

On the Plurality of Quantum Theories: Quantum theory as a framework, and its implications for the quantum measurement problem

David Wallace*

November 10, 2018

This is a draft of a chapter that has been accepted for publication by Oxford University Press in the forthcoming book 'Realism and the Quantum' edited by S. French and J. Saatsi, due for publication in 2019.

Abstract

'Quantum theory' is not a single physical theory but a framework in which many different concrete theories fit. As such, a solution to the quantum measurement problem ought to provide a recipe to interpret each such concrete theory, in a mutually consistent way. But with the exception of the Everett interpretation, the main extant solutions either try to make sense of the abstract framework as if it were concrete, or else interpret one particular quantum theory under the fiction that it is fundamental and exact. In either case, these approaches are unable to help themselves to the very theory-laden, level-relative ways in which quantum theory makes contact with experiment in mainstream physics, and so are committed to major revisionary projects which have not been carried out even in outline. As such, only the Everett interpretation is currently suited to make sense of quantum physics as we find it.

1 Introduction

Philosophy is what you do when you don't yet know what the right questions are to ask.

Daniel Dennett¹

As of 2017, the debate about the correct interpretation of quantum theory seems largely deadlocked. There is a reasonable consensus within philosophy of physics on what the main viable interpretative strategies are, which goes something like this (see subsequent discussion for references):

*Dornsife School of Letters, Arts and Sciences, University of Southern California. Email: dmwallac@usc.edu

¹In Blackmore (2005, p.91).

1. ‘Realist’ strategies, dominated by:
 - (a) Hidden-variable theories, most notably the de Broglie-Bohm theory (aka Bohmian mechanics);
 - (b) Dynamical collapse theories, most notably the Ghirardi-Rimini-Weber theory and Pearle’s continuous-state-localisation theory;
 - (c) The Everett interpretation, in its modern (decoherence-based) form.
2. ‘Non-realist’ strategies, notably (but not exhaustively):
 - (a) The (various forms of the) Copenhagen interpretation;
 - (b) Physics-is-information approaches, most notably the ‘Quantum Bayesianism’ or ‘QBism’ of Chris Fuchs and co-workers;
 - (c) Healey’s quantum pragmatism.

The debate over realist-vs-non-realist strategies has been fairly cursory in recent discussions and has largely turned on general disagreements about the legitimate aims of science: Maudlin (1995), for instance, simply takes as read that quantum theory ought to have a representational role, while according to Fuchs and Peres (2000a, 2000b) or Healey (2012, 2017b), a central lesson of quantum mechanics is that it does not have such a role and needs to be understood more as some kind of predictive or calculational tool. In the bulk of philosophy of physics (in particular in its more metaphysically inclined corners), indeed, the non-realist strategies are set aside almost without comment.²

Meanwhile, comparative assessment of the realist strategies has tended to turn on relatively detailed, and fairly metaphysical, concerns with those strategies. Can the Everett interpretation explain probability? Is Bohmian mechanics just Everett in denial? How do dynamical-collapse theories resolve the problem of tails?³ The form of the discussion normally takes as read that these various approaches would succeed in ‘solving the measurement problem’ if only these metaphysical problems could be resolved. Indeed, it is common (see, e.g., Cordero (2001), Egg (2014), Lewis (2015), Lyre (2010)) to describe the choice between realist interpretations as a classic case of *underdetermination of theory by data*, with Bohmian mechanics, the GRW theory, and (sometimes) the Everett interpretation all on a par empirically and to be distinguished only by assessments of ‘extra-empirical virtues’, again of a largely metaphysical nature: preservation of determinism, avoidance of ontological extravagance, conformity with intuition, conformity with the spirit of relativity. If so, there seems to be little realistic likelihood of consensus, or even significant progress, on the interpretation of quantum mechanics any time soon. But perhaps we can console

²Productive debate is not helped by the fact that most critics *and* defenders of what I here call ‘non-realist’ strategies regard ‘realist’ as a virtue term: witness in particular Fuchs’ (2017) plea to critics *not* to call QBism anti-realist, but also Healey’s description of his position as a ‘third way’ between realism and non-realism. I don’t myself think the label is particularly helpful (I will later argue that ‘non-representationalist’ is a more neutral description); my summation here is intended as descriptive of the sociology of the field.

³See Wallace (2008) for references in each case.

ourselves that even if the final answers to the questions of quantum interpretation will continue to elude us, at least we have a clear understanding of what the questions are and what the space of possible answers looks like.

But in philosophy, most of the work, and most of the controversy, lies precisely in stating and understanding the questions, and so apparent consensus on how to frame a problem is often a sign of hidden assumptions and communication failure.⁴ And so it is in quantum theory (I shall argue). Advocates of different interpretative strategies differ not just on technical details of those strategies, or on assessment of whether and how those strategies overcome their own specific problems, but on basic questions of philosophy of science: notably, on how theories confront experiment and on how they represent physical facts. And (relatedly) they differ on what the theory is that we are supposed to interpret: to a first approximation, advocates of different strategies are trying to interpret very different theories in the first place. Put succinctly, if the central question in philosophy of quantum theory is ‘what is the correct interpretation of quantum mechanics’, then as well as open disagreement about the assessment of ‘correct’, there is hidden disagreement about the meaning of ‘interpretation’ and, even more so, of ‘quantum mechanics’.

To expand slightly, the case I will make is that:

1. Quantum theory is a framework theory, under which many specific quantum theories stand. These theories are related by different (and only partially understood) instantiation relations (where the instantiation relation holds between two specific quantum theories, describing the same system at different levels of detail) but a satisfactory interpretation of quantum mechanics must be a strategy for interpreting each of these theories in its own terms. No satisfactory interpretation can be an interpretation of the ‘fundamental’ quantum theory, partly because we do not have any such theory to interpret at present but mostly because the way in which quantum theory makes contact with experiment is through this plurality of different theories and cannot be cleanly described in the vocabulary of any one quantum theory, however fundamental.
2. Neither the so-called ‘non-realist’ interpretations (I will argue that ‘non-representationalist’ is a fairer name), nor the strategies that aim to modify or supplement quantum theory, adequately deal with the theory as it is actually used and tested in science. Rather, they attempt to interpret either the abstract framework (in a way which is not at present adequate to recover the empirical success of particular instances of that framework), or else they interpret a particular theory within that framework, under the fiction that it is by itself adequate to all quantum phenomena (and

⁴For examples elsewhere in philosophy, consider free will, or the mind-body problem. In the former case, careful formal proofs to the effect that we do or do not have free will tend to pack most of the philosophical work into apparently-uncontentious premises (Dennett 1984, p.3); in the latter case, even deciding to frame the issues in terms of an ‘easy’ and ‘hard’ problem (Chalmers 1995) concedes most of the ground on which hardline functionalists will want to fight (Dennett 2005, ch.6).

as such rely on a once-and-for-all account of how quantum theories make contact with empirical data which is again not adequate to recover the full empirical successes of quantum theory). Only the Everett interpretation attempts to provide a general recipe to interpret the various particular quantum theories in a way which is compatible with their interrelations.

These conclusions are not neutral: they give strong support for the Everett interpretation, of which I have been a longstanding advocate. That's not a coincidence: the paper is a codification and development of what I have long thought to be the strongest reasons for that interpretation (building on preliminary versions of these arguments presented in Wallace (2008, pp.83-85) and Wallace (2012, pp.33-35)). But I hope even readers sceptical of this particular conclusion might be persuadable of the general point that the deadlock in discussions of interpretation can be broken — or at least, our understanding of the problem can be deepened — by focussing less on specifics of the particular interpretative strategies, and more on what we should reasonably require of a solution to the measurement problem in the first place.

The structure of the paper is as follows: in sections 2 I develop the idea of quantum theory as a framework, and in section 3 I consider how to think about inter-theoretic relations between concrete theories in that framework. In sections 4–6 I consider, sequentially, Everettian quantum mechanics, so-called non-realist strategies, and strategies that try to supplement or modify the quantum formalism, and explore how each fits into the conception of quantum mechanics developed in sections 2–3. Section 7 is the conclusion. The physics I discuss in the paper is for the most part standard and well-established, and I do not attempt to give original references.

