Interpreting the quantum mechanics of cosmology

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 'Philosophy of Cosmology: an Introduction' edited by A. Ijjas and B. Loewer, due for publication in 2019.

Abstract

Quantum theory plays an increasingly significant role in contemporary early-universe cosmology, most notably in the inflationary origins of the fluctuation spectrum of the microwave background radiation. I consider the two main strategies for interpreting (as opposed to modifying or supplementing) standard quantum mechanics in the light of cosmology. I argue that the conceptual difficulties of the approaches based around an irreducible role for measurement - already very severe - become intolerable in a cosmological context, whereas the approach based around Everett's original idea of treating quantum systems as closed systems handles cosmological quantum theory satisfactorily. Contemporary cosmology, which indeed applies standard quantum theory without supplementation or modification, is thus committed - tacitly or explictly - to the Everett interpretation.

> Cosmology is the killer app for the Everett interpretation.
>
> Jim Hartle.[1]

1 States, dynamics, and the Objective view of physics

Pretty much any classical-mechanical theory, be it Newtonian particle mechanics, or the behaviour of coiled springs, or classical field theory, or general relativity, has a standard mathematical form — or at any rate can be put into that standard form. It is specified by:

1. A *state space* $\mathcal{S}$: a set of points with at least enough structure to define concepts of differentiation.

---

[1] Presentation to "Everett@50" conference: Oxford 19th-21st July 2007.

2. A *dynamics* on $\mathcal{S}$: a first-order differential equation that determines, through each point in $\mathcal{S}$, a path satisfying that equation. (Furthermore, the path is either unique, or — as in the case of gauge theories — its non-uniqueness is taken to indicate a many-to-one relation between the mathematics and the physics.)

In Newtonian particle mechanics, for instance, the state space is the space of N-tuples of points in 3-dimensional Euclidean space, and the dynamics is Newton's second law together with some expression for the force. In scalar field theory, the state space is the space of functions from Euclidean space to the real (or complex) numbers, and the dynamics is some field equation like the Klein-Gordon equation.

It is possible to have debates that might be called "interpretational" about these classical theories. Philosophers can debate — indeed, have debated — whether genuine point particles are objects in their own right or simply derivative on some primitive property of "occupation" that spacetime points do or do not possess; they have asked whether a "field" is an object whose parts occupy spacetime points, or simply a *facon de parler* for ascribing certain properties *to* spacetime points, or even whether the fact that ordinary language allows us to ask such questions shows that we need to approach metaphysics in a way more suited to physics. In the case of gauge theories their concerns have overlapped with physicists' in exploring the nature of the gauge principle and the status of spacetime concepts in general relativity.

But none of these questions amount to genuine dispute about 'interpretation' of the kind found in quantum theory. For the basic relation between theory and world is not in contention in any of these classical theories. In each case, a simple recipe applies for any physical system described by some classical theory:

1. At each instant of time, the physical features of the system are represented by the state. That is: the state *objectively represents* the way the physical system *actually is.*

2. The equation of motion of the system determines how those physical features change. That is, the trajectory in state space determined by the dynamics is a representation of the *actual history* of the system.

Call this the Objective view of physics. (I have elsewhere (Wallace 2013) called it, or something very similar, the *dynamical conception* of physics.) It is a view in which measurement and observation are sidelined: for sure, a sufficiently canny physicist may measure the system and so learn part or all of the truth about the system's history, but:

- That measurement process is external to the formalism of the theory; that formalism describes only the uninterrupted evolution of the system.

- The goal of a measurement process is to determine the objective, already-true facts about the system's state at any given time, and no barrier of principle prevents that goal being achieved arbitrarily well.

- While realistically a measurement process may disturb a system's evolution, that disturbance is not an essential part of measurement, and no barrier of principle prevents a measurement leaving a system undisturbed to arbitrarily high levels of accuracy.

The objective view is, in itself, silent about the actual process of measurement: measurement, in idealisation, is just some process whereby a physicist comes to know the facts about the system, while the system itself continues blithely on. Should we wish to go beyond this level of idealisation and actually ask how the measurement process works, the Objective view gives a simple recipe: measurement devices are just physical systems themselves, with their own state spaces and dynamics, and a "measurement" is just a dynamical coupling of the measurement device with the system being measured, with the result that the measurement device's state post-measurement ends up correlated with the measured system's state pre-measurement.

Of course, in most cases where we apply classical physics, representing the actual process of measurement inside the physics is impossible at anything beyond the most phenomenological level (not least because real measurement devices are made of atoms, which are unstable classically!) But the principles of application are nonetheless clear.

The same holds when we go beyond standard classical mechanics and consider stochastic theories, such as Brownian motion or the Langevin equation. The deterministic dynamics is replaced with a stochastic differential equation determining a probability distribution over a state's future histories, but the interpretation is unchanged: those future histories are the possible histories the system will actually have, and measurement will simply tell us which one actually obtains.

2 Quantum theory and the Objective View

A quantum theory — be it non-relativistic particle physics, or a collection of interacting oscillators, or a quantum field theory — *formally* has the same structure as a classical theory. Its state space is the space of rays in some Hilbert space, or if preferred, the space of density operators on that same Hilbert space Its dynamics is given by a semigroup of unitary operators determined by the Schrödinger equation.

To take the Objective view of quantum theory, though, would require us to interpret a point in the state space as representing a possible physical state of the world — and in quantum theory that is non-trivial. The problem is not that Hilbert space is a featureless mathematical space, in which no two rays can be distinguished. After all, the state space in classical mechanics — phase space — is, *qua* phase space, almost as featureless. A phase space is just defined as a symplectic manifold, and any two points in such a space related by a canonical transformation are indistinguishable with respect to the structure of such a space. But two so-related points in phase space can still represent intrinsically different physical possibilities, because any *particular* classical-mechanical theory has a great deal of structure on its phase space beyond that defined by the basic formalism of classical mechanics. The phase space of N-particle mechanics, for instance, is equipped with a preferred decomposition into one-particle phase spaces, each of which in turn comes with a preferred action of the symmetry group of spacetime on it; the phase space of a field theory comes with a preferred identification between points in physical space and functions on the phase space.[2]

Likewise, the Hilbert space of any *particular* quantum system is highly structured, with the structure typically given by a preferred relation between self-adjoint operators and either some spacetime symmetry group or else the spacetime itself. In a quantum field theory, in particular, the field operator is (speaking somewhat formally) a map from points $\mathbf{x}$ in space to operators $\widehat{\Psi}(\mathbf{x})$ on the Hilbert space.

The problem, rather, is that quantum theory also comes equipped with an interpretative rule — the *Born Rule* (or *probability postulate*) — that seems at odds with any use of this preferred structure to determine objectively-possessed properties of a quantum system. The Born Rule, in its most minimal form (to be expanded upon later) tells us that the result of measuring the possessed value of some quantity represented by operator $\widehat{X}$ is determined only probabilistically by the state $|\psi\rangle$, and that the expectation value of that measurement is $\langle\psi|\,\widehat{X}\,|\psi\rangle$.

This, of course, leads swiftly to trouble for the Objective View. From the Born rule it follows that measurements of (the quantity represented by) $\widehat{X}$ give a deterministic result iff the state $|\psi\rangle$ is an eigenstate of $\widehat{X}$. On the Objective View, measurement is simply a process of recording already-possessed values; hence, an Objective view of quantum mechanics tells us that if $\widehat{X}\,|x\rangle = x\,|x\rangle$, $|x\rangle$ represents a system with possessed value x of the quantity. But then it is obscure what to make of a superposition like

$$|\psi\rangle = \sum_i \lambda_i \,|x_i\rangle\,. \tag{1}$$

Measurement will give some x_i as a result, yet $|\psi\rangle$ itself tells us nothing of which x_i will result. And this looks to be in flat contradiction with the Objective View: if measurement simply returns an actually-possessed value of that which is measured, and if the state of the system is supposed to represent the physical features of the system being measured, then that state must represent the actually-possessed value — yet it manifestly does not, when one measurement of a system in state $|\psi\rangle$ can give one result and a second measurement on a system in the very same state can give quite another.

