Against wavefunction realism

David Wallace

February 14, 2017

Abstract

I argue that wavefunction realism — the view that quantum mechanics reveals the fundamental ontology of the world to be a field on a high-dimensional spacetime, must be rejected as (a) relying on artefacts of too-simple versions of quantum mechanics, and (b) not conceptually well-motivated even were those too-simple versions exactly correct. I end with some brief comments on the role of spacetime in any satisfactory account of the metaphysics of extant quantum theories.

1 Wavefunction realism: the outline case

The outline case for wavefunction realism begins with quantum mechanics, formulated (for N particles) something like the following:

1. The *wavefunction* of the system, at time t, is a function ψ_t from the $3N$-fold Cartesian product of the real numbers, $\mathbf{R}^{3N}$, to the complex numbers. (We can conveniently absorb the time index into the function, so that the wavefunction assigns a complex number $\psi(x_1, \ldots q_{3N}; t)$ to each $3N$-tuple of real numbers and each time t.)

2. The dynamical evolution of the wavefunction (at least when the system is not being observed) is given by the Schr ́odinger equation,

$$\frac{\partial}{\partial t}\psi = -\frac{i}{\hbar}\left(\sum_{k=1}^{N}\frac{\hbar^2}{2m_k}\nabla_k^2\psi + V(q_1, \ldots, q_{3N})\psi\right) \tag{1}$$

where m_k is the mass of the kth particle, $\hbar$ is Planck's constant, and ∇_k^2 is the Laplacian differential operator with respect to the kth triple of coordinates on $\mathbf{R}^{3N}$.

3. If the positions of the particles are simultaneously observed at time t, the probability of their $3N$ position coordinates lying in a small region of volume δV around the point $(q_1, \ldots, q_{3N})$ is

$$\Pr = |\psi(q_1, \ldots, q_{3N})|^2 \delta V. \tag{2}$$

The third of these premises suggests that ψ ought to be thought of as some sort of probability distribution or parametrisation of our restricted knowledge, but a wealth of arguments (beginning with elementary observations about interference, and proceeding through formal no-go theorems both classic (Kochen and Specker 1967; Gleason 1957) and modern (Pusey, Barrett, and Rudolph 2011) make clear that this strategy is not really viable, at least without concessions to operationalism (Fuchs and Peres 2000; Fuchs 2002) and/or pragmatism (Healey 2012) that philosophers of physics have by and large been loath to make. At least from a scientific realist's perspective, it looks as if the quantum state has to be taken as representational: different quantum states[1] represent different ways the world could be, not simply different levels of human information about the world.

Wavefunction realism — as advocated originally by Albert (1996) and since defended by, e. g. , Ney (2013b) and North (2013) (and also by Ney, this volume) then seems to follow straightforwardly from the formalism: what the wavefunction *is*, clearly, is that familiar mathematical object, a field on space, assigning a value independently to every point in that space. But the space on which the wavefunction lives is not familiar 3-dimensional space (represented mathematically by $\mathbf{R}^3$) but rather $3N$-dimensional space ('configuration space' as it is called in physics). And since N is the number of particles in the system — and since, from a metaphysician's perspective, 'the system' is presumably the whole universe, even the observable part of which contains $\sim 10^{80}$ electrons and atomic nuclei — this is a space of staggering dimension.

Quantum mechanics, then — says the wavefunction realist — reveals to us that ordinary space is illusory (Albert 1996) or at any rate non-fundamental (Ney 2013b). The *real* space shown to us by physics is wildly different.

The obvious questions raised by this revelation are: (i) why does space *appear* to us to be 3-dimensional, and, relatedly, (ii) how do we connect this radical take on fundamental reality with our observations and so empirically confirm QM? The pragmatic response to the latter is straightforward, via the probability rule (3) above: when the wavefunction is peaked around a point in $3N$-dimensional space, we should expect to observe the particles arranged with 3-dimensional coordinates corresponding to the coordinates of the point at which they are peaked. (The question of what we should observe when the wavefunction is not so peaked leads us to the related, but distinct, *quantum measurement problem*, which lies outside the scope of this essay.) But that pragmatic response, absent additional argument, is unprincipled: *why* should we expect such observations? How does this actually follow from the fundamental ontology and dynamics?

