantaneously across space, it is that it apparently operates between two distinct spaces.

There are two general approaches one might take to dealing with this residual problem. First, one might postulate that all the ontology of the relevant theory lives in the high-dimensional space. That is, the Bohmian particles for an N-particle system are really just a manner of speaking about a *single* point in a $3N$-dimensional space, a point that accounts for the outcomes of our measurements. Similarly, GRW collapses are events that concentrate the wave function amplitude around a single location in $3N$-space, again accounting for the results of our measurements (without any appeal to three-dimensional

ontology such as flashes or a mass density distribution). Albert (1996) has endorsed such an account. The main challenge here is to explain how events in a fundamentally $3N$-dimensional space can yield the *appearance* that the measurement outcomes we observe are situated in 3-dimensional space. Albert is happy to take on that explanatory burden, although the extent to which he succeeds is a matter of ongoing debate (Ney and Albert 2013).

The other approach is to postulate that all the ontology of the relevant theory lives in three-dimensional space. The most prominent proposals for achieving this point beyond the confines of non-relativistic quantum mechanics. For instance, Goldstein and Zanghì (2013) postulate that in the context of a quantum theory that can incorporate gravity, the time-dependence of the wave function may drop out, and hence the wave function may function as a constant law rather than as a time-evolving entity. This opens up the possibility that the ontology of Bohm's theory consists entirely of particles in 3-space evolving according to this law. If tenable, this proposal rescues three-dimensional local beables (in Bell's ontological sense) in the context of Bohm's theory. But quantum gravity is rather speculative territory: there is no settled theory yet, so it is unclear whether quantum gravity really has the consequence of a static wave function, and it is doubtful whether the laws that would push around the Bohmian particles in such a theory could be made causally local.[5]

Putting aside such speculation, Wallace and Timpson (2010) and Myrvold (2014) argue that the wave function of non-relativistic quantum mechanics reduces to local properties of the three-dimensional field in quantum field theory, and hence the ontology of quantum mechanics is fundamentally three-dimensional, with the configuration-space wave function simply acting as a convenient shorthand representation in the non-relativistic limit.[6]

Except for Goldstein and Zanghì's proposal, what the remaining responses to the residual dimensionality problem have in common is that they essentially give up on the demand for ontologically local beables—beables that can be assigned to a bounded region of 3-space. If the ontology of quantum mechanics is fundamentally high-dimensional, then the beables are assigned to a bounded region of the high-dimensional space, but not necessarily to a bounded region of the three-dimensional world as it appears to us. For example, in Bohm's theory our measurement outcomes are explained via a single point in the high-dimensional space, but in general this point can correspond to locations

---

[5] Bell's theorem suggests that any theory of ontologically local beables must be causally non-local (Bell 1964). Many worlds theories and retrocausal theories might evade this conclusion; they are discussed below.

[6] It is worth noting, though, that neither Wallace and Timpson nor Myrvold are attempting to defend either a Bohmian or a spontaneous collapse theory, since their sympathies lie with the Everettian approach described in the next section.

indefinitely far apart in the three-dimensional space of experience. Similarly in the GRW theory, our measurement outcomes are explained by a wave function that is strongly peaked around a point in the high-dimensional space, but which may correspond to locations indefinitely far apart in three-dimensional space. And even if the fundamental ontology is three-dimensional, entanglement means that some of the properties of quantum systems are irreducibly relational, applying equally to two or more locations in three-dimensional space (Wallace and Timpson 2010, 713).

The lesson that Wallace and Timpson (2010) and Myrvold (2014) draw is that *local* beables are unnecessary: this is explored later. But there may be philosophical reasons to prefer beables that are ontologically local in a high-dimensional space rather than no local beables at all, and hence to prefer the high-dimensional ontology to the three-dimensional one. For example, Barry Loewer (1996) defends a high-dimensional ontology for quantum mechanics on the grounds that it allows us to retain Humean supervenience—David Lewis's doctrine that all the properties of a system supervene on the local properties of its smallest parts. However, there is one further problem to consider, and it suggests that Loewer's defense of Humean supervenience may be misguided.

