

Decoherence and its role in the modern measurement problem

David Wallace*

October 31, 2011

Abstract

Decoherence is widely felt to have *something* to do with the quantum measurement problem, but getting clear on just *what* is made difficult by the fact that the “measurement problem”, as traditionally presented in foundational and philosophical discussions, has become somewhat disconnected from the conceptual problems posed by real physics. This, in turn, is because quantum mechanics as discussed in textbooks and in foundational discussions has become somewhat removed from scientific practice, especially where the analysis of measurement is concerned.

This paper has two goals: firstly (sections 1–2), to present an account of how quantum measurements are actually dealt with in modern physics (hint: it doesn’t involve a collapse of the wavefunction) and to state the measurement problem from the perspective of that account; and secondly (sections 3–4), to clarify what role decoherence plays in modern measurement theory and what effect it has on the various strategies that have been proposed to solve the measurement problem.

1 The myth of the “conventional interpretation”

Foundational discussions of quantum theory have a tendency to talk about the “conventional” or “orthodox” or “standard” interpretation of quantum mechanics. It’s not generally very clear what’s meant by this, but a rough stab might be that it consists of these two principles:

1. (the Measurement Algorithm) Observable quantities are represented by self-adjoint operators; the possible outcomes of a measurement of some observable are the eigenvalues of the associated operator; the probability of a given measurement outcome obtaining is given by the usual (Born) probability rule.
2. (the Projection Postulate) While in the absence of measurement a system evolves unitarily and deterministically, according to the Schrödinger

*Balliol College, Oxford University; david.wallace@balliol.ox.ac.uk

equation, when a measurement is made the system evolves stochastically, with its state vector being projected onto the eigensubspace corresponding to the actual measurement outcome. As such, the dynamics of quantum theory have a dual nature, with one evolution rule for non-measurement situations and one for measurement situations.

Often (including by Dirac (1930)), the second is derived from the first by consideration of repeat measurements: presumably (so goes the reasoning) two measurements in succession had better give the same result, and this can only be guaranteed if the Projection Postulate holds.

It's generally held that this "conventional interpretation" is profoundly unsatisfactory conceptually and philosophically, essentially because it treats "measurement" as a primitive term. Measurement, so the argument goes, is a physical process like any other, so (i) our theory should not contain assumptions which can be stated only by treating measurement as some primitive process, and (ii) whether a quantum system evolves according to one rule or another should not depend on whether we classify a physical process as a "measurement" or not.

My concern, however, isn't what makes the orthodox interpretation *unsatisfactory*; it's what makes it *orthodox*. To be sure: the orthodox interpretation is pretty much what we find in von Neumann's (1955) and Dirac's (1930) original presentations of the subject. To be sure: it's pretty much what we find in most textbooks on quantum mechanics.¹ And to be sure: it's pretty much what we find in most discussions of the quantum measurement problem. But it's not something which we actually find in use, much or at all, in mainstream applications of quantum theory.

We can see this pretty straightforwardly in the case of the Projection Postulate. The postulate tells us how a quantum system evolves *during measurement*, and this tells us immediately that it can only play a role in applications of quantum physics in situations where we want to analyse *repeated measurements*. If all we care about is the outcome of single measurements, the Measurement Algorithm tells us all we need to know.

But if the point of the Projection Postulate is to analyse repeated measurements, there's an embarrassing problem: the Postulate tells us successive measurements of the same quantity always give the same result, and this is straightforwardly false in many — perhaps most — experimental contexts. (We can see this particularly dramatically in quantum optics: when we make a measurement of a photon, the photon is typically *gone* at the end of the measurement.)²

The case of *continuous* measurements (like those performed by a Geiger counter) also demonstrates that the Projection Postulate is unacceptable as an analysis of quantum dynamics during the measurement process. The only way to analyse continuous measurements via the Postulate would be to treat it

¹See, for instance, Cohen-Tannoudji, Diu, and Laloë (1977), Landau and Lifshitz (1977), Rae (1992), and Sakurai (1994); Ballentine (1990) is an interesting exception.

²Bell and Nauenberg (1966) pointed this out some while ago, and introduced a distinction between "moral" measurements which are compatible with the projection postulate, and "immoral" measurements, which aren't; see Home and Whitaker (1997) for further discussion.

as the limiting case of increasingly frequent measurements. But the quantum Zeno effect (Misra and Sudarshan (1977); see Home and Whitaker (1997) for discussion) tells us that the result of any such limiting-case measurement is to freeze the system's evolution altogether, in flat contradiction of observed fact. (This was the original reason that Misra and Sudarshan, who did analyse continuous measurement in exactly this way, called the quantum Zeno effect a 'paradox'.)