The recent literature has largely explored two answers. The strategy of Albert, Ney and North has sought to recover the three-dimensional world as an emergent, higher-level description, in analogy with the general pattern by which

[1] There is a caveat: ψ and $e^{i\theta}\psi$ make precisely the same predictions and so are empirically indistinguishable even in principle. The norm in physics is to regard them as the same state (and thus to accept a slight redundancy in the formalism, which can be eliminated by moving to a density-operator or ray formulation of quantum mechanics. See Maudlin (2013) for further consideration of this point in the context of wavefunction realism.

higher-level ontology emerges from the lower level. The primitive-ontology strategy of Allori *et al*(2008, cf also Allori 2013) and Maudlin (2013) instead explicitly supplements the formalism of quantum mechanics with new, fundamental three-dimensional ontology (typically this additional ontology is also designed to solve the measurement problem).

Both strategies agree on this much: that *if* unsupplemented quantum mechanics (with or without some dynamical collapse to solve the measurement problem) is correct, then we do not fundamentally live in a three-dimensional world. But (I will argue) this is a misreading of the quantum formalism. On both technical and conceptual grounds, the move from unmodified quantum mechanics to an ontology with a single fundamental entity — the wavefunction — should be rejected.[2] Technical, because the account relies on features of a certain simplified version of quantum mechanics that does not generalise; conceptual, because even if that simplified version of quantum mechanics is accepted uncritically, the move to wavefunction realism is unmotivated.

2 The technical case against wavefunction realism

'Quantum mechanics', like classical mechanics, is a framework theory: within that framework, a great many different physical systems can be described, ranging from simple two-state systems, though collections of interacting particles, through relativistic fields and even (speculatively) including the dynamics of spacetime itself. The version of quantum mechanics described above, by contrast, is a *specific* theory: to be precise, the theory of a finite number of spinless particles, interacting through long-range potential forces at nonrelativistic fields. We might call the latter theory 'toy nonrelativistic quantum mechanics', or 'toy NRQM'. The 'toy' epithet is a little harsh from the point of view of practical physics — with tweaking, it is widely used in applications — but strictly interpreted its scope is very narrow, being insufficient to analyse (for instance) any process involving light, or any atom more complex than hydrogen. Metaphysical conclusions based on it can be of value only insofar as they generalise —either to arbitrary theories falling under the quantum framework or, at a minimum, to more realistic quantum theories with a wider domain of applicability. This is not the case, as we will see.[3]

The first thing to note is that even for toy NRQM, the particular presentation given above is only one of many mathematically equivalent ways in which the theory can be formulated. It could equally well, for instance, have been given in the *momentum space representation*, where the wavefunction obeys a quite different differential equation and where $|\psi|^2$ gives the probabilities for

[2] Monton (2013) and Lewis (2013) also pursue a strategy somewhat along these lines, though differing significantly in the details (and focussed on the conceptual rather than the technical issues).

[3] The quantum mechanics I discuss in this article is in all cases established work and I do not attempt to give original references.

momentum, rather than position, measurements. Abstractly speaking, modern quantum mechanics is formulated on a (usually infinite-dimensional) space of vectors, called *Hilbert space*. Wavefunctions are just one way of describing these abstract vectors (usually called the 'position representation'); momentum-space wavefunctions are another, and there are indefinitely many more; physicists shift between them freely according to which is most helpful in solving a particular problem. And so wavefunction realism already seems to involve an arbitrary choice; though, to be fair, in toy NRQM the position representation has a central role in the formulation of the theory, so it is not difficult to imagine making a case for preferring it metaphysically.