4. Non-locality

The remaining problem is Myrvold's non-locality problem—the problem that a non-zero squared wave function amplitude in one location has implications for the squared wave function amplitude at distant locations. As Myrvold notes, the move to a high-dimensional space is no help here: if the squared wave function amplitude is non-zero in some region of configuration space, then it cannot be the case that it integrates to 1 over some disjoint region of configuration space (2014, 4). Wave function properties may be ontologically local in their own high-dimensional space, but even in that space, they are not *definitionally* local.

The reason that this is significant is that Humean supervenience is ordinarily taken to be a commitment to both ontological and definitional locality: beables should be ascribable to one region of the relevant space without reference to other regions. Wave function properties fail this requirement: you cannot ascribe amplitude properties to points in *this* region of $3N$-space however you like, because the amplitude properties in some other region might rule out some assignments to this region. If the integral of the squared wave function amplitude over some distant region is close to 1, then it cannot also be close to 1 over this region.

To the extent that Bohm's theory and the GRW theory rely on the wave function as part of their ontology, they suffer from this problem. If the single point representing the Bohmian particles is located in a particular region of configuration space, then the squared wave function amplitude must be non-zero in that region, and hence can't integrate to

1 elsewhere. If a GRW collapse occurs to a given region of configuration space, then the squared wave function amplitude is again non-zero in that region, and hence cannot integrate to 1 elsewhere.

You might think that the violation of definitional locality is just a truism: if my desk is *here*, then trivially my desk is not anywhere else. But the point is that Humean supervenience requires that local properties carry no implications for other regions of space. The fact that my desk is here carries no implications for whether there is a desk (or anything else) in the next office over. But the fact that the squared wave function amplitude is non-zero here *does* restrict its values elsewhere.

It is tempting to think that the non-locality identified by Myrvold rests on a mistaken understanding of the normalization of the wave function. Because of its connection to probability via the Born rule, the squared amplitude of the wave function must integrate to 1 over the whole of space. But if the wave function is regarded as an entity spread over $3N$-space, presumably it describes the distribution of some kind of *stuff* over that space, and then one might think that normalization is just a fact about the *proportion* of wave function stuff in a particular region, not the absolute quantity. Suppose that there is some quantity of wave function stuff located in *this* region of $3N$-space. This carries no implications, one might think, for how much wave function stuff is located elsewhere, so Humean supervenience is safe. The only implication is that *if* there is a lot of wave function stuff elsewhere, then the *proportion* of squared wave function amplitude in this region is low, and if there is only a little elsewhere then the proportion here is high.

But this hope is short-lived. Such a proposal wouldn't give you any kind of Humean supervenience worth having, because the beables in a region would be radically disconnected from what you should expect to observe if you look at the region. If there is a lot of wave function stuff elsewhere, then the probability of finding the system in this region is low, and if there is a little, then it is high. We want the beables to explain what we observe, and the current proposal fails that test.

Along similar lines, you might propose that the normalization condition is a contingent initial condition (Albert 1996, 278): it just so happens that there is a certain amount of wave function stuff distributed over configuration space, and the Schrödinger equation means that the squared wave function amplitude is conserved over time. But while this might serve as a kind of explanation of definitional non-locality, it doesn't eliminate it. Compare this with ordinary conservation laws. In classical mechanics, if the universe is created with a certain amount of mass in it, then the amount of mass in one region carries implications for the amount of mass elsewhere. There is no violation of Humean supervenience: there *could have been* more mass elsewhere without affecting the amount of mass we can observe here. But in the quantum case, as explained above, if there had been more wave function stuff elsewhere, then the probabilities associated with our observations in this region would have been different. That is, keeping fixed the probabilities of various observations *here*, there

couldn't have been more wave function stuff elsewhere.