So if the only way the Projection Postulate could figure into the practice of physics is via the case of repeated measurements, and if it is patently unacceptable for that task, then there can be *no* practical role for the Postulate in physics. And for what it's worth, crude sociological measures seem to back this up: a search for "projection postulate", "wave-function collapse" and the like in the archives of Physical Review turns up only a few hundred references, nearly all of which turn out to be (a) foundational discussions, (b) discussions of proposed alternatives to quantum theory, or (c) theoretical quantum-computation discussions, where "measurement" does indeed get treated as primitive. (For comparison, searches for terms like "state vector" or "Hilbert space" or "Schrodinger equation" typically turn up several tens of thousands of references.) It's also notable that in *relativistic* quantum mechanics, even defining the Projection Postulate is a delicate matter, because of relativity of simultaneity. One might then expect that textbooks on quantum field theory would need to address this point, but no such textbook of which I am aware does so; we can infer that the Projection Postulate is not much missed by practicing particle physicists or quantum field theorists.

However, I do not wish to suggest that we have *no* effective way to analyse the process of repeated measurement. Actually, we can analyse it just fine (where here "just fine" means "just fine if our goal is getting the empirically observed answer", not necessarily "just fine conceptually speaking".) Doing so, though, requires us to drop the idea that measurement, as conceived in the Measurement Algorithm, is a primitive process.

This should have been obvious in any case. Treating "measurement" as a primitive term is not just philosophically unsatisfactory: it's a non-starter if you actually want to build and use a measurement device. And since we do actually build and use measurement devices, we can't just be treating measurement processes as primitive!

To expand: textbook and foundational discussions of quantum mechanics often give the impression that measurement devices are inexplicable black boxes which we find scattered across the desert, each one stamped with a self-adjoint operator and equipped with a funnel to insert quantum systems and a dial displaying the outcomes. But of course, real measurement devices are designed very carefully by experimental physicists. It's not a primitive feature of them that they measure whatever it is that they measure: it's a deliberate design feature. Perhaps at one time that design proceeded by handwaving appeals to quantum-classical correspondence, but in modern experimental physics the design of a quantum-mechanical measurement device relies extensively on quantum theory itself.

How can this be done? The answer, pretty much invariably, is that we apply quantum theory — unitary quantum theory, with no Projection Postulate — to the measurement process itself, and then apply the Measurement Algorithm to the combined system of original-system-plus-measuring-device. This might seem like an infinite regress — doesn't applying the Algorithm presuppose yet one more measurement? — but in practice the regress can be ended when the system is macroscopically large, so that we can treat its macroscopic features — say, its centre-of-mass position or momentum — as straightforwardly observable properties.

In the particular case of repeated measurements, a multiply-conducted measurement will need multiple readout displays. The measurement algorithm, applied to the combined original-system-plus-apparatus quantum system, delivers a determinate probability prediction over those displays. In some situations, it predicts that the displays are perfectly correlated; in *these* situations, we could have got away with using the Projection Postulate. In most situations, the correlation will be either imperfect or absent altogether. Similarly, in the case of continuous measurements, analysing the actual measurement dynamics — using *unitary* quantum mechanics — and then applying the Measurement Algorithm will give us the probability of a decay (say) being registered as a function of time (and will, in general, tell us that the decay rate is not especially sensitive to how closely the system is monitored, except in extreme cases).

We are left with an understanding of the measurement problem that is a better fit to contemporary scientific practice. There is no conventional *interpretation* of quantum theory. What there is, is the *formalism* of unitary quantum theory, shorn of any interpretation. We go from quantum theory to empirical predictions, not by understanding unitary quantum theory as a description of physical reality, but by applying a Measurement Algorithm, in principle at the level of microphysics but, when pushed, ultimately at the level of macroscopic, straightforwardly observable, facts. But that Algorithm (a) seems ill-defined and shot through with terms that don't belong in fundamental physics (“macroscopic” and “straightforwardly observable” are only slightly more satisfactory than “measurement” as primitives), (b) relatedly, runs into the problem that even macroscopic observation processes are just more physical processes, and so shouldn't have special status, and (c) seems to block any attempt to understand quantum theory as a theory which tells us about Nature, and not just about what the dials and computer screens in our lab are going to do in any given situation.

From this point of view, it's also reasonably clear what would count as *solving* the problem. We either need to find some way of *understanding* unitary quantum theory so that it can give experimental predictions directly and not via the unsatisfactory Measurement Algorithm — or we need to *modify or supplement* unitary quantum theory so as to obtain a new theory which does not need the Measurement Algorithm.

2 Two concepts of state space

The account of the previous section may look unfamiliar. Normally (at least in foundational and philosophical discussions) the argument for a measurement problem is more direct: unitary quantum theory predicts macroscopic superpositions, which aren't observed, so unitary quantum theory must fail when applied to macroscopic systems (and presumably something like a Projection Postulate must be applied, albeit perhaps at the macroscopic level rather than for microscopic systems).

The basic assumption here is that Schrödinger-cat states like

$$\alpha |\text{cat alive}\rangle + \beta |\text{cat dead}\rangle \tag{1}$$

have to be understood as somehow describing a cat in an indefinite state of aliveness — a cat that somehow is alive and dead at the same time. But why are we required to think this way?

In the philosophy literature, this reading of (1) is often justified by appeal to the so-called

Eigenvalue-eigenvector link (E-E link): A quantum system possesses a definite value of an observable if and only if its state is an eigenstate of the associated operator; if so, the definite possessed value is the eigenvalue associated with that eigenstate.