---

[2] Sometimes 'phase space' is defined more richly, as the cotangent bundle to a differentiable manifold (the 'configuration space'). But the same basic point goes through: the structure *qua* phase space is preserved by any cotangent lift of a diffeomorphism on the configuration space, so any points related by such a cotangent lift are indistinguishable with respect to the phase-space structure; in reality, any concrete example of a classical-mechanical theory places a great deal of structure on the configuration space, and it is largely that structure which allows different points in phase space to represent distinguishable physical possibilities.

What to do? At least in foundational circles, there is a strong temptation to modify or supplement quantum mechanics in some manner so as to preserve as much as possible of the Objective View. But while much can be said in favour of this temptation, it must be acknowledged that physics seems to have got by for 85 years without any apparent need of such modification. Furthermore, the most fully developed such modifications — the hidden-variable theory of de Broglie and Bohm (cf Bohm 1952; Bohm and Hiley 1993; Holland 1993; Cushing, Fine, and Goldstein 1996) and the dynamical-collapse theories of Ghirardi, Rimini, and Weber (1986) and Pearle (1989) — are effectively toy models, restricted to the domain of non-relativistic particle mechanics in the absence of electromagnetic radiation. While there have been serious attempts to extend these theories to the relativistic domain (Colin and Struyve 2007, Struyve and Westman 2006, 2007, Tumulka 2006, Bedingham 2010, 2011), it's fair to say that there is so far no such generally-accepted extension to an empirically adequate effective field theory like electroweak theory or the full Standard Model, let alone to the post-Standard-Model physics typically appealed to in contemporary inflationary cosmology.[3]

As such, the urgent priority for advocates of the "modify or supplement" strategy is (or ought to be) to develop robust extensions of their theories to the full range of contemporary physics. There is a place for philosophical explorations of these theories, but that exploration needs to be tentative and provisional, and not overcommitted to precise details of the existing models; in the meantime, the main problems of these approaches are technical rather than philosophical (they were, after all, designed precisely so as to resolve the main philosophical problems with quantum theory).

By contrast, if we want to understand the apparently-successful application of *extant* quantum theory in contemporary cosmology (and indeed in contemporary physics more generally), our problem is essentially philosophical. We seem in practice to be able to apply the theory despite the apparent failure of the Objective View: we need to understand if this practice can be underpinned by a conceptually sound understanding of quantum theory.

3 The Lab View

The Born Rule is phrased in terms of measurement, and from the earliest days of quantum physics it has been interpreted in a way which gives measurement a central conceptual role, in sharp contrast with the passive, external-monitoring view of measurement seen in the Objective View. We can make this role explicit in an alternative approach to a physical theory, which might be called the Lab View: it can fairly be said to underpin many textbook discussions of quantum

---

[3] There have been highly interesting attempts to develop *testable alternatives* to that contemporary cosmology using principles from hidden-variable physics (where those theories are understood phenomenologically and schematically); see, e. g. , Valentini (2010), Underwood and Valentini (2015), Goldstein, Struyve, and Tumulka (2016). This work lies outside the scope of this paper; its ultimate test is experimental.

theory (here I follow in particular Peres (1993)).

In the Lab View, any application of a physical theory should be understood as applying to some experimental setup. That setup in turn is broken into three processes:

1. State preparation;
2. Dynamics;
3. State measurement.

The first and last of these are *primitive*: the question of how the system is prepared in a given state, and how it is measured, are external to the experiment and so not modelled in the physics. Only the second is regarded as a modelled physical process.

In quantum mechanics, in particular:

1. The system is prepared in a state represented by some (pure or mixed) Hilbert-space state;
2. It evolves under the Schrödinger equation for some fixed period of time;
3. The quantity being measured is represented by a self-adjoint operator, and the measurement outcome is given probabilistically by the Born rule, with respect to the time-evolved state and self-adjoint operator.

The Lab View itself does not force a unique interpretation of the underlying physics, but those interpretations compatible with it inevitably give a special status to measurement. Particularly prominent are:

Straight operationalism: There is nothing more to quantum mechanics than a calculus that connects preparation processes (conceived of macroscopically and phenomenologically) to measurement processes (likewise conceived of); physics neither needs, nor can accommodate, any microscopic story linking the two.

Straight operationalism is perhaps the closest realisation in mainstream physics of the old logical-positivist conception of the philosophy of science; it seems to have been more or less Heisenberg's preferred approach, and has been advocated more recently by Peres (Peres 1993, pp.373–429, Fuchs and Peres 2000). The 'quantum Bayesianism' or "QBism" of Fuchs *et al* (Fuchs 2002; Fuchs, Mermin, and Schack 2014; Fuchs and Schack 2015) has much in common with straight operationalism, although it holds out for some objective physical description at a deeper level (see Timpson (2010, pp.188–235) for a critique).

Complementarity: It *is* possible to describe a physical system at the microscopic level, but the appropriate description depends on the experimental context in question. Is an electron a wave or a particle? If you're carrying out a two-slit experiment, it's a wave; if you change experimental context to check which slit it went down, it's a particle.

Niels Bohr is the most famous proponent of complementarity, though he tended to describe it in qualitative terms and engaged little with modern (Schrödinger-Heisenberg-Dirac) quantum mechanics. Saunders (2005) provides a rational reconstruction of complementarity in modern terminology; the approaches of Omnes (1988, 1992, 1994) and Griffiths (1984, 1993, 1996) are very much in the spirit of complementarity.

Measurement-induced collapse: The system can be described in microscopic terms, and in a way independent of the measurement process: the physical quantities of the system are represented by the state, and the system has a definite value of the quantity represented by operator $\widehat{X}$ iff its state is an eigenvector of $\widehat{X}$ (the 'Eigenvalue-Eigenvector link'), or perhaps iff it is very close to an eigenvector (the 'Fuzzy link'; cf Albert and Loewer (1996)). But at the final moment of measurement at the end of the experiment, the state discontinuously jumps ('collapses') into an eigenstate of the operator being measured, with the probability of a jump being given by the Born rule.

Dirac (1930) and von Neumann (1955) both interpreted quantum mechanics via measurement-induced collapse, and it is the approach most often seen in introductory textbooks today. Physicists often call it the Copenhagen interpretation, though historically that term is better used to describe the philosophies of Bohr and Heisenberg (philosophers often call it the Dirac-von Neumann interpretation). Strictly speaking it is not a pure *intepretation* of quantum theory, since the collapse of the wavefunction is a one-off modification of the dynamics, but since it is introduced for purely philosophical reasons and is undefined beyond the statement that it happens 'at measurement', it has more in common with other pure interpretations than with genuine modifications of quantum theory.

The Lab View is often criticised on philosophical grounds, as part of the general rejection of operationalism in contemporary philosophy of physics (see Newton-Smith (1981, ch.II), Psillos (1999, chs.1-2), and references therein for philosophical discussion, and Deutsch (1997) and Bell (2004) for physicists' versions of the same critique). But these criticisms — though quite correct! — are not my concern here. In fact there are straightforwardly *scientific* grounds to reject the Lab View (or at least to recognise it as a special case of something more general). Firstly, it fails to do full justice to the practice of experimental physics; secondly, and of particular relevance to cosmology, it fails to allow for applications of quantum theory in situations which don't fit the 'experiment' framework at all.

4 Limitations of the Lab View: real experiments

A measurement apparatus or a state-preparation device (my experimentalist colleagues assure me) is not an unanalysed gift from God: it is made of atoms, and designed and built on the assumption that its behaviour is governed by physical laws.