Alternatively, one might think that Myrvold's non-locality problem is just an artifact of considering only non-relativistic systems with a fixed number of particles. If there are N particles in the system, and there are N particles in this region (either because of Bohmian particle beables or GRW wave function beables), then it cannot be the case that there are any particles elsewhere. But it is unfair to suggest that this is a violation of Humean supervenience, one might think, because the constraint that there are exactly N particles in the world is a *global* fact about the world, and combining local beables with a global fact can certainly have non-local implications.

But in fact the normalization of the quantum state carries straightforward non-local implications even when there is not a fixed number of particles (Myrvold 2014, 16). Suppose the (Bohmian or GRW) beables are such that there is exactly one particle in a given region of space.[7] This means that at least some of the squared amplitude is associated with one particle being in this region, and this rules out the possibility that all the squared amplitude is associated with finding exactly three particles (or whatever) in some distant region of space. Hence the beables still carry non-local implications: normalization is the culprit, not the assumption of a fixed number of particles.

The only way I can see for avoiding this kind of definitional non-locality—for making quantum mechanics over $3N$-space consistent with Humean supervenience—is the possibility that we may be able regard the wave function as a law rather than an entity, as described in the previous section. But it is worth noting that in that case the move to a high-dimensional space does no work: if the wave function is a law governing the evolution of beables, those beables may as well live in three dimensions. The move to a high-dimensional space is not a way to recover Humean supervenience.

So it looks like attempts to introduce local beables don't get us as far as we might have liked. While Bohm's theory and the GRW theory offer a direct solution to the incompleteness problem, they do so at the expense of introducing causal non-locality. And the dimensionality problem and Myrvold's non-locality problem resist solution, unless some speculation about the time-dependence of the wave function in quantum gravity pays off. Many commentators are already convinced that the price for local beables—namely instantaneous action at a distance—is too high. If we add to that Myrvold's point that the beables we get aren't even fully local (in the definitional sense), then the price starts to look like money for nothing.

5. Who needs local beables?

---

[7] I don't wish to imply here that it is straightforward or even possible to extend Bohm's theory or the GRW theory to the relativistic domain.

Even Bell, the biggest champion of local beables, concedes that "we may be obliged to develop theories in which there *are* no strictly local beables" (2004, 53). Perhaps the thing to do at this point is to concede that no quantum mechanical theory in terms of local beables is possible—that Humean supervenience is dead. Indeed, if the motivation for local beables is primarily philosophical—to save the doctrine of Humean supervenience—then it's hard to see that much is lost: this is just another example of a philosophical intuition that falls to empirical science. But Bell and Einstein were not primarily motivated (if at all) by such intuitions; their concern was with the physical adequacy of the theory in light of the incompleteness problem. Bohm and GRW deliver this much at least.

Can we do better? So far I have said nothing about the many worlds theory—the third of the "big three" interpretations. According to its advocates, the many worlds theory can solve the incompleteness problem without recourse to Bohmian particles or GRW collapses. The trick is that a structure of decoherent branches is identified in the wave function, and beables representing the outcomes of measurements are identified within each branch. That is, if the z-spins of two particles the singlet state $2^{-1/2} \left(\vert{\uparrow_z}\rangle_1 \vert{\downarrow_z}\rangle_2 - \vert{\downarrow_z}\rangle_1 \vert{\uparrow_z}\rangle_2 \right)$ are measured, then decoherent branches are produced, relative to some of which the state is close to $\vert{\uparrow_z}\rangle_1 \vert{\downarrow_z}\rangle_2$, and relative to others of which the state is close to $\vert{\downarrow_z}\rangle_1 \vert{\uparrow_z}\rangle_2$. Hence (its advocates conclude), the wave function itself can provide all the beables we need, and there never was an incompleteness problem in quantum mechanics.