2 Quantum theory: frameworks and instances

Asked what ‘quantum theory’ is, two initially-plausible answers might be:

Abstract quantum mechanics: A quantum theory is specified by:

1. A Hilbert space \mathcal{H} (the rays of which represent possible states of the system);
2. A collection of self-adjoint operators on \mathcal{H} which represent the *observables* of the system (this is a term of art; a more neutral term might be ‘physical quantities’).
3. A preferred observable \hat{H} , the *Hamiltonian*, which generates the system’s dynamics via the *Schrödinger equation*

$$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = \hat{H} |\psi(t)\rangle. \quad (1)$$

Physical content is extracted from the theory by the *Born rule*, which states that the expectation value, on measurement, of the observable

corresponding to operator \widehat{O} , if the system's quantum state is $|\psi\rangle$ is

$$\langle \widehat{O} \rangle_\psi = \langle \psi | \widehat{O} | \psi \rangle. \quad (2)$$

Quantum particle mechanics: The quantum theory of N particles of mass m_1, \dots, m_n is specified by a function ψ from the $3N$ -dimensional configuration space of those N particles to the complex numbers, satisfying the *Schrödinger equation*

$$i\hbar \frac{\partial \psi}{\partial t}(\mathbf{x}, t) = - \sum_{1 \leq i \leq N} \frac{\hbar^2}{2m_i} \nabla_i^2 \psi(\mathbf{x}, t) + \sum_{1 \leq i < j \leq N} V_{ij}(|\mathbf{x}_i - \mathbf{x}_j|) \psi(\mathbf{x}, t) \quad (3)$$

where V_{ij} is the interaction energy between particles i and j , \mathbf{x} represents, schematically, the N -tuple of coordinates $\mathbf{x}_1, \dots, \mathbf{x}_N$ on configuration space, and ∇_i is the gradient with respect to \mathbf{x}_i . Physical content is extracted from the theory by the *Born rule*, which states that the probability density, on measurement, of finding the particles in configuration \mathbf{x} at time t is

$$\text{Pr}(\mathbf{x}; t) = |\psi(\mathbf{x}, t)|^2. \quad (4)$$

As accounts of quantum mechanics as a scientific theory, both accounts are for different reasons very deficient (even before we address more philosophical concerns like the quantum measurement problem). To begin with ‘abstract quantum mechanics’: in an important sense this is not really a scientific theory at all. By itself it makes no predictions and explains no phenomena; by itself it cannot be tested or falsified. It is a *framework* within which concrete quantum theories can be stated: a given theory within the framework is specified (typically) by a specific choice of Hilbert space, a specific choice of dynamical variables (represented by self-adjoint operators), a specific dynamics (represented by some choice of Hamiltonian) and often a specific decomposition into subsystems. These *specific, concrete* theories can of course be tested, and those tests can indirectly confirm or falsify the viability of the framework, but without being supplemented by a particular concrete realisation of the framework, it can do only very little scientific work.

When we recognise this, we can also see what is deficient about ‘ N -particle quantum mechanics’: it is only one example of a quantum theory, and one of quite restricted applicability. Indeed, strictly speaking there is to my knowledge only one small class of applications of wave mechanics in this form: to the hydrogen atom, and to other ions with only one electron. If the theory is modified to include identical particles, background potentials, and spin, its applicability becomes much wider: most of atomic physics, and most of quantum chemistry, and a large part of solid state physics falls within its remit, taking the ‘particles’ to be electrons and the nuclei of atoms. But even there, the theory rarely suffices to describe *all* the relevant phenomena of a system: for instance, atomic energy levels are calculated within this theory, but the theory cannot incorporate the

de-excitation of those energy levels by photon emission which is the main route by which we test the theory's energy-level predictions.⁵

A really satisfactory description of 'quantum mechanics' would have to characterise it as a large *collection* of theories, each fitting (more or less) within the abstract framework, and each applicable to different systems and at different energy- and length-scales. One of those theories would be N -particle quantum mechanics. Or more accurately: *many* of those theories would be different forms of N -particle quantum mechanics, with different numbers and statistics of particles, different background potentials, and different interpretations of the 'particles'. (Sometimes they are electrons, sometimes ions, sometimes the centres of mass of larger bodies — even (Zurek and Paz 1995) moons or planets.) Others would be quantum field theories of various kinds, applicable to various situations involving light, collective excitations of solid bodies, or relativistic effects. Others still would be discrete theories with finite-dimensional Hilbert spaces, applicable to the internal degrees of freedom of certain systems whose spatial degrees of freedom can be idealised away.

It should not be surprising to find that 'quantum theory' needs to be characterised this way. After all, the same is true for 'classical mechanics'. There we could have tried to characterise the theory as the abstract form of Hamiltonian or Lagrangian mechanics, or as the theory of point particles interacting under long-range forces. There too, the first characterises a framework for dynamical theories rather than a concrete theory; the latter characterises a theory which captures only a small part of classical mechanics (failing to capture, for instance, electromagnetic effects, or the physics of rigid bodies or fluids). And classical particle mechanics, like quantum particle mechanics, is best thought of not as a single theory but as a wide class of different theories, characterised by different particle masses, interactions, background potentials, and interpretations of the 'particles'. (Recall that in the most famous application of classical mechanics, the 'particles' are the centers of masses of the Sun and planets.) If we were to set out to 'interpret classical mechanics', our task would not be completed *either* by trying to interpret the abstract form of Hamiltonian mechanics as a concrete theory (which would simply be a category error), *or* by trying to interpret classical particle mechanics (which has no univocal interpretation, and in any case is only a small part of classical mechanics).⁶

At least *prima facie*, the same would seem to apply to quantum theory. Before developing this point, though, I need to develop further the picture of quantum theory I have sketched by considering how different versions of quantum theory relate to one another.

⁵The theory is also somewhat misrepresented in the form I have given it by its focus on configuration space and on the position representation, but that can more readily be fixed.

⁶Mark Wilson has developed (essentially) this view of classical mechanics in much more depth; see, in particular, Wilson (2013).

3 Quantum theory: inter-theoretic relations

A tempting and popular picture of inter-theoretic relations is that of a tower of theories, each approximating the theory below it in the appropriate limit.⁷ For physics, at the bottom of the tower would lie the Standard Model of particle physics (perhaps with its base shrouded in mist to leave room for the hoped-for theory of quantum gravity that *it* approximates). Above it, perhaps, would be quantum electrodynamics; above that, the quantum theory of photons and nonrelativistic atoms; above that, nonrelativistic quantum mechanics; above *that*, perhaps, classical particle mechanics, and then classical fluid mechanics.

But this picture is badly misleading. It implies there is some unified class of phenomena all simultaneously describable by the Standard Model, some subclass of that class simultaneously describable by quantum electrodynamics, and so forth. In reality in physics, modelling is local. A given system — say, the electrons of a metal, or the electrons and nuclei of an ion — might be describable by quantum particle mechanics, but it does not follow that in anything but the most formal sense there is a single system comprising both systems that is describable by it. Indeed, sometimes the same system is described by the same theory at different levels: for instance, the electrons in the metal might be treated as evolving against a classical background potential which, at a deeper level of description, is produced by a lattice of ions. Nor is there any consistent hierarchy of theories: in classical mechanics, for instance, a collection of lattices of masses connected by stiff springs could be describable at one level by point-particle mechanics (with the ‘particles’ being the lattice elements), at a higher level by rigid-body mechanics (with the ‘bodies’ being the lattices), and at a higher level still by point-particle mechanics again (with the ‘particles’ being the centers of masses of the lattices).

So a better picture is more like a patchwork than a tower: for any given system there are various levels of description at which various theories are applicable. And since the notion of ‘system’ is itself theory-laden, there is no theory-free starting language with which we can describe this picture. For instance, one ‘system’ is comprised of the centers of masses of the bodies in the Solar system, but ‘center of mass’ is itself theory-laden and theory-relative. (Again, see Wilson *ibid* for further development of this point in the classical context.)

What do inter-theoretic relationships look like on *this* picture? If they exist, they would have to look like this: if theory X describes a system at some level of description, and theory Y describes that same system at a more detailed level, then the description from theory X can in some sense be derived from that of theory Y. And if several systems S_1 through S_n can be described separately by theory X at a higher level, and collectively by theory Y at a lower level, we can derive from theory Y both the applicability of X to the systems separately and the validity, at that level, of the decomposition into subsystems. Ultimately we might hope to find a sufficiently fine-grained level of description at which the

⁷The classic statement of this view is Oppenheim and Putnam (1958).

whole Universe can be jointly regarded as a single system, and the applicabilities of all of the higher-level theories derived directly or indirectly from it.

It is not *a priori* obvious that things must work out this way: Nature could be inherently disunified, with many different descriptions on different length-scales and no systematic relations between them. Cartwright (1983, 1999) has argued, based on careful study of case studies in physics and other areas of science, that indeed this is the case. But (although there is far more to say here than space permits) this kind of disunified picture — whatever its metaphysical coherence or incoherence — does not seem to give adequate weight to the great many cases in which physicists have actually succeeded in constructing local inter-level relations. Most of solid state physics, for instance, and most of non-equilibrium statistical mechanics, and much of particle physics and astrophysics, is concerned with deriving higher-level descriptions from lower-level ones, and there are a great many successes that have led to novel empirical predictions, albeit they often involve approximations and assumptions, and contain conceptual puzzles of their own. And the progress of physics does seem to be a progression towards greater unification: the many different solid-state systems are all instantiations of the same nonrelativistic quantum theory; the physics of stellar interiors and of stellar atmospheres seem to be derivable from the same underlying physics under different assumptions. We have reached the point where one theory, the standard model of particle physics (with the spacetime metric treated as one more quantum field) is at least a candidate to underly all the various applications of high-level physics, and to provide the basis for explanation of all physical phenomena outside the extremes of the early universe and the singularities within black holes. So notwithstanding Cartwright's criticisms, I will continue for this article to assume — tentatively — that the various applications of physics are indeed interrelated by locally-valid derivations, and ultimately that all can be underpinned by the Standard Model.