Which laws? Back in the glory days of the Copenhagen intepretation, perhaps it was possible to suppose that the workings of lab equipment should be analysed classically, but in these days of quantum optics, superconducting supercolliders and gravity-wave-sensitive laser interferometers, we cannot avoid making extensive reference to quantum theory itself to model the workings of our apparatus. And now a regress beckons: if we can understand quantum theory only with respect to some experimental context, what is the context in which we understand the application of quantum theory to the measurement itself?

Relatedly: how should we apply quantum mechanics in the case of *repeated* measurements, where we measure a given quantity not once but twice or many times? It is sometimes suggested that the Dirac-von Neumann collapse postulate solves this problem, and indeed (not least by Dirac himself) that the collapse postulate is justified precisely to ensure that measurement is repeatable.

But this will not do. Typical measurement processes are *not* repeatable: standard methods of detecting a photon, for instance, absorb and destroy that photon (at least, such is the language we usually use, however dubious it might be on the Lab View). To suppose that the collapse postulate is there only to ensure that repeatable measurements are repeatable is to court tautology, and at any rate does not answer the general problem of how to analyse experiments involving multiple measurement processes.

The limit of repeated measurements is continuous observation. This is routine in physics: the Geiger counter clicks when the nucleus has decayed, but that counter does not operate in a discrete fashion. Trying to discretise this process, apply a collapse postulate, and take the limit leads to disaster: as Misra and Sudarshan (1977) observed, in that limit the quantum Zeno effect entirely halts evolution, and it was for exactly that reason that they initially called it a 'paradox'. (It is telling that while Schrödinger's original 'cat' thought experiment involved a decaying nucleus and so a continuously-changing amplitude for the cat to die, modern analyses usually discretise the process.)

For all this, physics clearly *does* succeed in applying quantum theory to these experimental contexts, and so must in some way go beyond the Lab View to do so. The method is in each case the same:

1. Expand the analysis to include the apparatus itself as part of the quantum system. (In quantum information this move has come to be known as 'the Church of the Larger Hilbert Space'.)

2. Avoid the regress by treating the Born-Rule-derived probability distribution over *macroscopic* degrees of freedom not as a probability of *getting*

> *certain values on measurement*, but as a probability of *certain values already being possessed.*

Schematically, a measurement process of, say, a spin-half particle looks like

$$(\alpha\,|{\uparrow}\rangle + \beta\,|{\downarrow}\rangle) \otimes |\text{`ready'}\rangle \rightarrow \alpha\,|\varphi_{\uparrow}\rangle \otimes |\text{`up'}\rangle + \beta\,|\varphi_{\downarrow}\rangle \otimes |\text{`down'}\rangle \tag{2}$$

where $|\text{`ready'}\rangle$, $|\text{`up'}\rangle$ and $|\text{`down'}\rangle$ represent states of a macroscopic measurement device whose macroscopic variables have approximately definite results picking out the appropriate measurement outcomes. (In a realistic analysis of the device's function, of course, there would actually be a high-dimensional subspace corresponding to each of these results, and not just a single state.) Applying the Born rule directly to these variables tells us that before measurement, the device had a 100% chance of being in the 'ready' state; after measurement, a $|\alpha|^2$ chance of being in the 'up' state and a $|\beta|^2$ chance of being in the 'down' state.

If we wish to measure the state of the particle twice, we need a second copy of the measurement device — and its outcome will depend on the post-measurement states $|\varphi_{\uparrow}\rangle$ and $|\varphi_{\downarrow}\rangle$. In the case where $|\varphi_{\uparrow}\rangle = |{\uparrow}\rangle$ and $|\varphi_{\downarrow}\rangle = |{\downarrow}\rangle$ — which *by definition* we call 'non-disturbing' — the result (assuming both devices work the same way) is

$$\begin{gathered}(\alpha\,|{\uparrow}\rangle + \beta\,|{\downarrow}\rangle) \otimes |\text{`ready'}\rangle \otimes |\text{`ready'}\rangle \rightarrow \\ \alpha\,|\varphi_{\uparrow}\rangle \otimes |\text{`up'}\rangle \otimes |\text{`up'}\rangle + \beta\,|\varphi_{\downarrow}\rangle \otimes |\text{`down'}\rangle \otimes |\text{`down'}\rangle \end{gathered} \tag{3}$$

and the Born rule tells us that the two devices have a 100% chance to agree on the measurement outcome. At the other extreme, if the measurement process leaves the state of the system in the spin-up state irrespective of its original state — that is, $|\varphi_{\uparrow}\rangle = |\varphi_{\downarrow}\rangle = |{\uparrow}\rangle$ — then the measurement evolution will give

$$\begin{gathered}(\alpha\,|{\uparrow}\rangle + \beta\,|{\downarrow}\rangle) \otimes |\text{`ready'}\rangle \otimes |\text{`ready'}\rangle \rightarrow \\ |{\uparrow}\rangle \otimes (\alpha\,|\text{`up'}\rangle + \beta\,|\text{`down'}\rangle) \otimes |\text{`up'}\rangle \end{gathered} \tag{4}$$

and the Born Rule tells us that the second measurement is certain to give result 'up' irrespective of the result of the first.

Similarly, continuous measurement processes (such as the Geiger counter) can be modelled as interactions whereby the measurement device evolves (at some fixed response rate) from a 'ready' to a 'triggered' state when the system's state is in some region of Hilbert space and remains static otherwise. If the system evolves from the static to the trigger region, then depending on how the speed of the system's evolution compares with the detector response rate, the detector may passively record the transition or may halt it almost entirely through Zeno-type effects with a smooth transition between the two regimes as the respective timescales are varied (cf Home and Whitaker (1997) for further discussion). In the particular case of the Geiger counter, its response rate is so slow compared to the transition rate of nuclear decay that Zeno freezing is negligible.

To sum up: both the general technological problem of constructing measurement devices, and the more specific problems that occur when we consider experimental contexts more general than the stylised ones used in the Lab View, force us to expand the system studied by quantum theory to include the measurement device itself; having done so, we no longer consider the supposed 'measurement context' of that larger system, but just apply the Born Rule directly to the *macroscopic* degrees of freedom.

This need for an objective, non-quantum, macroscopically applicable language to describe the physics of measurement was already recognised by Bohr (and is acknowledged in more sophisticated operationalist accounts of quantum theory; cf. Peres (1993, 423–427)). As we shall see, it is really a special case of a more general requirement.

5 Limitations of the Lab View: beyond the 'experiment' paradigm

It has long been suggested that Lab View quantum mechanics is unsuitable for cosmology simply because cosmology concerns the whole Universe, and so there is no 'outside measurement context'; indeed, it was for exactly these reasons that Hugh Everett developed his approach to quantum theory in the first place (Everett 1957).

However, this slightly misidentifies the problem. Cosmology is concerned with the Universe on its largest scales, but not with every last feature of the Universe: realistic theories in cosmology concern particular degrees of freedom of the universe (the distribution of galaxies, for instance) and we can perfectly well treat these degrees of freedom as being measured via their interaction with other degrees of freedom outside the scope of those theories (Fuchs and Peres 2000).

But there is a problem nonetheless. Namely: the processes studied in cosmology cannot, even in the loosest sense, be forced into the Lab View. They are (treated as) objective, ongoing historical processes, tested indirectly via their input into other processes; they are neither prepared in some state at the beginning, nor measured at the end, and indeed in many cases they are ongoing.