As a first move beyond toy NRQM (one sufficient to analyse atomic structure, though not to treat phenomena involving light), consider that almost all the particles analysed by quantum mechanics have *spin*, i. e. intrinsic angular momentum. A single particle with spin 1/2 (the most common case, and the one that describes the electron and the proton) is described not by a wavefunction taking values in the complex numbers $\mathbf{C}$, but by a wavefunction taking values in a two-dimensional Hilbert space — roughly, the Cartesian product $\mathbf{C}^2$ of two copies of $\mathbf{C}$. That in itself is scarcely problematic — a complex-vector-space-valued field is scarcely more metaphysically abstruse than a complex-valued field. But when we consider N particles rather than one, things get — literally — exponentially more complicated. It does not suffice to give one spin vector for each point in configuration space, or even N spin vectors. What is required, rather, is a vector in a 2^N dimensional Hilbert space, or put another way, the wavefunction for N spin-half particles is specified not by one complex number at each point of configuration space, but by 2^N complex numbers. So what wavefunction realism delivers for, say, a universe of 10^{80} spin-half particles is not 'merely' a function from a 3×10^{80}-dimensional space to the complex numbers, but a function from a 3×10^{80}-dimensional space to a $2^{10^{80}}$-dimensional complex vector space. This is a radical departure from the original conception of wavefunction realism. Or (if we are willing to choose an arbitrary direction in space, which defines a preferred set of coordinates for the $2^{10^{80}}$-dimensional space), we can instead see wavefunction realism as delivering not one complex-valued wavefunction, but $2^{10^{80}}$ of them; as metaphysical underdetermination goes, indeterminacy as to whether there is fundamentally one thing or $2^{10^{80}}$ things isn't bad going.

Now let's consider dropping the 'nonrelativistic' part of toy NRQM, and considering particles with relative velocities approaching light. There is a well-formulated (albeit limited) quantum theory of relativistic particles; however, that theory is formulated without direct reference to the position representation (it gives, instead, a central role to the momentum representation) and indeed there is no unproblematic way to *define* a position representation in relativistic particle physics. See (e.g.) Saunders (1998), Fleming (2000), or Halvorson (2001) for (fairly technical) details, but in essence, there are two different candidates. One has the property that it's impossible to define even in principle what it means for a particle to be sharply localised in a region; the other violates the

Principle of Relativity.

From the pragmatic point of view of applying relativistic quantum mechanics, this is all harmless: the 'position' measurements we make (via cloud chambers et al) are physical processes, and if they were ever precise enough to distinguish between different definitions of position (they're not) then their physical details would suffice to determine what is in fact being measured. But it leaves wavefunction realism without a clear definition in the relativistic regime.

This does not exhaust the problems with configuration-space representations of relativistic quantum mechanics (see Myrvold (2015) for a thorough discussion of relativistic configuration spaces in the context of wavefunction realism). But in any case these are the least of relativistic particle mechanics' problems. In the dawn of quantum mechanics it was recognised that the theory became *inconsistent* when interactions between particles were included, unless those interactions were permitted to create or destroy particles. An interim step to address this leads to *variable-particle-number* quantum mechanics, in which the Hilbert space of the theory is the direct sum of the N-particle Hilbert space for every value of N. Wavefunction realism for a theory of this kind (even setting aside spin and the ambiguity as to what the position basis is) requires an *infinite* number of configuration spaces, and a wavefunction on each, with interactions coupling the wavefunctions on different spaces — again, the position is radically transformed.

But even variable-particle-number quantum mechanics, it turned out, failed to fully realise the lesson of relativistic interactions. That lesson led physics in due course to *quantum field theory*, in which particles themselves become emergent, high-level entities: in quantum field theory, the 'right' particle description for a given system depends upon contingent facts about that system, in particular its energy density. For instance, 'the' mass of the electron in quantum field theory is not some lawlike feature of the theory, but a parameter adjusted to best fit the details of the physical situation being modelled. More radically, although popular science often presents the proton and neutron as simple agglomerations of quarks, a more accurate gloss on quark physics is that the proton/neutron description is the most perspicuous particle description of the system at low energies and gives way to the quark description at high energies. So a fundamental ontology based on the positions of particles looks forlorn in quantum field theory.