If we want to interpret quantum mechanics, then, should we simply interpret the Standard Model, and regard every other quantum theory as derivative upon it, or even just as useful calculational tools? There are two problems with so doing. Firstly, almost all our evidence for quantum mechanics is in the first instance evidence for some higher-level quantum theory, not for the Standard Model directly. In most concrete applications we do not have an explicit derivation of that higher-level theory from the Standard Model; even in the most favourable cases, we have a long and indirect chain of derivations. As such, the way in which the Standard Model confronts experiment is almost exclusively indirectly, via tests of higher-level theories.

For instance, quantum theory allows us to calculate the heat capacity of a crystalline insulating solid at low temperatures: it scales with the fourth power of the temperature. That prediction cannot even be stated, much less derived, in the Standard Model, where 'crystalline insulating solid' is not a well-defined term. Other quantum predictions concern the spectral lines of excited atoms, or the conductivity of metals, or the phase transitions of superfluids, or the neutrino emissions of stars; again, these cannot directly be stated, let alone tested, within the Standard Model. To understand quantum mechanics

well enough to recover its empirical success, we need to understand the various different quantum theories and their theoretical content. That understanding might in principle be derivative on some understanding of the Standard Model, but it cannot simply be skipped — not if our goal is to understand quantum mechanics as it is in fact used and tested.

Secondly, the standard model — like almost all empirically-relevant quantum field theories — is an *effective field theory*. What this means, in outline (see Wallace (2011, 2017) for a more detailed discussion) is that while the theory is *formally* a continuum field theory defined on arbitrarily short lengthscales, in fact it must be ‘cut off’ below some short lengthscale, which represents the point at which the theory breaks down and ceases to be accurate. (For the Standard Model, defined as I have to include the spacetime metric, that lengthscale is the Planck length.) Effective field theories are by their nature not candidates for a fundamental theory; indeed, they are not really theories at all in the philosopher’s usual sense, but equivalence classes of theories, with different cutoffs and different implementations of the cutoff, but the same larger-scale structure. Insofar as the interpretation of ‘quantum mechanics’ ought to be concerned with a single well-defined, universal theory, that theory cannot be the Standard Model but would have to be the yet-unknown theory of quantum gravity which, we hope, underlies it. But since we do not have that theory, it would be premature to try to interpret it, and more premature still to conflate the task of understanding that theory with that of making sense of quantum theory.

4 Interpretative recipes and the Everett interpretation

Given this account of quantum theory — or rather, given this account of the *many* quantum theories and the system-local, not-fully-hierarchical relations between them — what should we expect from an ‘interpretation’ of quantum mechanics? Here is one natural answer: we should expect an *interpretative recipe*, a set of instructions which tells us, for any given quantum theory, how to understand *that* theory. Furthermore, the recipe must be compatible with inter-theoretic relations: if theory X instantiates theory Y in certain regimes, our understanding of theory X ought to tell us how to understand theory Y , and it should deliver the *same* understanding as we would get by applying the recipe to Y directly. (I use ‘understanding’ here as a term neutral between pure interpretations of the quantum formalism, and strategies which supplement it with additional variables or modify its equations.)

The interpretation of *classical* mechanics has exactly this ‘recipe’ form. The interpretative recipe says, simply: interpret the classical phase space as a space of possible states that a physical system is in, so that a system with dynamical history $x(t)$ has, at time t , the physical features represented by $x(t)$. At this level of abstraction, there is little more to be said: in particular, it would

be a category error to ask for ‘the’ ontology of classical mechanics. For any *particular* classical theory, we can ask for the ontology of *that* theory (and at that level, there is space for controversy as to interpretative matters even in classical physics: the substantialist/relationist dispute is such a controversy, for instance). But that answer will have to be fairly structuralist, because the ontology of a theory instantiated by another theory must be compatible with the ontology of the latter theory, and inter-theoretic relations in classical mechanics are cashed out in terms of structure and dynamics and don’t fit into (for instance) any standard mereological form. If there were a classical *ur*-theory underlying all other classical theories, we could regard ontological questions about that particular theory as questions of ‘fundamental classical ontology’ and see the ontologies of all other classical theories as derivative from it; but there is no such theory, and so *all* questions of ontology in classical physics will have a somewhat local, scale-dependent nature. (For further development of these points see Wallace (2016a).)

Turning back to quantum mechanics, the Everett interpretation also has the recipe form. To be more specific, by ‘the Everett interpretation’ I have in mind here the interpretation in its modern form, with decoherence understood as providing an emergent branching structure for the macroscopic degrees of freedom, with higher-level ontology understood in structural terms, and with no modification of the quantum formalism, as developed by, e.g., Zeh (1973, 1993), Zurek 1991, 1998), Gell-Mann and Hartle (1989, 1993), Saunders 1993, 1995, 1997, and myself (Wallace 2003, 2010, 2012). (The extent to which this conforms to Everett’s own views is moot, though see Bevers (2011), Barrett (2011b, 2011a), and the commentary in Everett (2012) for some considerations.) Indeed, at this level of abstraction the Everett interpretation is pretty much the same as the interpretation of classical mechanics: the formalism is left unchanged, the state space of a system is interpreted as a space of possible physical states for that system, and the evolution of the quantum state under the Schrödinger equation is interpreted as describing the change over time of the system’s physical properties. The ‘emergence of the classical world from quantum mechanics’, in any particular physical situation, is just a special case of inter-theoretic reduction. (And to ask about ‘the’ ontology of Everettian quantum mechanics is again to commit a category error.)

Of course, this level of abstraction hides subtlety and controversy. Applied to microscopic systems, the Everett interpretation delivers an ontology which is not readily described in the categories of ordinary experience. Applied to macroscopic systems, it describes the state as representing not one, but an indeterminately-large multiplicity, of classical states (and it represents each of them only approximately and emergently). And to connect to experience in the latter case, we must interpret the branch-weights assigned to the classical states probabilistically. Each of these contains philosophical subtleties sufficient to allow a skeptic to reject the interpretation, even before any weight is given to the ontological extravagance of Everett-interpreted quantum theory.

But this is not the place to explore the viability of the Everett interpretation. (I defend it *in extenso* in Wallace (2012).) My point here is more modest:

because the Everett interpretation has the recipe form, it is at least a *candidate* for an interpretation of quantum mechanics suitable for physics as we find it. Among the commonly-discussed alternatives to the Everett interpretation, none have the recipe form, and so — I will argue — none is actually in a position to explain quantum theory as it is in fact used in physics. I begin with the so-called ‘non-realist’ strategies.

5 Non-representational approaches to quantum theory

As a starting point to understand the common core of ‘pragmatist’, ‘Copenhagen’ and ‘information-based’ approaches to quantum theory, consider the classic Schrödinger-cat state

$$\alpha |\text{live cat}\rangle + \beta |\text{dead cat}\rangle \tag{5}$$

which unitary quantum theory can straightforwardly produce. If the quantum state can be understood *representationally* — that is, if distinct quantum states correspond to distinct objective ways a physical system can be — and if the theory is unsupplemented by hidden variables, then it looks as if such a state must somehow represent a cat that is simultaneously alive and dead, or perhaps neither alive nor dead. Quite apart from the weirdness of such a thing, it seems in conflict with what quantum mechanics itself tells us, via the Born rule, that we should expect to observe given a system in such a state: namely, either a live cat, or a dead one, and with probabilities $|\alpha|^2$, $|\beta|^2$ respectively.