According to the E-E link, the state I wrote schematically as $|\text{cat alive}\rangle$ can be understood as a state of a live cat because it lies in some subspace of the cat's Hilbert space corresponding to the “aliveness=1” subspace of the “aliveness operator”; similarly, the state $|\text{cat dead}\rangle$ is in the “aliveness=0” subspace and so is the state of a cat that is determinately not alive. The superposition of the two, lying in neither subspace, represents a cat that is neither alive nor dead.

This argument should not worry us, though, because the E-E link has no place in any serious discussion of the measurement problem. Like the Projection Postulate, it plays basically no role in physical practice (searching Physical Review's archives for the phrase and its variants gives precisely two hits); unlike the Projection Postulate, it seems to be purely an invention of philosophers, which does not appear in any quantum physics textbook of which I am aware. And this is not surprising, since in certain contexts it is fairly absurd, as can be seen easily by considering measurements of position. No realistic quantum wave-packet will remain in a bounded region of space for any non-zero length of time, unless kept there by infinite (and hence unphysical) potentials: in any physically realistic situation, some small section of the wavefunction will tunnel to arbitrarily large distances, arbitrarily quickly. So realistic quantum wavefunctions are spread right across space. Of course, macroscopic bodies can have wavefunctions that are, and that remain, extremely small outside some narrow region; wavefunctions like these are straightforwardly used in physics to represent localised bodies. But the fact remains that no such state lies in the eigenspace corresponding to any outcome of a position measurement, however

coarse-grained.³ So according to the E-E link, all realistic quantum systems have completely indefinite positions. If the E-E link were an inextricable part of quantum physics, this would count as a *reductio ad absurdum* of quantum theory; as it is, it just counts as a *reductio ad absurdum* of the E-E link.

So much for the E-E link. But even if *that* argument for the unacceptable weirdness of Schrödinger-cat states does not go through, there might perfectly well be *other* arguments to the same conclusion. To get more perspective on states like (1), it helps to think a bit about the concept of a state space in general.

The paradigm example of a state space is classical phase-space (I could tell essentially the same story with configuration space). States — points in phase space — represent the physical features of the system in question; different points in phase space represent physically distinct systems. But this is not the only kind of state space definable in classical physics: the space of *probability functions* on phase space can also be treated as a state space (call it distribution space). Mathematically, phase space and distribution space are just two sets of abstract mathematical objects with deterministic evolution equations: Hamilton’s equations in the one case, Liouville’s equations in the other. The central difference is conceptual: distinct points in distribution space do not represent systems in physically distinct states, but distinct probability distributions over physical states. In other words, there are two conceptions of state space available: physical, or probabilistic.⁴

One way to understand the difficulty of making sense of quantum states is that we seem to shift between these two conceptions. When dealing with states like 1 in practice, we treat them as probabilistic, conveying nothing more or less than “the cat is either alive or dead, and it has chance $|\alpha|^2$ of being in the first state, and $|\beta|^2$ of being in the second”. On this reading of the quantum state, there is nothing at all mysterious about (1), and no need to invoke the Projection Postulate. Indeed, something like the Projection Postulate emerges spontaneously from the ordinary calculus of probability: it is a commonplace that probability distributions evolve *both* under the ordinary dynamics of the system, *and* via probabilistic updating when we acquire more information. From that perspective, the “collapse” of (1) onto either $|\text{alive cat}\rangle$ or $|\text{dead cat}\rangle$ upon measurement is no more mysterious than the collapse of the classical probability distribution “Heads with probability 1/2, tails with probability 1/2” onto “Heads” or “Tails” upon measurement.

When dealing with the quantum states of microscopic systems, however, this

³In my view, failure to appreciate this point has caused much confusion in discussions of the so-called “problem of tails” in the GRW dynamical-collapse theory (cf Wallace (2008) and references therein): the “problem”, as generally discussed, has little to do with the GRW theory per se, but is a general feature of quantum mechanics.

⁴An alternative terminology due to Harrigan and Spekkens (2010) calls the physical conception *ontic* and the probabilistic conception *epistemic*. (In the case of the quantum state, this becomes ‘ ψ -ontic’ and ‘ ψ -epistemic’ — hence Chris Fuchs’s wonderfully cruel term (in conversation) for those who take the ontic view: “ ψ -ontologists”.) This terminology, though, suggests a particular reading of probability as a quantification of our ignorance, which I want to avoid being committed to.