Now, quantum field theory has an *analogy* to the position representation: in some cases ('bosonic' fields, such as the electromagnetic field), the quantum state of the field theory can be represented as a wavefunction on a configuration space — albeit a space in which the points correspond (formally) to entire instaneous configurations of a field rather than to coordinates of N particles. Such a space is infinite-dimensional (and mathematically quite badly behaved) but at least *prima facie* the wavefunction realist can respond to the challenge of field theory by being a realist about the field-configuration-space wavefunction. (See Ney (2013a) for an explicit case for this approach.)

My immediate feeling about this move is: if what is really intended is a wavefunction on field configuration space, shouldn't we be discussing *that* metaphysics rather than being distracted by the red herring of wavefunctions on N-particle configuration space? Granted, the latter has the virtue of being simpler to talk about, but it has the vice of being inconsistent with our current best quantum theories, which seems more serious.

But in fact even field-configuration wavefunction realism has severe technical problems. For a start, as with relativistic quantum mechanics, it is difficult to reconcile it with relativistic symmetries; indeed, making straightforward sense of it seems to require a preferred choice of reference frame. (Admittedly, most advocates of wavefunction realism are sympathetic to resolutions of the quantum measurement problem which are already in severe tension with relativity, so this may be a fairly palatable bullet for them to bite.)

More seriously, 'the' field configuration basis is often not unambiguously defined in quantum field theory. Often the same operational content may be represented through radically different choices of field: this is the phenomenon of 'duality' (for a comparatively elementary example, see Coleman (1985, ch.6)). So the correct representation for which to express wavefunction realism is pretty radically underdetermined.

Most seriously of all, I noted that only bosonic field theories *can* be represented as wavefunctions on configuration space. Others — the 'fermionic' field theores that represent electrons and quarks (and so are central to our quantum-mechanical descriptions of ordinary matter) — possess no such representation.[4]

In conclusion, wavefunction realism seems to rely on features of toy NRQM which, far from being universal features of any realistic quantum theory, drop away as soon as we generalise. At least pending very substantial technical work, we should treat with grave scepticism any suggestion that a metaphysics based on these features of toy NRQM has any real bearing on the metaphysics of our Universe.

3 The conceptual case against wavefunction realism

Put aside all these technical objections, and consider the metaphysics of a possible world where toy NRQM is exactly true.[5] *Even so*, the move to wavefunction realism is unmotivated.

To illustrate this, consider an example from *classical* N-particle mechanics. One natural way to formulate this theory is as follows:

[4]A technical note: the possibility of a configuration-space representation in field theory relies on the fact that spacelike separated field operators commute, and so the collection of all such operators on a spatial hypersurface has a common set of eigenvalues. But in fermionic fields, spacelike separated operators *anticommute*. Introduction of Grassman numbers allows the *formal* introduction of something analogous to a configuration space representation, but it is at best unclear whether this has any significance beyond the purely calculational — at any rate, the burden of proof lies on the wavefunction realist here.

[5]I avoid saying 'suppose that we lived in such a world': it's pretty clear that toy NRQM, in which electromagnetic radiation is wholly absent, could not support complex organisms anything like us.

- The instantaneous state of the system is represented by N points X_1, $\ldots X_N$, in three-dimensional Euclidean space.
- The dynamics is given by the differential equations

$$m_k \frac{\mathrm{d}^2 X_k}{\mathrm{d}t^2} = \sum j \neq k F_{jk}(X_j, X_k) \tag{3}$$

  where m_k is the mass of the kth particle and F_{jk} is the force on the kth particle due to the jth.