But there are other roles than representation for a piece of mathematical formalism to play. Consider the probability distributions of classical statistical mechanics, for instance: mathematically they are functions on phase space but different such functions correspond not to different states of the actual world, but to different probability distributions over the possible states of the world (however the notion of probability is to be understood here).⁸ And in particular, distributions can have nontrivial support on regions of phase space corresponding to macroscopically distinct states of affairs: for instance, the probability distribution that statistical mechanics assigns to a ferromagnet cooled below its critical temperature is, in the first instance, an equally-weighted sum of distributions describing the magnet with magnetisation in all possible directions. But this does not represent a ferromagnet in some weirdly indefinite state of magnetisation, but merely a magnet with equal probability of being magnetised in each possible direction. Non-representational strategies take the quantum state as playing a role at least somewhat akin to the statistical-mechanical distribution:

⁸I have in mind here classical statistical mechanics as it was understood before the quantum revolution: that is, on the false view that classical physics was exact. I argue elsewhere (Wallace 2016b) that even supposedly-classical systems in the actual universe should not be understood in purely classical terms: rather, they should be understood as quantum-mechanical states (mixed or pure) in the decoherent limit.

as encoding the various probabilities associated to a physical system. Different strategies differ in the way they understand those probabilities: as encoding our information about a system, or our partial beliefs about it, for instance, or as providing a practical guide to how to use and interact the system. And just as it would be misleading to call classical statistical mechanics non-realist simply because the distribution function does not play a representational role, so would it be misleading to call these approaches to quantum theory non-realist *simply* because in those theories the quantum state does not play a representational role either.

However, we can reasonably ask: if the quantum state does not represent the physical state of a system, what — if anything — does? Classical mechanics provides a straightforward answer to its analogous question: the *points* of phase space represent physical states, and the classical observables — those functions that correspond to position, momentum and the like — represent the physical properties of those states. In doing so, classical statistical mechanics also offers a straightforward answer to the question of what the probabilities are probabilities *of*: they are the probabilities of the system being in one physical state or another. And this representational story will vary from one classical-mechanical system to another, and in each case the theory itself will provide the representational machinery. In this way, an ‘interpretation’ of classical statistical mechanics in terms of probability distributions straightforwardly is compatible with the idea of classical statistical mechanics as a whole being a framework theory: it provides, for each concrete instantiation of statistical mechanics, an interpretation of the distribution function in that instantiation. In the classical dilute gas it represents probabilities for the particles to have various positions and momenta; in a classical model of the ferromagnet it represents probabilities for the individual magnetic atoms to have one orientation or the other; and so forth through the indefinitely many physical systems to which classical statistical mechanics can be applied. (Note that this interpretation of classical statistical mechanics has the ‘recipe form’ which I discussed in section 4.)

In the dawn of quantum theory it was possible to imagine essentially the same story playing out. Quantum systems are equipped with algebras of ‘observables’ that correspond to the dynamical variables that describe a physical system, and the quantum state provides — and, conversely, is completely characterised by its description as — a probability distribution over those variables. So it would be extremely natural to suppose that the physical state of any quantum system is given by some definite value of each of its observables and that the quantum state is a probability distribution over those physical states. Such an interpretation would seem to be a very natural generalisation of classical statistical mechanics (and, again, to have the recipe form). A shame, then, that such an interpretation is impossible: von Neumann’s original no-hidden-variable theorem (von Neumann 1955) shows that it cannot be done straightforwardly, and the Kochen-Specker theorem (Kochen and Specker 1967) rules it out pretty much completely. So any viable non-representational account of quantum mechanics will owe us another account of what the probabilities encoded in the quantum state are probabilities of, of what the *physical* features of systems are.

Or, put in more pragmatic terms: what the non-quantum features of a system are such that the quantum state is a tool for answering questions about those features.

In the heyday of the Copenhagen interpretation, the standard answer (made explicit in, e. g., Landau and Lifshitz (1977)) was: the physical description of a system is a classical description; quantum mechanics cannot be understood except against a background of classical mechanics. This line is quite rarely defended in modern physics,⁹ for familiar reasons: separating off ‘classical’ from ‘quantum’ physics is not at all trivial in an era where the quantum-classical transition is a major topic of theoretical and empirical study, where the classification of a system as ‘classical’ or ‘quantum’ has become a matter of degree, and where the language of experimental physics is rich with terms — ‘laser’, ‘superconductor’, ‘LCD’ whose very definition is quantum-mechanical.

Richard Healey (2017a), in discussing his ‘pragmatist’ interpretation of quantum theory (developed in more detail in Healey ((2012, 2017b), this volume), is keenly aware of this issue:

A successful interpretation must explain how quantum mechanics may be formulated as a precise physical theory and unambiguously applied to real-life physical situations . . . by applying quantum mechanics we become able better to describe and represent those situations in non-quantum terms. I say non-quantum rather than classical to acknowledge that the progress of science naturally introduces novel language to describe or represent the world (Bose-Einstein condensate, Mott insulator, quark-gluon plasma). (Healey 2017a)

But *how* does science ‘naturally introduce’ this novel language? Take ‘quark-gluon plasma’, for instance. I can give you a verbal gloss on what that means — it’s a state of matter so hot that protons and neutrons — normally made up of three tightly-bound quarks — break up into individual quarks. But that verbal gloss isn’t the real physical description — it is a metaphor at best, a fiction at worst. Let’s see what happens if we ask for a proper (albeit still a bit simplified) account of ‘quark-gluon plasma’.

Q: What’s the quark-gluon plasma?

A: It’s the state of a quantum-chromodynamics (QCD) system above a certain temperature, at which a phase transition occurs to a state where the fermionic elementary excitations are associated to the quark field rather than to color-neutral products of that field.

Q: Slow down. What’s ‘temperature’ in QCD?

A: A quantum system, including a field-theoretic system, is at (canonical) thermal equilibrium when its quantum state is

$$\rho(\beta) \propto \exp(-\beta\widehat{H}) \tag{6}$$

⁹A notable exception is Peres (1993).

where \hat{H} is the Hamiltonian and β is a real number. For a system at thermal equilibrium — or that is reasonably close to thermal equilibrium — its temperature T is given by $\beta = 1/k_B T$.

Q: And what's an 'elementary excitation'?

A: Generally in quantum field theory, we can analyse systems in states reasonably close to the thermal equilibrium state as gases of weakly-interacting particles. Those weakly-interacting particles are the elementary excitations.

Q: 'Particles' as in classical point particles?

A: Not really. 'Particles' as in subsystems whose Hilbert space bears an irreducible representation of the Poincaré group, at least in the interaction-free limit.

Q: So the quark-gluon plasma is associated with one sort of particle, colder systems with another. Shouldn't I be able to say what the 'particles' are once-and-for-all?

A: Not in quantum field theory: the optimal choice of particle depends on the state of the system. Hot systems are described most naturally in terms of quarks, colder systems, in terms of protons and neutrons.

Q: Can't I just think of a proton or neutron as an agglomeration of three quarks?

A: Only heuristically. The more precise way to explain the relation between the protons and quarks is at the field level: the proton is associated with a certain triple product of the quark field.

Q: How is a particle supposed to be associated with a field?

A: If a quantum system is in thermal-equilibrium state $\rho(\beta)$, the 'two-point function' of that system with respect to field $\hat{\phi}(x)$ and that state is

$$G_2(x - y; \phi, \beta) = \text{Tr}(\rho(\beta)\hat{\phi}(x)\hat{\phi}(y)). \quad (7)$$

If the Fourier transform of that state has a pole, there's a particle associated with it.

Q: That's a weird postulate.

A: It's not a postulate; it's something you derive, by looking at the dynamics of states obtained by excitations of the thermal-equilibrium state. Where there's a pole, there's a subspace of states which can be interpreted as superpositions of singly-localised excitations and which is preserved under the dynamics.¹⁰

¹⁰Experts will recognise this as (a corollary of) the Callan-Lehmann representation of the two-point function; see, e.g., Peskin and Schroeder (1995, ch.7) for the details.

Q: What’s a ‘plasma’, anyway?

A: In ‘quark-gluon plasma’ it’s a bit metaphorical. This phase of QCD has a lot in common with ordinary electromagnetic plasmas, which are characterised by screening of the Coulomb force, or equivalently by the photon acquiring mass.

Q: I thought the photon was massless?

A: The mass of a particle depends on the quantum state. If the state is the vacuum or some small excitation of the vacuum, then yes, photons are massless. But in a state that’s an excitation of a hot, high-density system of protons and electrons, the photon acquires mass.

Q: What do you mean, ‘acquires mass’?

A: It’s a dynamical statement about elementary excitations again, derived from looking at the plasma dynamics. There’s again a formal statement in terms of poles of two-point functions.

It doesn’t matter if you followed all of the details of the above; they’re not the main point. (I could have carried out essentially the same exercise with ‘Bose-Einstein condensate’ or ‘Mott insulator’.) The main point is that I have *not the faintest idea* how to make sense of any of this without taking the quantum state of the QCD system, and its dynamical evolution under the Schrödinger equation, as representational. (Even the claim that the system has temperature T is a claim about its state.) I don’t know how to begin eliminating representational uses of the state from my account of the quark-gluon plasma — and, to the best of my knowledge, neither does Healey.