straightforward probabilistic reading of the quantum state breaks down. Partly this is because the democracy of bases in Hilbert space allow a multitude of ways of expressing a state as a superposition of eigenstates: is the spin state

$$\alpha |+_z\rangle + \beta |-_z\rangle \tag{2}$$

to be interpreted as a probabilistic mixture of a particle with z -spin up and a particle with z -spin down, or as a different probabilistic mixture of x -spin up and x -spin down? (Or, somehow, mysteriously, as both at once?) But more crucially, the probabilistic reading simply fails to make sense of interference phenomena. Suppose that we set up a Mach-Zender interferometer, in which a beam of (say) photons is split into left(L) and right(R) beams by a half-silvered mirror and then re-interfered by means of another such mirror, and the resultant beams are fed into two detectors, A and B . We can easily establish, by blocking the R beam, that if the photon is originally in the L beam, it has a 50% chance of ending up at detector A after the second half-silvered mirror, and a 50% chance of ending up at detector B . We can easily establish that the same is true if it is originally in the R beam. But in that case, which beam it is in makes no difference to the outcome probabilities, so *any* probabilistic mixture of the two possibilities should lead to a 50% chance of each of the two results. And of course, this is not what happens: depending on the relative phase of the two beams, any outcome from 100% A to 100% B is possible.⁵

The two objections are related, of course. Interference phenomena can occur in quantum mechanics precisely because amplitudes have phases as well as magnitude. The probabilistic reading of a quantum state like (1) or (2) interprets the magnitudes of the amplitudes α and β as probabilities, but provides no interpretation of their phases. On this reading, replacing α with $\alpha \exp(-i\theta)$ should have no physical significance — but of course, doing so not only has “significance”, it has empirically detectable consequences.

This provides an alternative way to state the measurement problem:

We cannot consistently understand the state space of quantum theory *either* as a space of physical states, *or* as a space of probability distributions. Instead, we have to use one interpretation for microscopic physics and another for macroscopic physics. Furthermore, both the point at which we have to transition between the physical and probabilistic interpretation, and the basis with respect to which the probabilistic interpretation is to be specified, are defined only in an approximate, rough-and-ready way, which seems to make

⁵Nor will it do to say “quantum probabilities just behave differently from classical ones”. The point of the argument is that we cannot explain the observed phenomena allowing the two “possibilities” to influence one another. But if something that can have a dynamical effect on a real thing doesn’t thereby count as real, we lose our grip on reality. (Cf discussions in Deutsch (1997, chapter 4), Wallace (2011, chapter 10), or Anandan and Brown (1995).) The bottom line is that a microscopic superposition cannot be understood as merely probabilistic: both terms in the superposition represent physical features of the world. But if the same is true for the Schrödinger-cat state, the paradoxical nature of that state is not alleviated by the supposedly “probabilistic” reading.

essential use of terms like “macroscopic” which have no place in a fundamental physical theory.

It also provides a ready way to understand the various strategies that have been proposed to resolve the measurement problem. Firstly, there are the attempts to *solve* the problem, by replacing quantum theory by a new theory which can be understood in a consistent way. There are two basic strategies:

Hidden-variable theories These theories hold on to the dynamics of unitary quantum theory, but remove the probabilistic interpretation of the quantum state. Instead, some new physical entities are introduced, which are dynamically influenced by the quantum state but not vice versa. Probability is then introduced via a probability distribution over these “hidden variables”. The paradigm example is the de Broglie-Bohm pilot-wave theory, also called Bohmian mechanics,⁶ in which the quantum state is supplemented by a collection of point particles whose collective evolution is determined by the quantum state. (So-called ‘modal interpretations’⁷ also fit this category.)

Dynamical-collapse theories These theories⁸ try to make sense of the transition from a physical to a probabilistic reading of the quantum state not as an interpretative shift but as an objective physical process. This is generally done by introducing some version of the Projection Postulate that objectively removes macroscopic superpositions, and does so in a probabilistic manner. According to dynamical-collapse theories, the quantum state is always to be interpreted physically and its stochastic evolution deviates from the deterministic Schrödinger equation in certain circumstances — but it does so in such a way that the resultant probabilities are very close to those defined by the probabilistic reading of unitary quantum mechanics.

Secondly, there are the attempts to *dissolve* the problem, by finding a consistent interpretation of unitary quantum theory. There are only two real options here:

The quantum state is always probabilistic This approach tries to interpret quantum state space as systematically like the space of classical probability distributions. This has proved extremely difficult, though, for essentially the reasons given above: interference phenomena don’t seem to be understandable as probabilistic phenomena. More rigorously, the theoretical results of Bell (1966), Kochen and Specker (1967), and Gleason (1957), make it clear that any such strategy will have the apparently pathological feature of *contextuality*.⁹ For this reason, most¹⁰ adherents of

⁶See Cushing, Fine, and Goldstein (1996) and references therein.

⁷See Dieks and Vermaas (1998) and references therein.

⁸Examples include the GRW theory (Ghirardi, Rimini, and Weber (1986); see Bassi and Ghirardi (2003) for discussion) and proposals due to Penrose (1989); cf also Philip Stamp’s discussion in this volume).

⁹For detailed discussion, see, e.g., Redhead (1987).

¹⁰Not all: Spekkens (2007) is a counter-example.

this strategy have given up on the idea of interpreting the “probabilities” as probabilities of any microphysical even, and just read them directly as probabilities of measurement outcomes, treated as primitive. That is, they fall back on an instrumentalist reading of quantum theory. (For a clearly articulated example, see Fuchs and Peres (2000).)