It is extremely natural to interpret this as the theory of N particles moving in space. Notice that essentially all the structure of the world, according to the theory, is encoded in the mathematical structure of the state — specifically, in the distances between the various X_k. The actual value of a given X_k, in isolation, encodes no real information about the system: Euclidean space is featureless, with no point and no direction distinguished from another. Mathematically speaking, this is because the automorphism group of Euclidean space is the full three-dimensional group $\mathcal{E}(3)$ of translations and rotations: any two points are related by some translation, any two directions by some rotation.

But there is another way to represent this theory. We can define the configuration space as the product of N copies of Euclidean three-space. Each N-tuple of points $(X_1, \ldots X_N)$ now corresponds to a *single* point in this $3N$-dimensional space, and the N coupled differential equations (3) to a single differential equation on that space.

If we compare the theory in this formulation to the original formulation, we observe that:

1. The mathematical state is now completely featureless: a mere point. Any two states are intrinsically identical.
2. Conversely, the configuration space is much more highly structured than three-dimensional Euclidean space. Beyond mere dimension it has virtually nothing in common with $3N$-dimensional Euclidean space, and indeed is a 'space' only in the mathematician's sense, not in any sense based on a physical analogy with ordinary space.. The latter's structure is characterised by the $3N$-dimensional translation/rotation group $\mathcal{E}(3N)$, whereas the symmetry group of configuration space is simply $\mathcal{E}(3)$, the same as for 3-dimensional Euclidean space.[6] In geometrical terms, the coordinate-free structure of configuration space is most perspicuously specified via a preferred identification of each point in the former with N points in 3-dimensional Euclidean space.

The move to configuration space has encoded all salient features of the system via the position of a maximally simple state (a single point) in a highly structured space, rather than via an intrinsically complex state (an Ntuple of points) located in a much less structured space. There's nothing wrong with

[6] Lewis (2013) explores this point further in his discussion of wavefunction realism.

this: it's a standard move in theoretical physics. Indeed, it's a completely general move: a state-space formalism, where all states are intrinsically identical and all structural features of the world are encoded in a state's location in a highly structured state space, can be straightforwardly defined for pretty much any dynamical theory.

I draw two morals:

1. When physical theories are presented to us as formulated on spaces with much more structure than Euclidean space, we should not rush to interpret them as *physical* spaces, rather than as mathematical devices to encode information about the physical state.

2. Conversely, when a state is represented mathematically by a comparatively simple entity living in a highly structured space, we should not rush to assume that the physical world is comparably simple. Indeed, we should not rush to assume a one-to-one correspondence between mathematical states and physical entities: doing so in the case of classical mechanics would lead us to assume there is one fundamental entity, not N.

What then of quantum mechanics? The most direct analogue of the classical discussion is the *Hilbert*-space formulation of quantum theory, in which states are normalised vectors. Any two vectors are intrinsically identical (being mere lines); all the physical information about a system is encoded in the location of that vector in Hilbert space. As such, it would be naive in the extreme to be a "Hilbert-space-vector realist': to reify Hilbert space, and take it as analogous to physical space.

North (2013), in making the case for wavefunction realism, actually discusses Hilbert-space-vector realism, and of course rejects it, but her reasons are instructive: she writes that "The Hilbert space formulation seems to contain too little structure from which to construct a picture of the world as we experience it. Hilbert space does not support an objective, structural distinction between positions and other physical properties, like spin, in the way that the wavefunctions space does". But of course the Hilbert space of any particular quantum theory does indeed have enough structure to do so: any particular quantum theory is given not by a bare, unstructured Hilbert space but by that space together with an algebra of preferred observables (to which I return later). If it were not so structured, Hilbert-space quantum mechanics would be *pragmatically* unsuited to quantum mechanics, which manifestly is not the case. The reason to reject Hilbert-space-vector realism, rather, is that the space on which the Hilbert space vector is defined is much *too* highly structured to be taken as analogous to physical space: it is a state space, just as configuration space is.