If Healey departs from Copenhagen-style use of classical physics in order to seek a broader conception of objectivity in physics, the ‘quantum Bayesianism’ (or QBism) of Chris Fuchs and co-workers moves radically in the opposite direction:¹¹

the primitive concept of experience is fundamental to an understanding of science. According to QBism, quantum mechanics is a tool anyone can use to evaluate, on the basis of ones past experience, ones probabilistic expectations for ones subsequent experience. (Fuchs, Mermin, and Schack 2014)

Taken at its most straightforward, quotes like this seem to suggest that the probabilities of quantum theory are probabilities of agents having certain experiences. It can be tempting for more conventionally-inclined physicists or philosophers to argue against such a view on the grounds that science should give us something more objective, less apparently solipsistic — but the real place to object is more mundane. Namely: the formalism of any particular quantum

¹¹For presentations, see, e.g., Caves, Fuchs, and Schack (2002), Fuchs (2002) or Fuchs, Mermin, and Schack (2014).

theory doesn't say anything *about* experiences. In any concrete instantiation of quantum theory, the observables over which probability distributions are defined are particle positions, field strengths, collective spins and the like. The only way to say anything non-circularly about an agent's experience in quantum mechanics is to characterise it externally, as an experience *of* something describable in a more physical language. And then the problems confronted by the Copenhagen strategy, and by Healey, reappear.

It's reasonable then to ask how it is that Fuchs *et al* (and Healey, for that matter) actually manages to develop and apply a non-representational version of quantum mechanics. (And it would be unfair not to recognise that they do indeed demonstrate a significant number of applications). The answer is that the explanations they provide of quantum phenomena are (I think without exception) explanations of features of quantum theory in the abstract — Bell inequality violations, EPR effects, quantum state tomography and the like. They treat measurement as primitive and make no use of details of any particular quantum system (sometimes a two-state system is described using spin variables but even then there is no particular connection to the physics of spin).

Non-representational strategies, we might say, provide an interpretation of the *abstract framework* of quantum theory. They do not appear to provide a method — at least at present — to interpret any particular instance of that framework. Explanations of, say, superconductivity, or the heat capacity of crystals, or the thermodynamic features of the quark-gluon plasma, or the colour of gold, or any of the thousands of concrete applications of quantum theory that form its real empirical base, seem out of reach for QBism, or for pragmatism, at least as they are currently stated. From this perspective it is unsurprising to find that advocacy of these approaches within physics is predominantly found in quantum information, a field whose power indeed comes from studying those features of quantum theory that are common to all instances of the theory, rather than those physical explanations which are irreducibly quantum but which require engagement with the features of one or other specific quantum theory.

Let me close this section with an observation drawn from more general philosophy of science. The non-representationalist strategy — interpreting some large part of the content of physics not as representing things in themselves but as an inferential or pragmatic tool used in describing some smaller class of physically-real phenomena — is not new, nor is it specific to quantum theory. It is, rather, the central idea in the logical-positivist and logical-empiricist pictures of science (see Creath (2017) and references therein for historical details). There, we make a principled distinction between the 'observation language' in which our observations are described, and the 'theory language' in which the non-observational parts of our scientific theories are stated. The observation language is to be understood literally, but the theory language is not; rather, it is an inferential or pragmatic tool to help us derive truths stated in the observation language.

It is almost universally accepted today that these approaches are not viable. But the predominant reason, historically, that they fell from grace was

not some awakening realisation among philosophers of science that scientific theories ‘should’ give a more robust account of the unobservable.¹² It was the increasingly clear realisation — notably (though by no means exclusively) by Kuhn (1962) and Quine (1951) and in the recognised failure of Carnap’s project in the *Aufbau* (1928) — that observation is theory-laden, that there *is* no clean separation of the vocabulary of science into a part which represents our observations, or the macroworld, or ‘medium-size dry goods’, and a part which does not. And so no actual scientific theory can be analysed as the logical positivists propose.

Non-representationalist strategies at least seem to be committed to making the same division, whether the analog of the ‘observation language’ is Copenhagen’s use of classical mechanics, or pragmatism’s ‘non-quantum’ description, or QBism’s appeal to direct experience. The problem with these approaches, as with positivism, is not that making such a distinction is *unreasonable* or *illegitimate*, but that — at least at present — we do not know how to do it.

6 Modifications of quantum theory

Non-representational strategies continue to have a significant following among physicists, but philosophers have generally given them short shrift. Their general attitude is well-captured by Tim Maudlin’s paraphrase of Hume:

[W]e can be clear on the questions that must be asked of an interpretation. Is it an additional variables interpretation whose dynamics guarantee solutions to the problem of statistics and the problem of effect? Is it a collapse theory that leads to appropriate outcome states with the right probabilities, and whose fundamental terms all have clear physical significance? If the answer in each case is “no”, then commit it to the flames, for it can contain nothing but sophistry and illusion. (Maudlin 1995)

To a substantial degree, philosophers of quantum mechanics have followed Maudlin’s line, focussing on solutions to the measurement problem that either (a) augment quantum theory with additional ‘hidden’ variables whose task it is to represent the physical world, or (b) modify the Schrödinger equation to introduce ‘dynamical collapse’ and so remove macroscopic superpositions (there is, perhaps, increasing willingness to consider the Everett interpretation alongside these two strategies).

In principle, either strategy could fit the ‘recipe model’ I laid out in section 4. A hidden-variable theory would be a fairly general recipe to assign additional variables to a quantum system, presumably picked out in some way in terms of the observables and the dynamics of that system; a dynamical collapse theory

¹²Which is not to say that arguments of this sort played no role in the fall of logical positivism: Putnam (1962), for one, gave arguments of this kind. (I am grateful to an anonymous referee for the observation.)

would be a recipe for modifying the Schrödinger equation given the same inputs. For the recipe to be consistent, it would have to behave properly in cases of inter-theoretic reduction: if $h(X)$ is the hidden-variable theory associated with quantum theory X , and some particular quantum theory $X1$ instantiates higher-level quantum theory $X2$ in some regime, then hidden-variable theory $h(X2)$ would have to be derivable from $h(X1)$ in the same regime (and similarly for dynamical-collapse theories). For the recipe to solve the measurement problem, it would have to handle the quantum-classical transition correctly: in those regimes where quantum theories reduce *formally* to classical theories, the hidden-variable or dynamical-collapse theory would need to reduce to the appropriate classical theory in a way which properly predicts unique classical outcomes.

There is even a concrete class of hidden-variable theories that actually aspire to realising this picture (I don't know of any dynamical-collapse strategy which aims to do so). The *modal* interpretation¹³ really does give a general recipe for hidden variables, applicable to any quantum system, and which at first sight does reproduce definite outcomes in the classical limit. The generally (not universally) accepted *failure* of that approach, conversely, can be attributed in large part to a mixture of (a) its failure to tell a fully consistent story at different levels of description (the hidden variables depend sensitively on the details of the system/subsystem decomposition; (Arntzenius (1990, 1998), (Clifton 1996)) and (b) the failure of its account of the quantum-classical transition to handle the necessary approximations involved in that transition (Bacciagaluppi, Donald, and Vermaas (1995), Donald (1998), Bacciagaluppi (2000)). (See Wallace (2008, s.2.6.3) for further discussion.)

But the bulk of discussion of dynamical-collapse and hidden-variable theories takes a very different form. It is almost entirely concerned with theories which supplement or modify one *specific* quantum theory: non-relativistic particle mechanics. Indeed, it is almost entirely concerned with two specific such theories: de Broglies and Bohm's pilot-wave theory, aka Bohmian mechanics (Bohm 1952; Bell 1987; Bohm and Hiley 1993) and the 'GRW' collapse theory (Ghirardi, Rimini, and Weber 1986), sometimes with lip service paid to the need to modify the latter along the lines of Pearle's 'CSL' theory (Pearle 1989) in order to account adequately for identical particles.

Furthermore, the *way* Bohmian mechanics, and GRW theory, are normally discussed in philosophy of physics (especially in more metaphysical contexts) is sharply at odds with the relatively humble role nonrelativistic particle mechanics plays in real quantum theory (where, recall, it is a useful model to describe various particular systems of nonrelativistic particles in the absence of radiation, normally with some phenomenologically-understood background fields and potentials). The only way I know to make sense of (most of) this literature is to interpret it as discussing nonrelativistic quantum particle mechanics *under the fiction that it is a fundamental and universal theory*. The literature is too large

¹³Originally proposed by van Fraassen (1991) and developed by many authors; see Dieks and Vermaas (1998) and references therein.

to demonstrate this exhaustively, but here are some illustrative recent examples (all drawn from Ney and Albert’s recent anthology, *The Wave Function* (2013)):

1. Albert (2013, p.53) describes the dimensionality of configuration space as “three times as large as the total number of elementary particles in the universe’, despite the fact that most such particles are not describable in the configuration-space formalism in most regimes.
2. Lewis (2013, p.111) follows Albert: “quantum mechanics represents the state of the world via a $3N$ -dimensional wave function, where N is the number of particles in the universe”.
3. Allori (2013), in her discussion of ‘primitive ontology’ — of which more later — is explicit (in her title, no less) that she is discussing “the structure of *fundamental* physical theories”, and then proceeds to apply her framework to nonrelativistic particle mechanics.
4. North (2013, pp.184-5) is refreshingly explicit: “[T]he fundamental structure of a world’s space(time) may be more properly given by a relativistic theory. Still, it is plausible that the fundamental theory of our world will be quantum mechanical. So it is worthwhile to think about what the world’s fundamental space would be if [nonrelativistic quantum particle mechanics] is its fundamental theory”.