The quantum state is always physical This approach, most famously associated with the name of Hugh Everett (1957), tries to interpret quantum state space as systematically like classical phase space. This deals straightforwardly with the problem of microscopic interference, but struggles with the problem of macroscopic superpositions. If Schrödinger-cat states like 1 are to be understood as physical states, what physical situation are they representing? The only coherent answer seems to be: a situation with two independently-evolving cats, one alive and one dead. But what justifies this interpretation, and where do the probabilities enter the theory, if they are not to be added directly through either a probabilistic reading of the state, or through additional hidden variables, or through stochastic dynamics?¹¹

However, once again it will be useful not to consider the conceptual question of how to *make sense* of a state space that seems sometimes to be a space of physical states and sometimes a space of probability distributions, but to ask, more practically, how it is that we sometimes *get away* with treating the state space as a space of probability distributions, given that fundamentally that interpretation does not seem consistent.

3 The role of decoherence

Let us consider more carefully just when the probabilistic interpretation can be consistently applied. Recall that it is *interference* that blocks a probabilistic interpretation at the microscopic scale; if a probabilistic interpretation can be applied at any scale, then, it must be because interference phenomena can be neglected.

This, of course, is precisely the question which decoherence theory is designed to answer. In the following, I apply the environment-induced decoherence framework¹² used by Zeh (1993), Zurek (1991) *et al*; essentially the same conclusions, though, can be derived in the decoherent-histories framework (See, e. g. ,Gell-Mann and Hartle (1993), Griffiths (1993), Halliwell (1998); see also my discussion in chapter 3 of Wallace (2011)).

Firstly: suppose that the physical system we are considering has a Hilbert space \mathcal{H} which can be decomposed as

$$\mathcal{H} = \mathcal{H}_M \otimes \mathcal{H}_E, \tag{3}$$

¹¹For detailed discussion of the Everett interpretation, see Saunders *et al*(2010) or Wallace (2011).

¹²See Joos *et al* (2003) and references therein for a review of this framework.

where \mathcal{H}_M is the Hilbert space of the degrees of freedom that interest us — say, the centre-of-mass degrees of freedom of a macroscopic object, or the low-momentum vibrational modes of an extended solid — and \mathcal{H}_E is the Hilbert space of all the other degrees of freedom that interact with those that interest us. \mathcal{H}_M is normally said to represent the “system”, and \mathcal{H}_E the “environment”, degrees of freedom, but it’s important to recognise that the “environment” need not be spatially external: for the extended solid, for instance, \mathcal{H}_E might be the Hilbert space of vibrational modes above some momentum cutoff.

Suppose also that there is some basis $\{|z\rangle\}$, of \mathcal{H}_M , labelled by some (discrete or continuous) variable z , with the following property: the state

$$\left(\int dz \alpha(z) |z\rangle \right) \otimes |\psi\rangle, \quad (4)$$

(where the integral sign schematically denotes summation or integration over z , as appropriate) evolves rapidly, for fairly generic $|\psi\rangle \in \mathcal{H}_E$, to

$$\int dz \alpha(z) |z\rangle \otimes |\psi(z)\rangle, \quad (5)$$

where (i) ‘rapidly’ means ‘rapidly as compared to the dynamical timescales of the system’s own dynamics, and (ii) $\langle\psi(z)|\psi(z')\rangle \simeq 0$ unless $z \simeq z'$. Effectively, dynamics like this consists of the environment ‘measuring’ the system state. The classic example has $\{|z\rangle\}$ as a wave-packet basis, reasonably localised in both centre-of-mass position and momentum. Since the system-environment interactions are local, two wave-packets with significantly distinct locations will cause significantly different evolution of the environment. And two wave-packets with significantly different momentum will swiftly evolve into two wave-packets with significantly different position. (Handling this case carefully requires us to deal with the fact that wavepackets form an overcomplete basis; I’ll ignore this issue for simplicity.)

If this occurs, we will say that the system is *decohered* by the environment, with respect to the $\{|z\rangle\}$ basis.

Now suppose we want to apply the probability interpretation with respect to the $\{|z\rangle\}$ basis: that is, we want to treat a state like (5) as having probability¹³ $|\alpha(z)|^2$ of having whatever physical property (call it Z) that z is supposed to represent (some given centre-of-mass position and/or momentum, say). If the density operator of the system at time t is $\rho(t)$, then we have

$$\Pr(Z = z, t) = \langle z | \rho(t) | z \rangle \quad (6)$$

where $\Pr(Z = z, t)$ is the probability at time t of the system having $Z = z$.

For this really to be interpretable as a *probability*, though, its evolution needs to be free from interference effects: that is, it needs to evolve *like a probability function*. So: let $\Pr(Z = z', t'; Z = z, t)$ be the probability¹⁴ that the system

¹³Or probability density, if z is a continuous variable.

¹⁴Again, this should be interpreted as a probability density if z is continuous.

has $Z = z'$ at time t' , *given* that it has $Z = z$ at time t . Then the standard laws of probability tell us that

$$\Pr(Z = z', t') = \int dz \Pr(Z = z', t' | Z = z, t) \times \Pr(Z = z, t). \quad (7)$$

The dynamics will have this form iff $\Pr(Z = z', t')$ is a linear functional of $\Pr(Z = z, t)$; that is, if $\langle z | \rho(t') | z \rangle$ (considered as a function of z) is a linear functional of $\langle z | \rho(t) | z \rangle$.