In fact, this is not an issue specific to cosmology. The luminosity of the Sun, for instance, is determined in part via quantum mechanics: in particular, via the quantum tunneling processes that control the rate of nuclear fusion in the Sun's core as a function of its mass and composition. We can model this fairly accurately and, on the basis of that model, can deduce how the Sun's luminosity has increased over time. Astrophysicists pass that information to climate scientists, geologists, and paleontologists, who feed it into their respective models of prehistoric climates, geological processes, and ecosystems. All good science — but only in the most Procrustean sense can we realistically regard a successful fit to data in a paleoclimate model as being a measurement of the nuclear fusion processes in the Sun a billion years ago.[4]

Issues of this kind abound whenever we apply quantum theory outside stylised lab contexts. (Is the increased incidence of cancer due to Cold War nuclear-weapons tests a quantum measurement of the decay processes in the fallout products of those tests?) In each case we seem to have extracted objective facts about the unobserved world from quantum theory, not merely to be dealing with a mysterious microworld that gets its meaning only when observed. But they are particularly vivid in cosmology, which is a purely observational science, and a science chiefly concerned not with repeating events in the present but with the historical evolution of the observed Universe as a whole.

As perhaps the most dramatic example available — and probably the most important application of quantum theory in contemporary cosmology — consider the origin of structure in the Universe. Most of that story is classical: we posit a very small amount of randomly-distributed inhomogeneity in the very-early Universe, and then plug that into our cosmological models to determine both the inhomogeneity in the cosmic microwave background and the present-day distribution of galaxies. The latter, in particular, requires very extensive computer modelling that takes into account astrophysical phenomena on a great many scales; it cannot except in the most indirect sense be regarded as a 'measurement' of primordial inhomogeneity. Quantum theory comes in as a proposed source of the inhomogeneity: the posited scalar field (the 'inflaton field') responsible for cosmic inflation is assumed to be in a simple quantum state in the pre-inflationary Universe (most commonly the ground state) and quantum fluctuations in that ground state, time-evolved through the inflationary era, are identified with classical inhomogeneities. Quantum-mechanical predictions thus play a role in our modelling of the Universe's history, but not a role that the Lab View seems remotely equipped to handle.

6 The Decoherent View

The last section may have read as a call to arms: the Lab View is inadequate, so we must urgently seek an alternative way to do quantum mechanics! But that would miss the point: as with the previous section's discussion of experimental physics beyond the Lab View, physicists manifestly *are* doing quantum mechanics in these regimes, so they must *already* have a method for applying it that goes beyond the Lab View.

In fact, the method is fairly obvious, and fairly similar to that we used in section 4. The probability distribution over certain degrees of freedom — solar energy density, radiation rate, modes of the inflaton field — is simply treated as objective, as a probability distribution over actually-existing facts, and not merely as something that is realised when an experiment is performed. So we

---

[4] Philosophers may recognise this as an instance of Quine's classic objection to logical positivism (Quine 1951) — the empirical predictions of particular applications of quantum mechanics cannot be isolated from the influence of myriad other parts of our scientific worldview.

can say, for instance, not merely that a given mode of the primordial inflaton field *would have had* probability such-and-such of having a given amplitude if we were to measure it (whatever that means operationally), but that it *actually did have* probability such-and-such of that amplitude.

Now, it's tempting to imagine extending this objective take on quantum probabilities to *all* such probabilities: to interpret a quantum system as having some objectively-possessed value of every observable, and the quantum state as simply an economical way of coding a probability distribution over those observables. But this cannot be done. A collection of formal results — the Kochen-Specker theorem (Kochen and Specker 1967; Bell 1966, Redhead 1987, pp.119–152, Mermin 1993); Gleason's theorem (Gleason 1957, Redhead 1987, pp.27–9, Peres 1993, pp.190–195 Caves *et al* 2004); the Bub-Clifton theorem (Bub and Clifton 1996; Bub, Clifton, and Goldstein 2000); the PBR theorem and its relatives (Pusey, Barrett, and Rudolph 2011; Maroney 2012; Leifer 2014)) — establish that reading quantum mechanics along these lines — as bearing the same relation to some underlying objective theory as classical statistical mechanics bears to classical mechanics — is pretty much[5] impossible.

In fact, the central problem can be appreciated without getting into the details of these results. To take an objective view of some physical quantity is to suppose that the quantity has a definite value at each instant of time, so that we can consider the various possible *histories* of that quantity (that is: the various ways it can evolve over time) and assign probabilities to each. But the phenomena of interference means that this does not generically work in quantum mechanics. The quantum formalism for (say) the two slit experiment assigns a well-defined probability $P_1(x)$ to the history where the particle goes through Slit One and then hits some point x on the screen, and a similarly-well-defined probability $P_2(x)$ to it hitting point x via Slit Two, but of course the probability of it hitting point x at all (irrespectively of which slit it goes through) is not in general $P_1(x) + P_2(x)$. So the 'probabilities' assigned to these two histories do not obey the probability calculus. And things that don't obey the probability calculus are not probabilities at all.

At a fundamental level, the problem is that quantum mechanics is a dynamical theory about amplitudes, not about probabilities. The *amplitudes* of the two histories in the two-slit experiment sum perfectly happily to give the amplitude of the particle reaching the slit, but amplitudes are not probabilities, and in giving rise to probabilities they can cancel out or reinforce.

However, in most physical applications of quantum theory — and, in particular, in cosmology — we are *not* working 'at a fundamental level', which is to say that we are not attempting the usually-impossible task of deducing (far less interpreting) the evolution of the full quantum state over time. Rather,

---

[5] A more precise statement would be "impossible unless that underlying objective theory has a number of extremely pathological-seeming features." It is not universally accepted that this rules out such theories, though; see, e. g. , Spekkens (2007) and Leifer (2014) for further discussion. From the perspective of this article, such strategies share with the modificatory strategies of section 2 the feature that they require us to redo post-1930s physics, and so are (to say the least) not suitable in their current form to make sense of contemporary cosmology.

we are interested in finding higher-level, emergent dynamics, whereby we can write down dynamical equations for, and make predictions about, certain degrees of freedom of a system without having to keep track of all the remaining degrees of freedom. In the examples of the previous section, for instance, we have considered:

- The robust relations between macrostates of measuring devices and states of the system being measured, abstracting over the microscopic details of the measuring devices
- The bulk thermal properties of the core of the Sun, abstracting over the vast number of microstates compatible with those bulk thermal properties
- The low-wavelength modes of the inflaton field which are responsible for primordial inhomogeneities, abstracting over the high-wavelength degrees of freedom and the various other fields present.

In each case, we can derive from the quantum-mechanical dynamics an autonomous system of dynamical equations for these degrees of freedom. In each case, we can also derive from the Born Rule a time-dependent probability distribution over the values of those degrees of freedom. And in each case, that probability distribution defines a probability over histories that obeys the probability calculus. In each case, then, we are justified — at least formally, if perhaps not philosophically — in studying the autonomous dynamical system in question as telling us how these degrees of freedom are actually evolving, quite independently of our measurement processes.

This view of quantum physics is, in effect, the Objective View, but applied not to our fundamental physics but to certain higher-level dynamical theories emergent from that physics. It allows us to derive the validity of the Lab View in the particular context of well-controlled experiments. And it is, in practice, what is used — tacitly or explicitly — in applications of quantum theory that go beyond the confines of the Lab View.

Historically speaking, the first explicit statement of this view was by Gell-Mann and Hartle (1993). They and others[6] and had already developed a formalism — the 'consistent histories' or 'decoherent histories' formalism — that gives a criterion for when the quantum probability distribution over a given degree of freedom obeys the probability calculus. But, as they recognised, this fact by itself is not nearly strong enough to justify interpreting that degree of freedom objectively (indeed, and as stressed by Dowker and Kent (1996), the decoherent-history condition will generically be satisfied by a very large number of observables, most of which are pathologically non-classical). What is required — again, speaking just pragmatically and postponing the deeper interpretative questions — is not just a kinematics that can be described probabilistically, but a probabilistic dynamics. And that requires the behaviour of the objectively-interpretable degrees of freedom to be autonomous, to satisfy its own self-pp

---

[6] See in particular Gell-Mann and Hartle (1989), Griffiths (1984) and Omnes (1988); for more detailed references and a general review, see Halliwell (1995).