But why work under this fiction? I think the most charitable assumption (which assumes authors realise that it *is* a fiction¹⁴) is that it is a warm-up exercise: ‘nonrelativistic quantum particle mechanics’ is a standin for the *real* fundamental quantum theory, and we hope that as many as possible of our conclusions carry over to hoped-for hidden-variables or dynamical-collapse theories formulated for the real fundamental theory.

In previous work, I have taken the referent of ‘real fundamental quantum theory’ to be something like the Standard Model, and have argued that this strategy is problematic both technically (Wallace 2008, pp.83-85, Wallace 2012, pp.33-35), because it is much harder than is generally recognised to construct a quantum-field-theory version of Bohmian mechanics or GRW theory and so confidence that such a theory even exists in premature, and conceptually (Wallace 2016a), because most of the features of nonrelativistic quantum theory appealed-to by metaphysicians of quantum mechanics are emergent approximations at best in QFT. But from this paper’s perspective, these criticisms do not get at the heart of the problem: that any strategy that works only when applied to a universal and fundamental theory does not seem to have the resources to explain the success of quantum mechanics *in general*, which is mediated through a very large number of explicitly *non*-fundamental physical theories. Absent any strategy for systematically constructing hidden-variable

¹⁴Philosophers of physics mostly do so realise, I think; certainly, all the authors I cite above do. But the very fact that the fiction is rarely spelled out explicitly creates a real risk of confusion, especially in more mainstream metaphysics discussions which get their physics second-hand from the philosophy of physics literature.

or dynamical-collapse versions of higher-level theories from lower-level theories, any hypothetical Bohmian or dynamical-collapse version of the standard model has a gap when it comes to connecting quantum theory with phenomenology and empirical confirmation.

Can the gap be filled? Doing so would seem to require us to find a way to provide information, *directly in the language of the fundamental*, about what can be observed. And the current thrust of work in the metaphysics of hidden-variable and dynamical-collapse theories aims to do exactly that. A particularly clear example is the ‘primitive ontology’ strategy of Allori *et al* (2008, 2013, this volume) (Maudlin’s (2013) ‘primary ontology’ is a close relative; Esfeld *et al* (2017) adopt essentially the same framework). The idea is that to any purportedly-fundamental theory must be associated a primitive ontology of spatially-localised matter (point particles or extended continua), such that the macroscopic world is identified with composites of the primitive ontology. (A table, for instance, is by definition a collection of primitive-ontology elements arranged in a table shape — the mereological sum of that collection, in metaphysical parlance.) We are assumed to have direct empirical access to at least coarse grainings of the primitive ontology, and so the condition for empirical success of a fundamental theory is that it generates a distribution of primitive-ontology elements whose coarse grainings match our observations of macroscopic matter. Non-primitive ontology might also be admitted to our theory (the electromagnetic field is supposed to be a classical example of non-primitive ontology; the quantum state represents non-primitive quantum ontology) but it has no direct role in empirical confirmation: we learn of it only indirectly, through its dynamical effect on the primitive ontology. Primitive ontologists seldom discuss *non-fundamental* theories, but so far as I can see, in their framework these are of purely instrumental value, providing calculational techniques to extract information about the primitive ontology but having no ontological significance of their own.

Four observations follow naturally. Firstly, although advocates of primitive ontology often write as if theirs is an unproblematic account of the physics-world relation before quantum mechanics (usually cashed out in terms of a primitive ontology of classical point particles), they do not provide historical arguments for this and it looks most implausible. The ‘particles’ to which classical particle mechanics was applied historically were typically the centers of mass of large bodies (planets, moons and comets in celestial mechanics, in particular). Other applications of classical mechanics were directly to extended objects like fluids and rigid bodies. Whatever the *hope* might have been that these large bodies could be understood as swarms of smaller bodies, that hope was never realised, and we now know that it *could not* be realised: classical microphysics does not support stable matter, and systems like comets or liquids (*pace* Allori (2013)) cannot be analysed as agglomerates of classical particles.

Secondly, if we want to look for a primitive ontology for extant physics, presumably we had better look in the Standard Model. But the task of finding a characterisation of the macroscopic *within the basic vocabulary of the Standard Model* looks monumentally difficult. Most of the connections between the

Standard Model and experiment are fantastically indirect, proceeding through layers upon layers of only-partially-understood emergence — so indirect, indeed, that philosophers like Cartwright can deny that the predictions of higher-level physics are grounded in the Standard Model at all with at least *prima facie* plausibility. Even those predictions which are ‘direct’ results of the Standard Model — the Higgs boson, say — are not stated in anything like a primitive ontology: the experimental signature of the Higgs boson is a certain resonance in the cross-section for hadron-hadron scattering, and of course ‘resonance’, ‘cross-section’, ‘scattering’ and ‘hadron’ are all highly theory-laden terms.

(Does the primitive-ontology strategy succeed in defining the observable in terms of the microscopic in non-relativistic particle mechanics, under the fiction that this theory is universal and fundamental? Never mind why the question is of any interest; it is not even well posed. For beings like us (who, *inter alia*, make most of our ‘observations’ using electromagnetic radiation) could not exist under the assumptions of that fiction. They who do not exist, do not observe.)

So the primitive ontologist is committed to forging a direct link between high-energy physics and observation, with little or no help from actual physical practice. And (my third observation) the task is made still harder by features of quantum field theories in particular. Recall that the Standard Model, like pretty much any empirically-relevant quantum field theory, is an *effective* field theory, regularised by some short-distance cutoff. The physics at the cutoff lengthscale is far from negligible, but it is possible to absorb its effects via ‘renormalisation’, whereby those degrees of freedom that describe the quantum field theory at large lengthscales are redefined in a complicated, cutoff-dependent way that absorbs most of the effects of the short-distance physics. (See Wallace (2017), and references therein, for a more detailed discussion from a conceptual point of view and for further details.) This means that the relation between the ‘fundamental’ and the empirically-relevant in quantum field theory is complicated, indirect, dynamically mediated and cutoff-dependent. It is very hard to see how this could be made compatible with the primitive-ontology approach, or indeed with any approach committed to a description of a theory’s empirical content directly in its microphysical vocabulary. It is notable that none of the various extant suggestions for Bohmian quantum field theories — based on associating the hidden variables to fermion number density (Bell 1984; Colin 2003; Colin and Struyve 2007), or particle number (Durr *et al* 2004), or field-configuration-strength (Struyve and Westman (2006, 2007), Struyve 2007) — have discussed renormalisation, or given more than a qualitative verbal plausibility argument for how these theories recover the macroworld. The acid test of such a theory is to demand a full model of how some nontrivial quantum-field-theoretic prediction — say, the cross-section for electron-electron scattering, calculated to loop order where renormalisation matters — actually plays out as a physical process, but no such theory is currently close to passing that test.)

The considerations of effective field theory are really just a reminder that even though the Standard Model is the closest we have to a fundamental physical theory, it is not such a theory, and could not be: its emergent, approximate status is built into its characterisation as an effective field theory. And this

brings me to the last observation: if the primitive-ontology strategy makes a direct connection between *fundamental* physics and observation, and higher-level, emergent physics is to be treated purely instrumentally, what — beyond idle curiosity — is the point of trying to develop such an approach until and unless we have a physical theory which we have reason to think *is* fundamental? (Not that I am sanguine about the prospects of developing a primitive ontology for that theory either: its ‘observational’ claims are likely to be even more indirect and theory-laden than the Standard Model’s.) Even in its own terms, the primitive-ontology strategy seems to be a strategy for postponing the measurement problem until the dust of fundamental physics has otherwise settled.

There is something rather ironic about this situation. Advocates of hidden-variable or dynamical-collapse theories are normally ardently committed to some form of scientific realism; to compare them to the logical positivists would be a killing insult. But what are the advocates of primitive ontology looking for, if not something like the observation language that the logical positivists sought in vain?¹⁵ And again, the problem with this strategy is not so much that the *metaphysical* distinction between primitive and non-primitive ontology is ill-defined or unmotivated; it is that we do not know how to make it, for realistic physics, in a way which achieves the task it is supposed to perform.