In general, this will not be the case. Assuming we can write down autonomous dynamics for ρ in the first place, then of course the overall linearity of unitary dynamics implies that $\rho(t')$ is a linear functional of $\rho(t)$. But in general, $\langle z | \rho(t') | z \rangle$ depends not only on diagonal terms of $\rho(t)$ like $\langle z | \rho(t) | z \rangle$, but on off-diagonal terms like $\langle z | \rho(t) | w \rangle$.

However, in the case we are considering, the interaction with the environment guarantees that off-diagonal terms in ρ are suppressed, and suppressed on timescales much quicker than those which characterise the dynamics of the system itself. In *this* situation, then, the effective evolution equation for ρ reduces to an evolution equation specifically for the diagonal elements of ρ in the $\{|z\rangle\}$ basis.

Let's sum up. If the total system we are studying can be decomposed into "system" and "environment" degrees of freedom, such that for some basis $\{|z\rangle\}$ of the system, the system is decohered by the basis with respect to that basis, then we can consistently treat the total system as a probabilistic mixture of states like $|z\rangle \otimes |\psi(z)\rangle$.

Furthermore, we have excellent reason — from general physical arguments, from specific models, and increasingly from experiment — to think that the macroscopic degrees of freedom of a system are decohered by the residual degrees of freedom with respect to a wave-packet basis for those macroscopic degrees of freedom. So macroscopic systems can consistently be treated as probabilistic mixture of states with different — but definite — values of those macroscopic degrees of freedom.

Note that nothing in this analysis relies on the environment being in any way discarded, except in the pragmatic sense that we're not terribly interested in what it's doing. The total system continues to evolve unitarily, but by virtue of the particular form of that unitary dynamics, it can consistently be given a probabilistic interpretation with respect to its macroscopic degrees of freedom.

Note also that it is the *dynamical* aspects of decoherence that are important here (a point also stressed by Zurek(1994, 1998)). The rapid suppression of the off-diagonal elements of the density operator, which is usually taken as the signature of decoherence, is significant not in itself but because it implies that the dynamics of the *diagonal* elements must be probabilistically interpretable. (The decoherent-histories framework makes this more explicit: in that framework, the possibility of interpreting the system's evolution probabilistically is definitional of decoherence.)

4 Decoherence and the measurement problem

Decoherence, then, explains why the measurement problem is a philosophical rather than a practical problem. Given the ubiquity of decoherence, the strategy of applying the probability interpretation of the state to decohered systems (and to the basis with respect to which the decoherence occurs) will not actually give us contradictory predictions — at least, not to the levels of accuracy which we have any hope of probing empirically. Those readers inclined towards the so-called “shut up and calculate interpretation”, therefore, can stop reading now. For the rest, we should now return to the previous taxonomy of strategies for solving, or dissolving, the problem, and see how they fare in the light of decoherence.

We begin with the strategies for *solving* the problem: the dynamical-collapse and hidden-variable strategies. Recall that dynamical-collapse strategies effectively turn the shift from a physical to a probability reading of the quantum state into a dynamical process, while hidden-variable strategies effectively hold on to a physical reading of the quantum state and add additional dynamical variables over which the probabilities are defined.

At first sight, both strategies are made straightforward by decoherence. We can fairly straightforwardly specify a dynamical collapse theory by stipulating that superpositions of states in the decoherence-preferred basis spontaneously collapse, with collapse probabilities given by the probability rule; we can fairly straightforwardly specify a hidden-variable theory by treating the variable that labels the decoherence-preferred basis as a hidden variable and stipulating that the probability distribution over that variable is given by the probability rule. Neither of these recipes completely specifies the theory in question (how quickly does collapse occur? what dynamics govern the hidden variables, and does it ensure that the probabilities remain in line with the quantum predictions after the initial stipulation?) but — it might seem — the hard work has been done.

Unfortunately, things are not so simple, for a straightforward reason: *decoherence is not a precisely defined process*. We can see this in several related ways:

1. Decoherence suppresses interference only approximately. Interference between alternative “possibilities” is far too low to be experimentally detectable — which is another way to say that to within the limits of experimental accuracy, the probability interpretation gives consistent predictions — but it is not zero.
2. The basis picked out by decoherence is itself only approximately given. In the standard examples, for instance, we might specify it as “a wave-packet, not too localised in position or momentum”. But “not too localised” is scarcely a precise criterion; nor is there anything particularly privileged about the Gaussians usually used for wave-packets, save their mathematical convenience. Criteria like Zurek’s “predictability sieve” (Zurek 1994) can be used to pick a particular choice of basis, but these have more a pragmatic than a fundamental character.