The wavefunction formulation of quantum mechanics lies intermediate between the elementary formulation of classical mechanics (in which essentially all the physical structure is coded in the intrinsic properties of the state, and in which the space of the state is largely featureless) and the state-space formulations of classical and quantum mechanics (in which all the world's structure is coded via the position of the state). The wavefunction is far from featureless — a given wavefunction might be extremely highly structured — but two wavefunctions that are related by (say) an arbitrary translation or rotation on configuration space will describe radically different physics, because configuration space is also highly structured. Indeed, the structure of the wavefunction only serves to encode *quantum-mechanical* facts about a system. States corresponding to fully classical states of affairs have trivial, wave-packet wavefunctions, and again everything that distinguishes one classical state from another is encoded by the location of that wave-packet.

So reading wavefunction realism from the existence of the configuration-space representation of quantum mechanics seems unmotivated, for largely the same reasons as reading Hilbert-space-vector realism from the vector representation of quantum mechanics. In both cases, the combination of the intrinsic and locational features of the state serves to encode all the physical structure of the system in question, but there is no reason to think that the formalism transparently displays anything like the appropriate metaphysical description of the system.

So what *is* "the appropriate metaphysical description"? Here's one possibility: an N-particle quantum state is uniquely specified by assigning a complex number to every N-tuple of points of space. We could perfectly well interpret these complex numbers as relational properties of N-tuples of spatial points, irreducible to monadic properties of individual points; that's a highly nonlocal ontology, but the phenomenology of entanglement is also highly nonlocal so that looks like a feature, not a bug. I don't want to claim that this is the clear *best* metaphysical description of quantum mechanics — I seriously doubt that it is, given the technical criticisms of the previous section — but it will do as an existence proof that there are ontologies for quantum mechanics that don't regard the highly-structured configuration space as a physical space.

I have argued so far that there is no good case for simply *reading off* wavefunction realism from the quantum formalism. But can it be *argued for* as the best way to think of quantum ontology? (Here I continue to put aside purely technical objections.) North attempts to do so, but her argument is that we should prefer that ontology which has just the right level of structure to support physics, and we have seen that this fails to distinguish between wavefunction realism, Hilbert-space-vector realism and the nonlocal, three-dimensional-space-based ontology of complex N-point relations — all are structurally isomorphic, but they are sharply different metaphysically.[7]

Ney (this volume) argues for wavefunction realism on the grounds that it delivers an ontology that is both separable and local: on wavefunction realism, the state of the whole system is given by the separate wavefunction values at each point of configuration space, and the dynamics are given by configuration-space-local equations.

[7] Or maybe not. Those sympathetic to the ontic structural realism of Ladyman and Ross (2007), Saunders (2003) et al — like me — might be sceptical that there is a true distinction here. But for the purposes of this essay I assume a more straightforward metaphysics, in keeping with the presumptions of most advocates of wavefunction realism.

The case that wavefunction realism preserves locality is somewhat unclear to me. It's true that the Schrödinger equation is configuration-space local, but if wavefunction collapse is included (as in the GRW theory) then it is nonlocal even on configuration space — and if it isn't, then we are effectively assuming the Everett interpretation, which has local dynamics even on more explicitly spatiotemporal formulations (Wallace 2012, ch.8). The case for separability is considerably clearer, albeit there are separable formulations of Everettian quantum mechanics on spacetime, notably that given by Deutsch and Hayden (2000; cf discussion in Wallace and Timpson 2010).

But in any case, we need a positive argument for why separability and locality are desirable features for our ontology. And as Ney herself persuasively argues, standard arguments for wanting these features really concern separability and locality in three-dimensional space (fundamental or emergent), and not separability or locality in a high-dimensional but phenomenologically distant fundamental space. She concludes that "[t]he case for a separable and local metaphysics for quantum mechanics then comes from more broadly philosophical considerations, special relativity, perhaps brute intuition, and additionally considerations of what provides a more coherent and stable picture."

Of these, I will pass over the "broadly philosophical considerations" and "considerations of what provides a more coherent and stable picture": ultimately they are considerations to assess when comparing wavefunction realism to other concrete ontological proposals, and those lie outside the scope of this article. The appeal to brute intuition seems problematic for reasons that Ney herself again provides: why expect our intuitive faculties, evolved as they are for the emergent classical world, to track truth about fundamental metaphysics?