Borrowing Gell-Mann and Hartle's terminology, I call this emergently-objective view of quantum physics the Decoherent View.

7 Conditions for decoherence

Physical experience has provided us with a fairly good understanding of how to treat the dynamics of some subset $\mathcal{S}$ of the degrees of freedom of a large complex system. There are three possibilities. At one extreme these degrees of freedom may be dynamically decoupled from the remaining degrees of freedom, so that the equations of motion decompose into entirely separate equations for the two sets of dynamical variables. This in turn might occur because the degrees of freedom in $\mathcal{S}$ represent some subsystem spatially isolated from other systems; it might also occur where special features of the dynamics disconnect them from other degrees of freedom (as occurs, for instance, for the various modes of a free field, or for the centre-of-mass degree of freedom of a collection of particles). Call this case Full Autonomy.

At the other extreme — No Autonomy — it may be that the behaviour of degrees of freedom in $\mathcal{S}$ cannot be determined without fine-grained knowledge of the remaining degrees of freedom. In this case, the degrees of freedom in $\mathcal{S}$ cannot be said to represent a dynamically autonomous subsystem at all, but are just a partial description of a complex interacting whole.

The intermediate case is that described by non-equilibrium statistical mechanics. Here, typically, the degrees of freedom in $\mathcal{S}$ represent some coarse-grained, collective properties of a very complex system, and while the influence of the residual degrees of freedom on $\mathcal{S}$ cannot be ignored, those degrees of freedom are so numerous, and individually so insignificant in their effects, that we can get away with taking a statistical average over their effects rather than tracking each one individually. In this case — Statistical Autonomy — it remains possible to write down a closed-form dynamical equation of motion for $\mathcal{S}$, but that equation of motion is not the same as would apply if the interactions between $\mathcal{S}$ and the rest of the system were ignored. For instance, the effect of air resistance on a moving body can be analysed this way by averaging over the air molecules; so can Brownian motion, by using a probabilistic discription of the molecules that kick the pollen grain around. The validity of statistical autonomy is usually not absolute: some general assumptions have to be made about the initial state of the degrees of freedom not included in $\mathcal{S}$ (indeed, the equations of motion of statistically autonomous systems are typically not time-reversal invariant, which as a matter of logic requires some kind of state restriction). But once such general assumptions are made, the dynamics of $\mathcal{S}$ can be studied without further fine-grained information about other degrees of freedom.

Both Full Autonomy and Statistical Autonomy can give rise, in the right circumstances, to decoherence.[7] To begin with Full Autonomy: if a system's

---

[7] For detailed references to the decoherence literature, see Joos *et al* (2003) or Schlosshauer (2006).

Hamiltonian is quadratic in canonical position and momentum variables (as for a free particle, a harmonic oscillator or system of coupled harmonic oscillators, or a free field) then the quantum probability distributions over these variables evolve exactly as would be predicted from the classical equations of motion. For instance, the spreading out of a free particle wavefunction is formally identical to the spreading out of a probability distribution over the position and momentum of a free classical particle. (A helpful way to visualise this is through the Wigner function representation (Wigner 1932) which represents a quantum state as a real (albeit sometimes negative) function on phase space; the Schrödinger equation, in this representation, is given by the Liouville equation of classical physics supplemented by quantum correction terms that vanish if the Hamiltonian is quadratic. See Zurek and Paz (1995) or Wallace (2012, ch.3) for further discussion.)

However, this route to decoherence is somewhat delicate. Non-quadratic terms in a system's self-Hamiltonian can break the quantum-classical correspondence on relatively short timescales (cf (Zurek and Paz 1995)), and interactions with other systems, such as measurement devices, will also break it (consider the apparently-classical free-particle wave-packet after it interacts with a double slit and then a particle-detector screen). Robust decoherence requires more.

More is provided in a large fraction of those contexts where Statistical Autonomy holds. The central idea here is that the large number of residual degrees of freedom being treated statistically are constantly recording the state of the system, in the sense that their own states are strongly affected by the state of the system. (In the case of Brownian motion, for instance, collisions between an individual gas particle and the pollen grain have only minor effects on the state of the grain, but dramatic effects on the state of the gas particle.) This constant measurement (a) picks out a particular 'preferred basis' with which the measurement occurs (typically a basis of narrow wavepackets in nonrelativistic particle mechanics, or a coherent-state basis in (bosonic) field theory), and (b) ensures that the probability distribution over the observables defining that preferred basis evolves in a way that can be interpreted as a classical stochastic process — i. e., in a way where interference is almost entirely suppressed.

This form of 'environment-induced' decoherence is highly robust against disruption by external measurement processes. If the environment is constantly measuring the system with respect to a given basis, on timescales much shorter than any realistic human interferences, that's the only basis with respect to which it's physically viable to measure the system, and that additional measurement will make no further difference to a system's evolution.

It's helpful to distinguish two sorts of 'environment' in environment-induced decoherence. The most commonly discussed is a physically-external environment — in nonrelativistic cases, the atmosphere, stellar radiation, or even the microwave background radiation are all appealed to. But it is equally possible for a system's own small-scale degrees of freedom to decohere it. Schrödinger's unfortunate cat, for instance, has autonomous dynamics (the general behaviour of cats can be predicted, at least qualitatively, without fine-grained information on their exact microstate) but the microscopic degrees of freedom of the cat are quite capable of recording, quickly and redundantly, whether it is dead or not, without any need for an *external* environment.

8 Two models of decoherence: mundane and cosmological

As an illustrative model where all these mechanisms of decoherence play out, consider the classic case of a needle (assumed rotationally invariant at least in its macroscopic shape) balanced exactly on its end. In (highly implausible[8]) idealisation, we can treat the quantum state of the needle's centre of mass as a rotationally-invariant Gaussian pure state, placed exactly at the unstable equilibrium point. Then:

1. The Hamiltonian for that centre of mass can be approximated, close to the equilibrium point, as an upside-down harmonic oscillator. As long as we are interested only in the evolution of the needle in isolation, then, it can formally be treated as classical. The wave-packet will evolve into an expanding ring, and the evolution of the probability distributions defined by the ring is exactly that of a classical ensemble of needles distributed, symmetrically, close to the equilibrium point. However, the system remains in a coherent superposition and subsequent interactions could reveal that — were it not that

2. The atomic lattice of the needle is not rotationally invariant at the micro level and (while I have not analysed an explicit model) it is realistically almost certain that higher-frequency vibrational modes in the needle will couple to its centre of mass so as to quickly and redundantly record the position of the latter in the states of the former; and

3. Even if the internal degrees of freedom of the needle are entirely neglected, interactions between the needle and the atmosphere, or with ambient light, wil record the position of the needle's centre of mass on very short timescales.

As a result, the needle can be treated for all intents and purposes (that is: for the purposes of any measurement we might make of it) just as if it were a classical needle, with all the probabilities indistinguishable from those that would arise from classical ignorance about its initial position.

In fact, pretty much this same structure can be seen if we look at decoherence in the inflation era. Standard treatments (see, e. g. , Weinberg 2008) consider a free scalar inflaton field, which can be analysed into modes each of which may be treated as an inverted simple harmonic oscillator. Then:

---

[8] For macroscopic needles, any realistic process will prepare the needle in a highly mixed state; if we somehow managed to prepare it in a pure state after all, we would not realistically be able to position it exactly centered on the equilibrium point.