3. The analysis of decoherence I gave above relies on a certain decomposition of the total Hilbert space into “system” and “environment” degrees of freedom. Varying that decomposition will vary the probability rules defined by the decoherence process. Granted, any remotely sensible decomposition can be expected to give essentially similar results, but “remotely sensible” and “essentially similar” are not really the kind of terms we’d like to see in the specification of a fundamental theory. But Dowker and Kent (1996) have provided strong arguments (in the rather different mathematical framework of decoherent histories) that the *mere* criterion that the probability calculus applies to a given decoherence basis, divorced from other considerations, is overwhelmingly too weak to pick out a unique basis, even approximately.

These observations all stem from the same feature of decoherence: that it is a high-level, dynamical process, dependent on details of the dynamics of the world and even on contingent features of a given region of the world. It is not the kind of thing that can be captured in the mathematical language of microscopic physics.

Note that this is not to say that decoherence is not a real, objective process. Science in general, and physics in particular, is absolutely replete with real, objective processes that occur because of high-level dynamical effects. Pretty much any regularity in chemistry, biology, psychology, . . . has to have this character, in fact. But *emergent* processes like this don’t have a place in the axioms of fundamental physics, precisely because they emerge from those axioms themselves.¹⁵

An analogy: it has been well understood for many years that particles — whether in relativistic quantum field theory, or in condensed matter physics — are themselves not fundamental. The particle spectrum of a quantum field theory is determined by the dynamical features of the theory, and in general it is determined only approximately, and in a way that varies according to the contingent features of the regime which interests us. (In some situations, we analyse quantum chromodynamics in terms of quarks; in others, in terms of protons and neutrons; not only the masses and charges of the particles, but which particles we use in the first place, vary according to the energy levels at which the theory is analysed.)

Adrian Kent puts it rather well (in a slightly different context):

It’s certainly true that phase information loss is a dynamical process which needs no axiomatic formulation. However, this is irrelevant to our very simple point: no preferred basis can arise, from the dynamics or from anything else, unless some basis selection rule is given. Of course, [one] can attempt to frame such a rule in terms of a dynamical quantity - for example, some measure of phase information loss. But an explicit, precise rule is needed. (Kent 1990)

¹⁵Emergence has a long and tangled history in philosophy of science; see, e.g., Butterfield (2011a, 2011b), and references therein, for details. In physics, perhaps the most influential discussion of recent years has been Anderson (1972).

But if decoherence cannot be used to *define* a set of hidden variables, or a collapse law, nonetheless it serves to *constrain* both concepts. For if the collapse rule does not cause collapse onto a basis that is approximately decohered, it will fail to reproduce the experimental predictions of quantum mechanics; likewise, if a probability function over the values of the hidden variables does not determine a probability function over the decoherence-selected basis, the hidden variables will not allow us to recover the empirical predictions of quantum theory.

In the case of *nonrelativistic* quantum theory, this is unproblematic. The decoherence-preferred basis is basically a coarse-graining of the position basis, so a collapse rule which collapses the system onto wavepackets fairly concentrated around a particular centre-of-mass position, or a choice of position as the hidden variable, will do nicely. And indeed, we find that the main examples of non-relativistic collapse and hidden-variable theories — the GRW theory and Bohm’s theory — do indeed select position in this way.

It is crucial to note what makes this possible. Position has a dual role in non-relativistic quantum theory: it is at one and the same time (a) one of the fundamental microphysical variables in terms of which the theory is defined, and (b) such that a coarse-grained version of it is preferred by the high-level, dynamical, emergent process of decoherence. As such, it is possible to formulate modifications or supplements to non-relativistic quantum theory that are both precisely defined in terms of the *microphysical* variables used to formulate quantum mechanics, and appropriately aligned with the *macrophysical* variables picked out by decoherence.¹⁶

Unhappily for modificatory strategies, there does not appear to be a variable in extant relativistic quantum theory — in QED, say, or in the Standard Model — that manages to play this dual role. In the fermionic sector, decoherence seems to prefer states of definite particle number — and, as we have already seen, “particle number” is itself a dynamically emergent concept, determined by the complex details of renormalisation theory and dependent, to a considerable degree, on the energy levels at which we wish to study the systems of interest to us. In the bosonic sector — at least where electromagnetism is concerned — decoherence seems instead to select out coherent states (Anglin and Zurek 1996), but these are coherent states defined with respect to the effective field operators governing the system at low energies. In neither case is there any remotely straightforward definition of the decoherence-preferred basis in terms of the quantities used to formulate the quantum field theory at the microphysical level. Indeed, it is a commonplace of renormalisation theory¹⁷ that these variables are largely hidden from view at the level of observable phenomena.

For this reason, I suspect — if it is accepted that modifications of quantum theory should not themselves be stated in a way which makes essential

¹⁶Incidentally, this is what makes testing dynamical collapse theories so difficult: the fact that they must succeed in reproducing the empirical predictions of quantum theory in normal circumstances pretty much guarantees that it will be very difficult to distinguish genuine collapse from mere decoherence. (Difficult, but not impossible; see Stamp, this volume.)