The many worlds theory also works just as well as Bohm and GRW (if not better) at tackling the dimensionality problem. The wave function in the many worlds theory is interpreted realistically, and taken at face value it occupies a high-dimensional configuration space. One could try to supplement the wave function with ontology that resides in three-dimensional space (Allori et al. 2011), but that would raise the worries about interaction mentioned in section 3. So the relevant options seem the same as before: either defend the idea that reality is fundamentally high-dimensional, or show how the wave function can be interpreted as representing properties in three-dimensional space.

So the many worlds theory gives us beables, beables that solve the incompleteness problem, and address the dimensionality problem as far as it needs addressing. But it doesn't give us *local* beables in the definitional sense: the non-locality problem applies just as readily to the many worlds theory as to Bohm and GRW. Suppose for example, that my branch of the wave function is such that there is a particle in a particular region of space. This requires that most of the wave function amplitude in this branch is contained in the relevant region in the coordinates of the particle. And this in turn rules out most of the wave function amplitude in this branch being contained in some distant region of space.[8]

---

[8]   Sometimes "my branch" will not be defined for sufficiently distant regions of space. But let us assume a

Humean supervenience is not regained.

Myrvold (2014) and Wallace and Timpson (2010) are quite sanguine about the non-local nature of the beables in the many worlds theory. They can certainly be purchased far more cheaply than the non-local beables in Bohm and GRW: there is (arguably) no need for instantaneous action at a distance (Wallace and Timpson 2010, 713). And it looks like no interpretation of quantum mechanics that connects the beables to the quantum state can do better: Myrvold's non-locality problem follows from the normalization of the quantum state, so any interpretation in which the beables carry implications for the quantum state will suffer from this problem.

Still, there is something quite strange about the source of the non-locality. Bohm's theory, the GRW theory and the many worlds theory are all realist about the quantum state: the wave function describes the distribution of something physical over (ordinary or configuration) space. Normalization is an odd requirement to impose on the distribution of physical stuff. Explanatory strategies such as regarding normalization as the product of a conservation law are inadequate, as shown in the previous section. What it looks like of course—given the Born rule—is an *epistemic* constraint, since our degrees of belief should always sum to 1. It is strange that an epistemic constraint should act on the world.

Perhaps this is a tendentious way of putting things. Contemporary many-worlders like Wallace (2012) might say that I get things backwards: the normalization of the wave function is a constraint on the world, and while it might be prima facie strange that it should correspond so directly to a constraint on my beliefs, there are decision-theoretic arguments why this should be so. Still, the decision-theoretic arguments remain controversial,[9] and the source of the normalization constraint remains mysterious.

6. Restoring local beables

Of course, strangeness is no real objection to a theory, especially one like many worlds quantum mechanics. But if we like to have things explained, then it would be better if we could construe the wave function epistemically, since then the normalization constraint has a straightforward explanation. As a side effect, this also restores the possibility of defending Humean supervenience, since Myrvold's non-locality problem doesn't arise. But epistemic construals of the wave function face formidable obstacles, most notably a number of no-go theorems (Bell 1964; Kochen and Specker 1967; Pusey, Barrett and Rudolph 2012).[10]

---

scenario and a distant region for which it is defined.

[9] See e.g. the essays in Part IV of Saunders et al. (2010).

[10] Of course, the no-go theorems apply equally to theories in which the wave function is an entity: they constrain property ascription in many worlds theories, for example. But Everettians can shrug their

One way forward is to exploit the so-called "independence loophole" in the no-go theorems. As Price (1994) and Leifer (2011) point out, the no-go theorems all assume that the properties of a system are independent of the measurements performed on it. This assumption might be violated if causation were a time-symmetric phenomenon—if particles could carry the effects of *later* measurements performed on them, just as they carry the effects of *earlier* measurements. Then there is no barrier, in principle, to the wave function playing a purely epistemic role, where the ontology consists of particles and their local properties.