1. The evolution of the probability distribution for level of excitation of that mode can be treated, as long as the free-field assumption is valid, as if it were a classical probability distribution, though this by itself is not a robust route to classicality (Guth and Pi 1985; Polarski and Starobinsky 1996);
2. Assuming that the inflaton field is not exactly free, then interaction terms between the long-wavelength inflaton-field modes that determine the observed microwave-background-radiation fluctuation and which control structure formation, and the far-more-numerous shorter-wavelength modes, can be expected to rapidly decohere the former with respect to a basis of coherent states — i. e., of wavepackets around definite configurations of the long-wavelength part of the inflaton field.(Lombardo and Nacir 2005)
3. Interactions with other dynamical degrees of freedom (such as electromagnetic or gravitational radiation scattered by inhomogeneities) will also serve to decohere the long-wavelength modes with respect to that basis.(Calzetta and Hu 1995; Keifer, Polarski, and Starobinsky 1998)

So the Decoherent View can be consistently applied to the low-wavelength part of the inflaton field, just as with the needle. The recipe for using quantum mechanics that is — I have claimed — tacit in non-cosmological uses of the theory also works fine in its most important cosmological application.

9 Interpreting quantum theory in the cosmological context

To sum up: the Lab View of quantum mechanics, in which measurement is a primitive and an external observer is essential, is inadequate even for non-cosmological applications of quantum theory and doubly inadequate in cosmology. In fact we make sense of quantum theory outside lab contexts by applying the Objective View not to the theory as a whole, but to higher-level theories derived from it, in which the various processes that lead to decoherence have suppressed interference and enabled us to interpret the evolving probabilities of those theories in a classical fashion. In this context, the traditional name for the problem of how to interpret quantum theory — the "Measurement Problem" is out of place, for we can state the theory formally, and extract empirical content from it, without any talk of measurement.

So what is that problem of interpretation? It is this: to get an understanding of quantum theory — or a modification thereof — as an objective account of the world — as something to which the Objective View can after all be applied — or else to find some satisfactory way of thinking about quantum theory other than through the Objective View. And either way, that understanding must reproduce the Decoherent View, since that is how quantum theory makes contact with experiment.

We have already considered the modificatory strategies. At least in cosmology, they have to be thought of as incomplete research programs, not as complete alternatives awaiting philosopihical consideration. The physics of the inflaton field, in particular, is an application of *quantum gravity* (albeit in its perturbative form): the non-dynamical part of the metric field fluctuates with the inflaton field (cf Weinberg (2008, pp.470–474)), and the scattering of gravitons off the inflaton field is a coupling of gravitational degrees of freedom to matter. No extant dynamical-collapse or hidden-variables theory, to the best of my knowledge, has even come close to a (non-phenomenological) account of this physics.

What about alternatives to the Objective View? Sometimes the famous[9] instruction to "shut up and calculate" is taken as an alternative interpretation. But this is to misread the instruction. It is not a proposed way to understand quantum mechanics, but an instruction not to try — or, more charitably, as a statement that the speaker does not want to try. There is nothing wrong with that: not everyone working in quantum mechanics needs to try to solve its conceptual problems, and indeed we would understand the theory conceptually much less well than we do if no-one had just ploughed ahead and calculated. But as far as the interpretation of quantum mechanics is concerned, true adherents of "shut up and calculate" are properly silent. They are calculating, and shutting up.

A genuine alternative to the Objective View would be a positive philosophy of science that could be successfully applied to quantum theory and from which the validity of the Decoherent View would follow. One possibility would be to make sense of the Decoherent View on its own terms. In particular, the 'dappled world' approach of Nancy Cartwright (1983, 1999) abandons any idea that theories in physics should be seen as all derivable from one underlying theory, and treats the subject in piecemeal fashion, with each particular physics model applicable to certain systems and not in general extendible beyond those systems. (Ladyman and Ross (2013) explore a similar approach.) In approaches of this kind, the search for a systematic interpretation of quantum mechanics may simply be misconceived: we should instead interpret, on their own merits, each of the various higher-level theories that emerge from decoherence. Here we maintain the virtues of the Objective View at the cost of abandoing the search for unification.

Alternatively, we could look for some unified understanding of quantum theory that avoids the observer-independent, third-party-science approach of the Objective View but that breaks loose from the lab-based confines of earlier attempts at such an understanding. Carlo Rovelli's relational approach to quantum theory, for instance (Rovelli 2004, pp.209–222) holds onto the idea that quantum theory is always a description of one system from the perspective of another system, but drops the Copenhagen requirement that that second system is automatically a classical measurement device.

---

[9] It is usually attributed to Feynman, but David Mermin appears to be the real author; cf Mermin (2004).

Space does not permit me to do justice to these subtle and well-motivated ideas — but they share a common problem. Namely: this doesn't actually seem to be how we do quantum mechanics in practice. The various high-level theories derived from underlying physics via decoherence do indeed seem to be *derived*, not simply postulated independently of the physics: the whole point of inflation, in particular, is that it offers an underlying quantum derivation for the already-known phenomenological description of primordial inhomogeneity — and a derivation, indeed, that makes quantitative predictions for the free parameters in that description. And that "underlying physics" seems to be treated as a closed physical system, evolving unitarily and described without any appeal to other systems.

And that brings us back to the Objective View — and to the Everett interpretation. For a Universe objectively described by unitary quantum mechanics, and with a dynamics like that of our Universe, will be such that various of its subsystems will have autonomous high-level dynamics with the formal structure of a classical probabilistic theory. And in particular, the dynamics of any process of measurement will have to be such an 'autonomous high-level dynamics', either because (as in cosmology) the system being measured is thoroughly decohered by already-present interactions, or because (as in lab-based physics) the act of coupling a microscopic degree of freedom to a macroscopic lab-device degree of freedom brings into being such a dynamics. Everett's proposal — which originated precisely in the desire for an approach to quantum mechanics adequate for cosmology (Everett 1957) — is simply to treat the process of measurement physically in this way, so that a unitarily-evolving Universe appears, with respect to those regimes accessible to any physically-realised observer or measurement process, as a classical system with probabilistic dynamics. That is: the Everett interpretation is simply the reaffirmation of the Objective View of physics as valid for unmodified quantum physics, together with a detailed analysis of the physics of the quantum-classical transition and of the measurement process so as to show that the apparent impossibility of that Objective view is *only* apparent.

This is not a cost-free move. The reason why the state of the low-wavelength modes of the inflaton field can be treated as a classical probabilistic mixture is because they are thoroughly entangled with innumerable other degrees of freedom. But there is no precise point in the system's evolution where that entanglement becomes so thorough that interference exactly vanishes, no preferred instant at which we can imagine adding a non-unitary 'collapse' to the theory that eliminates all but one term in the superposition of low-wavelength field values. So taking the Objective View towards the quantum theory of the early Universe requires us to take the *full* quantum state of that early Universe as a faithful representation of the underlying physics. Since that full state is a superposition of rapidly-decohering terms only one of which we take to describe the actually-observed universe, but the formalism in no way picks that one term out from the others, this amounts to treating quantum theory as describing a vast ensemble of classical universes — an 'emergent multiverse', as I have elsewhere called it (Wallace 2012), since the 'universes' emerge from the formalism only in the regime in which decoherence applies and in which we can formally treat their mod-squared amplitudes as probabilities.

Since on a formal, technical level 'the Everett interpretation' is just the observer-independent, decoherence-based quantum theory, almost by definition its problems are philosophical rather than technical. The most-discussed include:

Micro-ontology: I have spoken of classical physics (or, more properly, the many instances of classical physics applicable to many different degrees of freedom) as 'emerging' from the underlying quantum description. But both the cohesion of that underlying description, and of the process of emergence, have been questioned; in particular, a number of philosophers (Allori *et al* 2008, Maudlin 2010; Allori 2013, Esfeld *et al* 2014, Esfeld, Deckert, and Oldofredi 2015) have advocated a very conservative take on 'fundamental' physics in which it necessarily makes contact with empirical data through claims about bodies with definite locations in space and time. On such an account, superpositions of different matter distributions — let alone of spacetime structure itself — are difficult to make sense of. (See Wallace (2010), or Wallace (2012, chapters 2 and 8), for my approach to these concerns; see also Bacciagaluppi and Ismael (2015, pp.146–7).)