For expository clarity I have focussed here on primitive-ontology strategies, but I believe the observations generalise. The mainstream modificatory approaches to the measurement problem are committed to (1) developing modifications specifically to ‘fundamental’ quantum theory with only an instrumentalist attitude to non-fundamental quantum theories; (2) finding a direct way to characterise the observational evidence for that fundamental theory in its micro-physical vocabulary. These commitments disconnect modificatory approaches from physics as it is practiced, and leave it opaque at best how they can hope to account for the empirical predictions of quantum theory writ large.

7 Conclusions

I have argued that when we recognise the real structure of quantum theory — an abstract framework realised by indefinitely many concrete theories that are realised by indefinitely many concrete systems and whose relations one with another are complicated and not really hierarchical — then most extant approaches to the quantum measurement problem should be recognised as inadequate to that real structure. Of currently-extant approaches, only the Everett interpretation, in its modern decoherence-based form, provides the interpretative *recipe* to make sense of the multiplicity of quantum theories in a self-consistent way. Other approaches either try to make sense of the abstract structure of quantum

¹⁵The parallels extend to attempts to treat the non-primary ontology as lawlike (Dürr, Goldstein, and Zanghi 1997; Goldstein and Teufel 2000; Goldstein and Zanghi 2013), or to eliminate it entirely at the metaphysical level via a Humean account of laws (Esfeld *et al* 2014, Miller 2014). See Dewar (2017) for further discussion.

theory and thus fail to give an adequate account of scientific representation in concrete applications, or else analyse one particular quantum theory as if it alone could exhaust the theoretical content of the subject, and as if all experimental predictions could be described directly in its vocabulary. If this is correct, there is no underdetermination in quantum mechanics: either the Everett interpretation is viable, in which case it alone provides an adequate interpretation of quantum theory; or else it must be rejected for some philosophical or technical reason, in which case there is at present *no* adequate interpretation of quantum theory.

To the reader who resists this conclusion, there is a straightforward way to prove me wrong. Take any moderately-complicated, moderately-concrete application of quantum theory in a regime which is not fully covered by nonrelativistic particle mechanics: take the BCS model of superconductivity, for instance (with electromagnetic radiation present), or take the quark-gluon plasma, or take the Higgs mechanism, or take the colour of gold; come to that, take the photoelectric effect, or the (photon) two-slit experiment. Explain, in reasonable detail, how that application is to be described, understood, and tested inside your preferred interpretation of quantum mechanics. And check, in particular, that your preferred interpretation can be used to make the quantitative calculations whose match with experiment is the reason why those applications are deemed successful. The Everett interpretation can do this, at least if its account of probability and the structuralist notions of emergence that it relies on are deemed adequate, simply by working through the standard accounts of these applications in physics and interpreting each particular version of quantum mechanics appealed to in Everettian terms. If your preferred alternative to Everettian quantum mechanics cannot do it, you have not yet solved the quantum measurement problem.

This might seem like cheating. What I have called an ‘Everettian’ way of making sense of concrete quantum applications is, after all, really just the explanation of those applications found within mainstream physics. But this simply underlines the fact that Everettian quantum mechanics — ontological extravagance at the macro-level notwithstanding — is a modest, conservative project, aimed at *legitimizing* and *making sense of* the ordinary practice of quantum theory. Other approaches to the measurement problem, by and large, are less conservative: the ordinary practice of quantum theory is fundamentally confused and in large part needs to be reformulated or replaced, not ‘made sense of’. But the great virtue of modesty and conservatism, in this context, is that it minimises the need to redo from scratch much of the last century of physics.

Let me finish, as did Maudlin (*ibid*), by paraphrasing Hume: We can be clear on the questions that must be asked of an interpretation. Does it provide a way to legitimate and make sense of the actual practice of quantum physics, across the various interrelated domains to which quantum theory has been applied? Does it set out to reform the practice of quantum physics, and does it provide evidence that this is more than a bluff by actually doing the hard work in some non-trivial, concrete examples across multiple instantiations of the quantum framework (and not just in an abstract quantum-information setting, or

in nonrelativistic particle mechanics)? If the answer in each case is “no” . . . , well, maybe don’t hastily commit it to the flames, as it may contain valuable insights and be the seed of a yet-to-be-completed research program, but don’t kid yourself that it is *at present* a viable interpretation of quantum mechanics, and maybe be a bit cautious exploring all its metaphysical implications until you’ve done some more work to see if it plausibly might be made viable.

Acknowledgements

I am grateful to Neil Dewar, Richard Healey, James Ladyman, Juha Saatsi, Simon Saunders and Chris Timpson for useful conversations, and to two anonymous referees for helpful comments and corrections.

References

- Albert, D. Z. (2013). Wave function realism. In A. Ney and D. Z. Albert (Eds.), *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, pp. 52–57. Oxford: Oxford University Press.
- 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.
- Arntzenius, F. (1990). Kochen’s interpretation of quantum mechanics. *Proceedings of the Philosophy of Science Association* 1, 241–249.
- Arntzenius, F. (1998). Curiouser and curiouser: A personal evaluation of modal interpretations. See Dieks and Vermaas (1998), pp. 337–377.
- Bacciagaluppi, G. (2000). Delocalized properties in the modal interpretation of quantum mechanics. *Foundations of Physics* 30, 1431–1444.
- Bacciagaluppi, G., M. J. Donald, and P. E. Vermaas (1995). Continuity and discontinuity of definite properties in the modal interpretation. *Helvetica Physica Acta* 68, 679–704.
- Barrett, J. A. (2011a). Everett’s pure wave mechanics and the notion of worlds. *European Journal for Philosophy of Science* 1, 277–302.
- Barrett, J. A. (2011b). On the faithful interpretation of pure wave mechanics. *British Journal for the Philosophy of Science* 62, 693–709.
- Bell, J. S. (1984). Beables for quantum field theory. CERN preprint CERN-TH 4035/84. Reprinted in Bell (1987), pp. 173–180.
- Bell, J. S. (1987). *Speakable and Unsayable in Quantum Mechanics*. Cambridge: Cambridge University Press.

- Bever, B. M. (2011). Everett's "many-worlds" proposal. *Studies in the History and Philosophy of Modern Physics* 42, 3–12.
- Blackmore, S. (2005). *Conversations on Consciousness*. Oxford: Oxford University Press.
- Bohm, D. (1952). A Suggested Interpretation of Quantum Theory in Terms of "Hidden" Variables. *Physical Review* 85, 166–193.
- Bohm, D. and B. J. Hiley (1993). *The Undivided Universe: An Ontological Interpretation of Quantum Theory*. London: Routledge and Kegan Paul.
- Carnap, R. (1928). *Der Logische Aufbau der Welt*. Leipzig: Felix Meiner Verlag.
- Cartwright, N. (1983). *How the Laws of Physics Lie*. Oxford: Oxford University Press.
- Cartwright, N. (1999). *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.
- Caves, C. M., C. A. Fuchs, and R. Schack (2002). Quantum probabilities as Bayesian probabilities. *Physical Review A* 65, 022305.
- Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of consciousness studies* 2, 200–219.
- Clifton, R. (1996). The properties of modal interpretations of quantum mechanics. *British Journal for the Philosophy of Science* 47, 371–398.
- Colin, S. (2003). Beables for quantum electrodynamics. Available online at <http://arxiv.org/abs/quant-ph/0310056>.
- Colin, S. and W. Struyve (2007). A Dirac sea pilot-wave model for quantum field theory. *Journal of Physics A* 40, 7309–7342.
- Cordero, A. (2001). Realism and underdetermination: Some clues from the practices-up. *Philosophy of Science* 68, S301–S312.
- Creath, R. (2017). Logical empiricism. In *The Stanford Encyclopedia of Philosophy* (Fall 2017 Edition), Edward N. Zalta (ed.), online at <https://plato.stanford.edu/archives/fall2017/entries/logical-empiricism/>.
- Dennett, D. C. (1984). *Elbow Room: the Varieties of Free Will Worth Wanting*. Oxford: Oxford University Press.
- Dennett, D. C. (2005). *Sweet Dreams: Philosophical Objections to a Science of Consciousness*. Cambridge, MA: MIT Press.
- Dewar, N. (2017). La bohème. <http://philsci-archive.pitt.edu/13235/>.
- Dieks, D. and P. E. Vermaas (Eds.) (1998). *The Modal Interpretation of Quantum Mechanics*, Dordrecht. Kluwer Academic Publishers.
- Donald, M. J. (1998). Discontinuity and continuity of definite properties in the modal interpretation. See Dieks and Vermaas (1998), pp. 213–222.
- Dürr, D., S. Goldstein, R. Tumulka, and N. Zanghi (2004). Bohmian mechanics and quantum field theory. *Physical Review Letters* 93, 090402.