So for example, when we describe a pair of particles using the singlet state $2^{-1/2} (|\uparrow_z\rangle_1 |\downarrow_z\rangle_2 - |\downarrow_z\rangle_1 |\uparrow_z\rangle_2)$, this simply means that owing to the way the particles were produced, we don't know (prior to z-spin measurements on the individual particles) whether particle 1 is z-spin-up and particle 2 is z-spin-down or vice versa. Nevertheless, one of these is the case: the particles have well-defined individual spin properties. The reason this doesn't violate Bell's theorem is that the spin values depend on the measurements performed: if spin measurements in different directions had been performed on the particles, then their earlier spin properties would have been different.

A retrocausal interpretation of quantum mechanics of this kind means that the local beables are precisely the properties of the particles revealed on measurement—in this case, their spins. This has a number of advantages. Clearly, the incompleteness problem doesn't arise: particles have pre-existing properties corresponding to the results of our measurements, and the fact that the wave function doesn't represent those properties is of no consequence, because the wave function just represents our knowledge of the system. Similarly, the particles and their properties reside in three-dimensional space, so the dimensionality problem doesn't arise either. The fact that the wave function is defined over configuration space simply reflects the complexities of our knowledge of quantum systems: for entangled systems like the pair of particles in the singlet state, we not only know the possible spin properties for each particle individually, we also know the correlations between them—in this case that they have opposite spins when measured in any given direction. This information is most readily represented in a configuration space.

Finally, Myrvold's non-locality problem doesn't arise in a retrocausal theory. Suppose a particle is located in a particular region of space. Nothing follows about the beables in distant regions of space: maybe there is a particle there, maybe there isn't. The

---

shoulders—of course you can't ascribe pre-measurement properties to systems corresponding to their unique outcomes, because measurements typically don't have unique outcomes. Epistemic views typically entail that the wave function represents our knowledge of pre-existing properties that are revealed on measurement; this is precisely the picture that the theorems make trouble for. Griffiths (2011) argues that Bell's theorem doesn't threaten epistemic views, understood correctly, but Maudlin (2011) responds that Griffiths misidentifies the flaw in Bell's reasoning.

normalization of the wave function is irrelevant here, because there is no general connection between the location of the particle (a fact about the world) and the wave function (a description of our knowledge). It is true that if we restrict ourselves to systems with a fixed number of particles (as in non-relativistic quantum mechanics), then the fact that there is a particle here does have implications for how many particles there are elsewhere—but in the retrocausal case it is clearly the global assumption about the number of particles that introduces the non-locality. Without the assumption of a fixed number of particles (as in quantum field theory) the location of a particle carries no implications for distant regions. The fact that there is a particle here does not by itself entail that there is a non-zero wave function amplitude here, because I might be convinced that there is no particle here.

However, suppose I do ascribe a non-zero probability to the particle being located here. Then the wave function amplitude will be non-zero in this region, which carries the implication that the squared wave function amplitude does not integrate to 1 over some distant region. But this apparent non-locality is just a matter of what credences I can simultaneously entertain: if I have a non-zero credence that there is a particle here, then I cannot also have credences totalling to 1 in possibilities that exclude a particle being here. Nothing follows about whether or not there is a particle in any distant region. Normalization applies to my credences, not to the world.

All these advantages are purchased without paying the price of instantaneous action at a distance: all action is along time-like lines, although some of that action is in the reverse temporal direction. So if a genuine, fully-fledged retrocausal theory of quantum mechanics were available, it would be an attractive contender. But unfortunately there is as yet no such thing, although there are a number of ongoing research programs. The sticking place, as one might expect, is interference: if the wave function is purely epistemic, how can it exhibit interference effects? While some suggestions concerning the origin of interference effects in wave-function-epistemic theories have been made (e.g. Price 1996, 255; Spekkens 2007), others have concluded that waves have an ineliminable role even in retrocausal theories (Wharton 2010; Kastner 2012). It remains to be seen how this all plays out, and whether local beables remain viable.

Acknowledgments

References

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

[2] Allori, Valia, Sheldon Goldstein, Roderich Tumulka, and Nino Zanghì (2008), "On the common structure of Bohmian mechanics and the Ghirardi–Rimini–Weber theory," *British Journal for the Philosophy of Science* 59: 353–389.