A conflict between special relativity and alternatives to wavefunction realism would indeed be a strong reason in favour of the position, but I'm unsure how the conflict is supposed to go. In philosophy of spacetime, 'compatibility with relativity' usually means that the theory is formulated on Minkowski spacetime with no additional structure; in mainstream physics, it more usually means that the theory has the Lorentz group as a symmetry group. But neither definition favours wavefunction realism over alternatives. As for the former: the spacetime of wavefunction realism (configuration space $\times$ time) is not Minkowski spacetime, and in an important sense is explicitly nonrelativistic, incorporating as it does a preferred sense of simultaneity (to specify a configuration is to specify the locations of several particles *at the same time*, and the formal features required to define the relation between configuration space and Minkowski spacetime continue to include this preferred simultaneity even as the interpretation of 'configuration' space changes. As for the latter: the equations of a theory have the same (dynamical) symmetry group however they are interpreted metaphysically, and so the question of whether they are Lorentz-covariant is independent of wavefunction realism. (The answer? 'Yes' for the Everett interpretation; 'No' for most proposed relativistic generalisations of Bohmian mechanics; 'Unclear' in the case of dynamical-collapse theories). If there is any further sense in which relativity favours wavefunction realism, it has not yet been developed.

I conclude that even if the particular features of toy NRQM that permit wavefunction realism to be formulated could be found in more sophisticated and realistic quantum theories, wavefunction realism would not be a well-motivated approach to the ontology of quantum theory: its supposed advantages are unpersuasive and the highly structured 'space' on which it is defined is not properly analogous to ordinary physical space.

4 Epilogue: spacetime in quantum theory

So what *is* the correct ontology of quantum theory?

I don't know. The advocates of wavefunction realism certainly deserve credit for recognising that there is a significant metaphysical question to be answered here, even if their proposed answer falls short: the ill-defined metaphysics of 'eigenvector-eigenvalue links' and 'indefinite properties' look unlikely to survive in any realist solution to the quantum measurement problem. Chris Timpson and I sketch one possibility in Wallace and Timpson (2010), and I review others in Wallace (2012, ch.8), but these are, at best, essays in the craft; more generally, I suspect that looking for 'the' ontology of a framework theory is a category error (Wallace 2017) and that we would do better to reformulate the question in terms of the ontology of specific quantum theories, such as the standard model of particle physics (and also to recognise that these are unlikely to be *fundamental* theories, so that hopes to learn about fundamental ontology from those theories are probably vain).

However, one striking theme in the formalism of pretty much every empirically successful *specific* quantum theory I know is that space (or rather, spacetime) concepts play a central role. In particle mechanics, the theory is normally defined in terms of the spacetime symmetry group (Galilean for nonrelativistic theories, Poincaré for relativistic): indeed, a one-particle quantum theory is frequently defined as an irreducible representation of the appropriate spacetime symmetry group (whatever the metaphysical status of that approach). In quantum field theory, the connection is significantly tighter: the theory's structure is specified via operator-valued fields, i. e. maps from spacetime points into the algebra of operators on Hilbert space. So dynamical quantities get their operational significance at least partly via their association with spacetime points. Indeed, at least according to the 'algebraic' formulation of quantum field theory, the operational significance of dynamical quantities is exhausted by their spacetime associations: a quantum field theory, in the algebraic setup, is specified completely by a map from spacetime regions to the algebra of dynamical variables associated to each region, with no further specification of which operator in a region corresponds to which physical value. Spacetime is thus mathematically required in the formulation of a quantum field theory.

Does this mean that spacetime is fundamental? It is too early to tell. Quantum field theories (which presume a fixed, background spacetime) must sooner or later give way to some quantum theory of gravity, and we do not yet have that theory, so metaphysical speculation about spacetime's status in it is probably premature. What we can say is that spacetime plays a fundamental, nonderivative role in our current best quantum theories, and so extant quantum physics gives no reason at all to expect its elimination.