Probability: "Probabilities" in Everett-interpreted quantum mechanics are Born-rule-calculated weights of terms in a superposition, in those regimes where decoherence means that those weights evolve without interference effects being significant. It's fair to say that this is radically different from any previously-entertained approach to probability (such as probabilities as relative frequencies, or as expressions of subjective ignorance, or as primitive) and critics (such as Albert (2010), Price (2010) and Kent (2010)) have claimed that it is unintelligible to treat these weights as probabilities in any more than a formal sense. Responses fall into two categories: positive responses (Farhi, Goldstone, and Gutmann 1989; Deutsch 1999; Zurek 2005; Greaves and Myrvold 2010, Wallace 2012, ch.5, Carroll and Sebens 2013) that attempt to demonstrate that the branch weights are (or at least play the role of) probability, and negative responses (Saunders 1995, 1998, Papineau 1996, 2010, Wallace 2012, ch.4) that argue that probability is at any rate *no more* conceptually puzzling in quantum than in classical mechanics.

Extravagance: The most common jibe about the Everett interpretation in informal discussions is that it is unreasonable and even unscientific for a theory to postulate such an extravagant ontology, especially as so little of it is observable. (Though, as I discuss in more detail in Wallace (2012, chapter 1), taking unobservable consequences of our best theories seriously is routine in science, and as Tegmark (2007) stresses, the extravagance of the Everettian multiverse is actually pretty mild by the standards of modern cosmology!)

10 Conclusion: The conservatism of the Everett interpretation

Detailed engagement with any of the objections above lies well beyond this article and I can only refer the reader to the references I provide. I want to conclude this discussion on a different point: that whether or not the Everett interpretation *itself* makes sense, it is about the only approach to quantum mechanics that makes sense *of* the way quantum mechanics is used in cosmology. Physicists working on the inflaton field — or on any other application of quantum theory outside some stylised lab contexts — do not work with (usually do not even have!) alternative theories to quantum mechanics; they do not adopt perspectival approaches to quantum theory where the degrees of freedom studies are always relativised to others; they do not treat higher-level processes as phenomenological primitives. They treat the system they are studying, in the first instance, as unitarily evolving and closed; they extract from that theory, more or less explicitly, higher-level dynamics and higher-level descriptions to which the Born rule can be consistently applied, and they connect to observations via that application.

Philosophers of science sometimes talk of the philosophical position *tacitly* held by practising scientists. They do not mean the answers that those scientists give to explicit questions about philosophy; as Tim Williamson (private conversation) once observed, the answers given tend to be whatever gets rid of the inquring philosopher most expediently! They mean the philosophical position that those practicing scientists would have to hold if their scientific activity were to make sense. (In this sense, it is fairly widely accepted — even by philosophers who are themselves sceptical about scientific realism — that scientists themselves are realists in at least some extended sense.)

The practice of those physicists who apply quantum theory to cosmology *makes sense* if those physicists are tacitly committed to the Everett interpretation (whether or not that commitment fails to recognise some deep pathology in that interpretation). So far as I can see, it *does not make sense* on any other interpretation of quantum mechanics. So cosmologists, when applying quantum mechanics, are tacitly committed to the Everett interpretation. (And in some cases, such as Hawking (1976),Weinberg (2002, p.233), or the quotation from Hartle that began this chapter, the commitment becomes explicit.)

Are they right to be so committed? In my view, yes, but I have not argued for it here. But to think otherwise seems to commit one to advocating a significant revision of physical practice, and perhaps of physics itself, on the basis of *a priori* philosophical objections and in the absence of the more naturalistic reasons that drove Everett, Zeh, Zurek, Gell-Mann, Hartle, Halliwell and many others to develop an approach to using quantum physics that avoids the deficiencies of the Lab View. My confidence in that approach's philosophical coherence is grounded only partially in the specific arguments, and partially in the observation that history has not been kind to attempts to revise scientific practice on *a priori* grounds.

References

Albert, D. Z. (2010). Probability in the Everett picture. In S. Saunders, J. Barrett, A. Kent, and D. Wallace (Eds.), *Many Worlds? Everett, Quantum Theory and Reality*. Oxford: Oxford University Press.

Albert, D. Z. and B. Loewer (1996). Tails of Schrödinger's Cat. In R. Clifton (Ed.), *Perspectives on Quantum Reality*, pp. 81–92. Dordrecht: 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.

Bacciagaluppi, G. and J. Ismael (2015). Review of the emergent multiverse. *Philosophy of Science 82*, 129–148.

Bedingham, D. J. (2010). Relativistic state reduction dynamics. Online at http://arxiv.org/abs/1003.2774.

Bedingham, D. J. (2011). Relativistic state reduction model. Online at http://arxiv.org/abs/1103.3974.

Bell, J. (2004). Against measurement. In *Speakable and Unspeakable in Quantum Mechanics (Second Edition)*, pp. 213–231. Cambridge University Press.

Bell, J. S. (1966). On the problem of hidden variables in quantum mechanics. *Reviews of Modern Physics 38*, 447–452. Reprinted in Bell (1987), pp. 1–13.

Bell, J. S. (1987). *Speakable and Unspeakable in Quantum Mechanics.* Cambridge: Cambridge 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.

Bub, J. and R. Clifton (1996). A uniqueness theorem for "no collapse" interpretations of quantum mechanics. *Studies in the History and Philosophy of Modern Physics 27*, 181–219.

Bub, J., R. Clifton, and S. Goldstein (2000). Revised proof of the uniqueness theorem for 'no collapse' interpretations of quantum mechanics. *Studies in the History and Philosophy of Modern Physics 31*, 95.

Calzetta, E. and B. Hu (1995). Quantum fluctuations, decoherence of the mean field, and structure formation in the early universe. *Physical Review D 52*, 6770–6788.

Carroll, S. M. and C. T. Sebens (2013). Many worlds, the born rule, and self-locating uncertainty. In D. Struppa and J. Tollaksen (Eds.), *Quantum Theory: a Two-Time Success Story, Yakir Aharonov Festschrift*, pp. 157. Berlin: Springer-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., C. Fuchs, K. Manne, and J. Renes (2004). Gleason-type derivations of the quantum probability rule for generalized measurements. *Foundations of Physics 34*, 193.

Colin, S. and W. Struyve (2007). A Dirac sea pilot-wave model for quantum field theory. *Journal of Physics A 40*, 7309–7342.

Cushing, J. T., A. Fine, and S. Goldstein (Eds.) (1996). *Bohmian Mechanics and Quantum Theory: An Appraisal*, Dordrecht. Kluwer Academic Publishers.

Deutsch, D. (1997). *The Fabric of Reality*. London: Penguin.

Deutsch, D. (1999). Quantum theory of probability and decisions. *Proceedings of the Royal Society of London A455*, 3129–3137.

DeWitt, B. and N. Graham (Eds.) (1973). *The many-worlds interpretation of quantum mechanics*. Princeton: Princeton University Press.

Dirac, P. (1930). *The Principles of Quantum Mechanics*. Oxford University Press.

Dowker, F. and A. Kent (1996). On the consistent histories approach to quantum mechanics. *Journal of Statistical Physics 82*, 1575–1646.

Esfeld, M., D.-A. Deckert, and A. Oldofredi (2015). What is matter? the fundamental ontology of atomism and structural realism. In A. Ijjas and B. Loewer (Eds.), *Philosophy of Cosmology*. Oxford: Oxford University Press.

Esfeld, M., D. Lazarovici, V. Lam, and M. Hubert (2014). The physics and metaphysics of primitive stuff. Online at http://arxiv.org/abs/1411.7545.

Everett, H. I. (1957). Relative State Formulation of Quantum Mechanics. *Review of Modern Physics 29*, 454–462. Reprinted in DeWitt and Graham (1973).

Farhi, E., J. Goldstone, and S. Gutmann (1989). How probability arises in quantum-mechanics. *Annals of Physics 192*, 368–382.