- Dürr, D., S. Goldstein, and N. Zanghi (1997). Bohmian mechanics and the meaning of the wave function. In R. S. Cohen, M. Horne, and J. Stachel (Eds.), *Potentiality, Entanglement and Passion-at-a-Distance — Quantum Mechanical Studies in Honor of Abner Shimony*. Dordrecht: Kluwer. Available online at <http://arxiv.org/abs/quant-ph/9512031>.
- Egg, M. (2014). *Scientific realism in particle physics: a causal approach*. de Gruyter.
- Esfeld, M., D. Lazarovici, M. Hubert, and D. Dürr (2014). The ontology of Bohmian mechanics. *British Journal for the Philosophy of Science* 65, 733–796.
- Esfeld, M., D. Lazarovici, V. Lam, and M. Hubert (2017). The physics and metaphysics of primitive stuff. *British Journal for the Philosophy of Science* 68, 133–161.
- Everett, H. (2012). *The Everett Interpretation of Quantum Mechanics: Collected Works 1955-1980 with Commentary*. Princeton: Princeton University Press. Edited by Jeffrey A. Barrett and Peter Byrne.
- Fuchs, C. (2002). Quantum mechanics as quantum information (and only a little more). Available online at <http://arXiv.org/abs/quant-ph/0205039>.
- Fuchs, C. and A. Peres (2000a). Quantum theory needs no “interpretation”. *Physics Today* 53(3), 70–71.
- Fuchs, C. A. (2017). On participatory realism. In I. T. Durham and D. Rickles (Eds.), *Information and Interaction: Eddington, Wheeler, and the Limits of Knowledge*, pp. 113–134. Springer. Online version at <https://arxiv.org/abs/1601.04360>.
- Fuchs, C. A., N. D. Mermin, and R. Schack (2014). An introduction to QBism with an application to the locality of quantum mechanics. *American Journal of Physics* 82, 749–754.
- Fuchs, C. A. and A. Peres (2000b). Fuchs and Peres reply. *Physics Today* 53, 14.
- Gell-Mann, M. and J. B. Hartle (1989). Quantum Mechanics in the Light of Quantum Cosmology. In W. H. Zurek (Ed.), *Complexity, Entropy and the Physics of Information*, pp. 425–459. Redwood City, California: Westview Press.
- Gell-Mann, M. and J. B. Hartle (1993). Classical equations for quantum systems. *Physical Review D* 47, 3345–3382.
- Ghirardi, G., A. Rimini, and T. Weber (1986). Unified Dynamics for Micro and Macro Systems. *Physical Review D* 34, 470–491.
- Goldstein, S. and S. Teufel (2000). Quantum spacetime without observer: ontological clarity and the conceptual foundations of quantum gravity. In C. Callender and N. Huggett (Eds.), *Physics meets Philosophy at the Planck Scale*, pp. 275–289. Cambridge: Cambridge University Press.

- Goldstein, S. and N. Zanghi (2013). Reality and the role of the wave function in quantum theory. In A. Ney and D. Z. Albert (Eds.), *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, pp. 91–109. Oxford University Press.
- Healey, R. (2012). Quantum theory: A pragmatist approach. *British Journal for the Philosophy of Science* 63, 729–771.
- Healey, R. (2017a). how pragmatism reconciles quantum mechanics with relativity etc (interview by richard marshall). <http://www.3ammagazine.com/3am/how-pragmatism-reconciles-quantum-mechanics-with-relativity-etc/>.
- Healey, R. (2017b). *The Quantum Revolution in Philosophy*. Oxford: Oxford University Press.
- Kochen, S. and E. Specker (1967). The problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics* 17, 59–87.
- Kuhn, T. (1962). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Landau, L. and E. Lifshitz (1977). *Quantum Mechanics (Non-Relativistic Theory)* (3rd ed.). Butterworth Heinemann.
- Lewis, P. (2013). Dimension and illusion. In A. Ney and D. Albert (Eds.), *The Wave Function*, pp. 110–125. Oxford: Oxford University Press.
- Lewis, P. J. (2015). Internet encyclopedia of philosophy. <http://www.iep.utm.edu/int-qm/>.
- Lyre, H. (2010). Why quantum theory is possibly wrong. *Foundations of Physics* 40, 1429–1438.
- Maudlin, T. (1995). Three measurement problems. *Topoi* 14, 7–15.
- 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.
- Miller, E. (2014). Quantum entanglement, Bohmian mechanics, and Humean supervenience. *Australasian Journal of Philosophy* 92, 567–583.
- Ney, A. and D. Albert (Eds.) (2013). *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*. 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.
- Oppenheim, P. and H. Putnam (1958). The unity of science as a working hypothesis. *Minnesota Studies in the Philosophy of Science* 2.
- Pearle, P. (1989). Combining Stochastic Dynamical State-Vector Reduction with Spontaneous Localization. *Physical Review A* 39(5), 2277–2289.
- Peres, A. (1993). *Quantum Theory: Concepts and Methods*. Dordrecht: Kluwer Academic Publishers.

- Peskin, M. E. and D. V. Schroeder (1995). *An introduction to Quantum Field Theory*. Reading, Massachusetts: Addison-Wesley.
- Putnam, H. (1962). What theories are not. In E. Nagel, P. Suppes, and A. Tarski (Eds.), *Logic, Methodology, and Philosophy of Science*, pp. 240–251. Stanford: Stanford University Press.
- Quine, W. V. O. (1951). Two dogmas of empiricism. *Philosophical Review* 60, 20–43.
- Saunders, S. (1993). Decoherence, Relative States, and Evolutionary Adaptation. *Foundations of Physics* 23, 1553–1585.
- Saunders, S. (1995). Time, Decoherence and Quantum Mechanics. *Synthese* 102, 235–266.
- Saunders, S. (1997). Naturalizing Metaphysics. *The Monist* 80(1), 44–69.
- Struyve, W. (2007). Field beables for quantum field theory. Available online at <http://arxiv.org/abs/0707.3685>.
- Struyve, W. and H. Westman (2006). A new pilot-wave model for quantum field theory. *AIP CONFERENCE PROCEEDINGS* 844, 321.
- Struyve, W. and H. Westman (2007). A minimalist pilot-wave model for quantum electrodynamics. *Proceedings of the Royal Society of London A* 463, 3115–3129.
- van Fraassen, B. C. (1991). *Quantum Mechanics*. Oxford: Oxford University Press.
- von Neumann, J. (1955). *Mathematical Foundations of Quantum Mechanics*. Princeton: Princeton University Press.
- Wallace, D. (2003). Everett and Structure. *Studies in the History and Philosophy of Modern Physics* 34, 87–105.
- Wallace, D. (2008). The interpretation of quantum mechanics. In D. Rickles (Ed.), *The Ashgate Companion to Contemporary Philosophy of Physics*, pp. 197–261. Burlington, VT: Ashgate. Available online at <http://arxiv.org/> under the title “The Quantum Measurement Problem: State of Play”.
- Wallace, D. (2010). Decoherence and ontology: Or: How I learned to stop worrying and love FAPP. In S. Saunders, J. Barrett, A. Kent, and D. Wallace (Eds.), *Many Worlds? Everett, Quantum Theory, and Reality*, pp. 53–72. Oxford: Oxford University Press.
- Wallace, D. (2011). Taking particle physics seriously: a critique of the algebraic approach to quantum field theory. *Studies in the History and Philosophy of Modern Physics* 42, 116–125.
- Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford: Oxford University Press.
- Wallace, D. (2016a). Lessons from realistic physics for the metaphysics of quantum theory. Forthcoming.

- Wallace, D. (2016b). Probability and irreversibility in modern statistical mechanics: Classical and quantum. To appear in D. Bedingham, O. Maroney and C. Timpson (eds.), *Quantum Foundations of Statistical Mechanics* (Oxford University Press, forthcoming).
- Wallace, D. (2017). Quantum field theory. Forthcoming in Knox and Wilson (ed.), *Handbook of the Philosophy of Physics* (Routledge).
- Wilson, M. (2013). What is “classical mechanics” anyway? In R. Batterman (Ed.), *The Oxford Handbook of Philosophy of Physics*, pp. 43–106. Oxford: Oxford University Press.
- Zeh, H. (1973). Toward a quantum theory of observation. *Foundations of Physics* 3, 109–116.
- Zeh, H. D. (1993). There are no quantum jumps, nor are there particles! *Physics Letters A* 172, 189.
- Zurek, W. and J. P. Paz (1995). Quantum chaos: a decoherent definition. *Physica D* 83, 300–308.
- Zurek, W. H. (1991). Decoherence and the transition from quantum to classical. *Physics Today* 43, 36–44. Revised version available online at <http://arxiv.org/abs/quant-ph/0306072>.
- Zurek, W. H. (1998). Decoherence, einselection, and the quantum origins of the classical: the rough guide. *Philosophical Transactions of the Royal Society of London A* 356, 1793–1820.