References

Albert, D. Z. (1996). Elementary quantum metaphysics. In J. T. Cushing, A. Fine, and S. Goldstein (Eds.), *Bohmian Mechanics and Quantum Theory: An Appraisal*, Dordrecht, pp. 277–284. Kluwer Academic Publishers.

Allori, V. (2013). Primitive ontology and the structure of fundamental physical theories. In *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, pp. 58–75. New York: Oxford University Press.

Allori, V., S. Goldstein, R. Tumulka, and N. Zanghi (2008). On the common structure of Bohmian mechanics and the Ghirardi-Rimini-Weber theory. *British Journal for the Philosophy of Science 59*, 353–389.

Coleman, S. (1985). *Aspects of Symmetry*. Cambridge, UK: Cambridge University Press.

Deutsch, D. and P. Hayden (2000). Information flow in entangled quantum systems. *Proceedings of the Royal Society of London A456*, 1759–1774.

Fleming, G. N. (2000). Reeh-Schlieder meets Newton-Wigner. *Philosophy of Science 67*, S495–S515.

Fuchs, C. (2002). Quantum mechanics as quantum information (and only a little more). Available online at at http://arXiv.org/abs/quant-ph/0205039.

Fuchs, C. and A. Peres (2000). Quantum theory needs no "interpretation". *Physics Today 53*(3), 70–71.

Gleason, A. (1957). Measures on the closed subspaces of a Hilbert space. *Journal of Mathematics and Mechanics 6*, 885–893.

Halvorson, H. (2001). Reeh-Schlieder defeats Newton-Wigner: On alternative localisation schemes in relativistic quantum field theory. *Philosophy of Science 68*, 111–133.

Healey, R. (2012). Quantum theory: A pragmatist approach. *British Journal for the Philosophy of Science 63*, 729–771.

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

Ladyman, J. and D. Ross (2007). *Every Thing Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press.

Lewis, P. (2013). Dimension and illusion. In A. Ney and D. Albert (Eds.), *The Wave Function*, pp. 110–125. Oxford: Oxford University Press.

Maudlin, T. (2013). The nature of the quantum state. In A. Ney and D. Albert (Eds.), *The Wave Function*, pp. 126–153. Oxford University Press.

Monton, B. (2013). Against 3N-dimensional space. In A. Ney and D. Albert (Eds.), *The Wave Function*, pp. 154–167. Oxford: Oxford University Press.

Myrvold, W. (2015). What is a wavefunction? *Synthese 192*, 3247–3274.

Ney, A. (2013a). Introduction. In *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, pp. 1–51. Oxford: Oxford University Press.

Ney, A. (2013b). Ontological reduction and the wave function ontology. In A. Ney and D. Albert (Eds.), *The Wave Function*, pp. 168–183. Oxford: Oxford University Press.

North, J. (2013). The structure of a quantum world. In A. Ney and D. Albert (Eds.), *The Wave Function*, pp. 184–202. Oxford: Oxford University Press.

Pusey, M. F., J. Barrett, and T. Rudolph (2011). On the reality of the quantum state. *Nature Physics 8*, 476. arXiv:1111.3328v2.

Saunders, S. (1998). A dissolution of the problem of locality. *Proceedings of the Philosophy of Science Association 2*, 88–98.

Saunders, S. (2003). Physics and Leibniz's principles. In K. Brading and E. Castellani (Eds.), *Symmetries in Physics: Philosophical Reflections*, pp. 289–308. Cambridge: Cambridge University Press.

Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford: Oxford University Press.

Wallace, D. (2017). Lessons from realistic physics for the metaphysics of quantum theory. Forthcoming.

Wallace, D. and C. Timpson (2010). Quantum mechanics on spacetime I: Spacetime state realism. *British Journal for the Philosophy of Science 61*, 697–727.