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.

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 R. Schack (2015). QBism and the Greeks: why a quantum state does not represent an element of physical reality. *Physica Scripta 90*, 015104.

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.

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

Goldstein, S., W. Struyve, and R. Tumulka (2016). The Bohmian approach to the problems of cosmological quantum fluctuations. In A. Ijjas and B. Loewer (Eds.), *Philosophy of Cosmology.* Oxford University Press.

Greaves, H. and W. Myrvold (2010). Everett and evidence. In S. Saunders, J. Barrett, A. Kent, and D. Wallace (Eds.), *Many Worlds? Everett, Quantum Theory and Reality.* Oxford: Oxford University Press.

Griffiths, R. (1993). Consistent interpretation of quantum mechanics using quantum trajectories. *Physical Review Letters 70*, 2201–2204.

Griffiths, R. B. (1984). Consistent histories and the interpretation of quantum mechanics. *Journal of Statistical Physics 36*, 219–272.

Griffiths, R. B. (1996). Consistent histories and quantum reasoning. *Physical Review A 54*, 2759–2773.

Guth, A. H. and S.-Y. Pi (1985). Quantum mechanics of the scalar field in the new inflationary universe. *Physical Review D 32*, 1899–1920.

Halliwell, J. J. (1995). A review of the decoherent histories approach to quantum mechanics. *Annals of the New York Academy of Sciences 755*, 726–740.

Hawking, S. W. (1976). Black holes and thermodynamics. *Physical Review D 13*, 191–197.

Holland, P. (1993). *The Quantum Theory of Motion.* Cambridge: Cambridge University Press.

Home, D. and M. A. B. Whitaker (1997). A conceptual analysis of quantum Zeno: Paradox, measurement and experiment. *Annals of Physics 258*, 237–285.

Joos, E., H. D. Zeh, C. Kiefer, D. Giulini, J. Kupsch, and I. O. Stamatescu (2003). *Decoherence and the Appearence of a Classical World in Quantum Theory* (2nd Edition ed.). Berlin: Springer.

Keifer, C., D. Polarski, and A. E. Starobinsky (1998). Quantum-to-classical transition for fluctuations in the early universe. *International Journal of Modern Physics D7*, 455–462.

Kent, A. (2010). One world versus many: the inadequacy of Everettian accounts of evolution, probability, and scientific confirmation. In S. Saunders, J. Barrett, A. Kent, and D. Wallace (Eds.), *Many Worlds? -Everett, Quantum Theory, and Reality*, 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.

Ladyman, J. and D. Ross (2013). The world in the data. In D. Ross, J. Ladyman, and H. Kincaid (Eds.), *Scientific Metaphysics*, pp. 108–150. Oxford: Oxford University Press.

Leifer, M. (2014). Is the quantum state real? an extended review of ψ-ontology theorems. *Quanta 3*, 67–155.

Lombardo, F. C. and D. L. Nacir (2005). Decoherence during inflation: The generation of classical inhomogeneities. *Physical Review D 72*, 063506.

Maroney, O. (2012). How statistical are quantum states? Available online at http://arxiv.org/abs/1207.6906.

Maudlin, T. (2010). Can the world be only wavefunction? In S. Saunders, J. Barrett, A. Kent, and D. Wallace (Eds.), *Many Worlds? Everett, Quantum Theory, and Reality.* Oxford: Oxford University Press.

Mermin, N. D. (1993). Hidden variables and the two theorems of John Bell. *Reviews of Modern Physics 65*, 803–815.

Mermin, N. D. (2004). Could Feynman have said this? *Physics Today 57*, 10.

Misra, B. and E. C. G. Sudarshan (1977). The Zeno's paradox in quantum theory. *Journal of Mathematical Physics 18*, 756.

Newton-Smith, W. S. (1981). *The Rationality of Science.* London: Routledge.

Omnes, R. (1988). Logical reformulation of quantum mechanics. i. foundations. *Journal of Statistical Physics 53*, 893–932.

Omnes, R. (1992). Consistent interpretations of quantum mechanics. *Reviews of Modern Physics 64*, 339–382.

Omnes, R. (1994). *The Interpretation of Quantum Mechanics.* Princeton: Princeton University Press.

Papineau, D. (1996). Many Minds are No Worse than One. *British Journal for the Philosophy of Science 47*, 233–241.

Papineau, D. (2010). A fair deal for Everettians. In S. Saunders, J. Barrett, A. Kent, and D. Wallace (Eds.), *Many Worlds? Everett, Quantum Theory, and Reality.* Oxford: Oxford University Press.

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.

Polarski, D. and A. A. Starobinsky (1996). Semiclassicality and decoherence of cosmological perturbations. *Classical and Quantum Gravity 13*, 377–391.

Price, H. (2010). Probability in the Everett picture. In S. Saunders, J. Barrett, A. Kent, and D. Wallace (Eds.), *Many Worlds? Everett, Quantum Theory and Reality.* Oxford: Oxford University Press.

Psillos, S. (1999). *Scientific Realism: How Science Tracks Truth.* London: Routledge.

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

Quine, W. V. O. (1951). Two dogmas of empiricism. *Philosophical Review 60*, 20–43.

Redhead, M. (1987). *Incompleteness, Nonlocality and Realism: A Prolegomenon to the Philosophy of Quantum Mechanics.* Oxford: Oxford University Press.

Rovelli, C. (2004). *Quantum Gravity.* Cambridge: Cambridge University Press.

Saunders, S. (1995). Time, Decoherence and Quantum Mechanics. *Synthese 102*, 235–266.

Saunders, S. (1998). Time, Quantum Mechanics, and Probability. *Synthese 114*, 373–404.

Saunders, S. (2005). Complementarity and scientific rationality. *Foundations of Physics 35*, 347–372.

Schlosshauer, M. (2006). Experimental motivation and empirical consistency in minimal no-collapse quantum mechanics. *Annals of Physics 321*, 112–149.

Spekkens, R. W. (2007). In defense of the epistemic view of quantum states: a toy theory. *Physical Review A 75*, 032110.

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.

Tegmark, M. (2007). The multiverse hierarchy. In B. Carr (Ed.), *Universe or Multiverse?* Cambridge: Cambridge University Press.

Timpson, C. (2010). *Quantum Information Theory and the Foundations of Quantum Mechanics.* Oxford: Oxford University Press.

Tumulka, R. (2006). Collapse and relativity. In A. Bassi, T. Weber, and N. Zanghi (Eds.), *Quantum Mechanics: Are There Quantum Jumps? and on the Present Status of Quantum Mechanics*, pp. 340. American Institute of Physics Conference Proceedings. Available online at http://arxiv.org/abs/quant-ph/0602208.

Underwood, N. G. and A. Valentini (2015). Quantum field theory of relic nonequilibrium systems. *Physical Review D 92*, 063531.

Valentini, A. (2010). Inflationary cosmology as a probe of primordial quantum mechanics. *Physical Review D 82*, 063513.

von Neumann, J. (1955). *Mathematical Foundations of Quantum Mechanics.* Princeton: Princeton University Press.

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

Wallace, D. (2013). Inferential vs. dynamical conceptions of physics. Available online at http://arxiv.org/abs/1306.4907.

Weinberg, S. (2002). *Facing Up: Science and its Cultural Adversaries.* Cambridge, MA: Harvard University Press.

Weinberg, S. (2008). *Cosmology.* Oxford: Oxford University Press.

Wigner, E. (1932). On the quantum correction for thermodynamic equilibrium. *Physical Review 40*, 749–759.

Zurek, W. and J. P. Paz (1995). Quantum chaos: a decoherent definition. *Physica D 83*, 300–308.

Zurek, W. H. (2005). Probabilities from entanglement, Born's rule from envariance. *Physical Review A 71*, 052105. Available online at http://arxiv.org/abs/quant-ph/0405161.