[3] 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: 1–27.

[4] Bell, John S. (1964), "On the Einstein-Podolsky-Rosen Paradox," *Physics* 1: 195–200. Reprinted in Bell (2004): 14–21.

[5] Bell, John S. (2004), *Speakable and Unspeakable in Quantum Mechanics, Second Edition*. Cambridge: Cambridge University Press.

[6] Bohm, David (1952), "A suggested interpretation of the quantum theory in terms of "hidden" variables," *Physical Review* 85: 166–193.

[7] Brown, Harvey R., and David Wallace (2005). "Solving the measurement problem: De Broglie-Bohm loses out to Everett," *Foundations of Physics* 35: 517–540.

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

[9] Einstein, Albert, Boris Podolsky, and Nathan Rosen (1935), "Can quantum-mechanical description of physical reality be considered complete?" *Physical Review* 47: 777–780.

[10] 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.

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

[12] Goldstein, Sheldon and Nino Zanghì (2013), "Reality and the role of the wave function in quantum theory," in Ney and Albert (2013): 91–109.

[13] Griffiths, Robert B. (2011), "EPR, Bell, and quantum locality," *American Journal of Physics* 79: 954–965.

[14] Kastner, Ruth E. (2012), *The Transactional Interpretation of Quantum Mechanics: the Reality of Possibility*. Cambridge: Cambridge University Press.

[15] Kochen, S., and E. P. Specker (1967), "The problem of hidden variables in quantum mechanics," *Journal of Mathematics and Mechanics* 17: 59–87.

[16] Leifer, Matt (2011), "Can the quantum state be interpreted statistically?" URL: http://mattleifer.info/2011/11/20/can-the-quantum-state-be-interpreted-statistically/

[17] Lewis, Peter J. (2007), "Empty waves in Bohmian quantum mechanics," *British Journal for the Philosophy of Science* 58: 787–803.

[18] Loewer, Barry (1996), "Humean supervenience," *Philosophical Topics* 24: 101–127.

[19] Maudlin, Tim (2011), "How Bell reasoned: A reply to Griffiths," *American Journal of Physics* 79: 966–970.

[20] Myrvold, Wayne C. (2014), "What is a wave function?" *Synthese* (forthcoming). URL: http://philsci-archive.pitt.edu/id/eprint/11117

[21] Ney, Alyssa, and David Z. Albert (2013) (eds.), *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*. Oxford: Oxford University Press.

[22] Price, Huw (1994), "A neglected route to realism about quantum mechanics," *Mind* 103: 303–336.

[23] Price, Huw (1996), *Time's Arrow and Archimedes' Point: New Directions for the Physics of Time*. Oxford: Oxford University Press.

[24] Pusey, Matthew F., Jonathan Barrett, and Terry Rudolph (2012), "On the reality of the quantum state," *Nature Physics* 8: 475–478.

[25] Saunders, Simon, Jonathan Barrett, Adrian Kent, and David Wallace (2010) (eds.), *Many Worlds?: Everett, Quantum Theory, & Reality*. Oxford: Oxford University Press.

[26] Schrödinger, Erwin (1935), "Die gegenwärtige Situation in der Quantenmechanik," *Naturwissenschaften* 23: 823–828.

[27] Spekkens, Robert W. (2007), "Evidence for the epistemic view of quantum states: A toy theory," *Physical Review A* 75: 032110.

[28] Tumulka, Roderich (2006), "A relativistic version of the Ghirardi-Rimini-Weber model", *Journal of Statistical Physics* 125: 821–840.

[29] Wallace, David (2012), *The Emergent Multiverse: Quantum Theory According to the Everett Interpretation*. Oxford: Oxford University Press.

[30] 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.

[31] Wharton, K. B. (2010), "A novel interpretation of the Klein-Gordon equation," *Foundations of Physics* 40: 313–332.