¹⁷See Peskin and Schroeder (1995), or any other modern textbook on quantum field theory, for technical details and references.

reference to dynamically emergent and high-level features of the theory — that the prospects of solving the measurement problem in the relativistic domain by modifying quantum theory are dim. It is notable that, to my knowledge, there is no dynamical-collapse theory even purporting to be applicable to relativistic quantum theory in the presence of interactions,¹⁸ and those hidden-variable theories that have been proposed in the relativistic domain¹⁹ are largely silent about renormalisation.

Turning to the strategies for *dissolving* the measurement problem, recall that there are again two: treat the wavefunction as always probabilistic, or treat it as always physical. The former strategy makes essentially no contact with decoherence: the point of decoherence (as I have presented it here) is to give an account of when a probabilistic reading of the wavefunction is consistent, but the probabilistic strategy treats it as *always* consistent. It does so, in general, by retreating from any attempt to interpret the probabilities as probabilities of anything except measurement outcomes.

This strategy, in effect, is a retreat to the idea that measurement is a primitive. Insofar as that makes sense, it suffices to resolve the puzzles of interference without any concern about decoherence (as it would have to: after all, it has to explain why the probabilistic reading is *always* possible, even in situations where decoherence is negligible). But, as I have alluded to earlier, it does *not* make sense, so far as I can see, partly on philosophical grounds but largely on the straightforward grounds that experimental physicists, and theorists who study experiment, cannot treat measurement as primitive but invariably fall back on the need to analyse it, using quantum theory itself.

To be fair, this objection has received at least some attention from advocates of the strategy (for recent examples, see Peres (1993) and Fuchs (2002)).²⁰ I leave it to readers to judge for themselves whether these responses really do justice to physical practice. It is worth pointing out, though, that in general its advocates tend to work in quantum information, a field whose *raison d'être* is to abstract away the messy details of quantum-mechanical processes and look at their abstract structure. This strategy has been remarkably successful, yielding deep insights about quantum theory that would not have come easily if we had kept the messy details in play; for all that, it's possible to worry that some quantum-information approaches to the measurement problem mistake the map for the territory.

The final strategy is Everett's: treat the quantum state as always giving the physical state of the system. If decoherence can contribute directly to any (dis)solution to the measurement problem, it is here. For we have seen that

¹⁸The nearest thing to such a theory is Tumulka's theory (Tumulka 2006), which is explicitly formulated for a multi-particle theory, on the assumption that there are no interactions. It was, of course, precisely the need to incorporate interactions that drove the pioneers of relativistic quantum mechanics to field theory.

¹⁹The main examples are Durr *et al* (2004, 2005), who take the hidden variables to be particle positions; Struyve and Westman (2007), who take them to be bosonic field configurations, and Colin and Struyve (2007), who take them to be local fermion-number densities.

²⁰And, as I mentioned earlier, it is not *universally* accepted that only a primitivist view of measurement can justify the strategy.

decoherence is an emergent process; what it tells us, interpreted as Everett suggests, is that even if the Universe is *fundamentally* a unitarily evolving whole, at the *emergent* level it has the structure of a probability distribution over states each of which describes approximately classical goings on. There is no mechanism by which one of those states is preferred (is ‘actual’, if you like) and the others are mere possibilities: at the fundamental level, all are components in a single, unitarily-evolving state, and no one is preferred over another. (Any such mechanism would amount to a dynamical-collapse theory, as discussed previously.) But at the emergent level, each term evolves independently of the others; furthermore, their mod-squared amplitudes behave as if they were probabilities.

Does this mean that decoherence suffices to solve the measurement problem, provided that we understand the quantum state as a physical state as Everett proposed? The answer turns on two problems:

1. **(The ontological problem):** For something to *be* a collection of quasi-classical worlds, does it suffice for it to have the *structure* of a collection of quasi-classical worlds — or is more needed?
2. **(The probability problem):** For something to *be* a probability measure over a set of quasi-classical worlds, does it suffice for it to have the *structure* of a probability measure over a collection of quasi-classical worlds — or is more needed?

Space does not permit an extensive engagement with these questions. My own view²¹ is that neither problem is truly problematic. But it is notable that both problems are essentially philosophical in nature: in the light of decoherence, if an Everettian solution to the measurement problem is to be rejected then it will have to be for subtle philosophical reasons rather than any structural deficiency in quantum theory.

5 Conclusion

In twenty-first-century physics, the “measurement problem” is best understood, not as an illegitimate intrusion of a primitive “measurement” postulate into physics, but as a conceptual incoherence in our interpretation of quantum states: it seems impossible to understand the macroscopic predictions of quantum mechanics without interpreting the state probabilistically, yet because of interference, quantum states cannot systematically be thought of as probability distributions over physical states of affairs. We can attempt to resolve that incoherence either by philosophical methods (thinking hard about how to understand quantum states so as to come up with a non-incoherent way) or by modifying the physics (replacing quantum mechanics with some new theory that doesn’t even *prima facie* lead to the conceptual incoherence).

²¹Developed in full in Wallace (2011); for earlier versions, see Wallace (2003, 2010a) in the first case, and Wallace (2007, 2010b) in the second. For further discussion — on both sides — see Saunders *et al* (2010) and references therein.

Decoherence explains why it is that quantum theory nonetheless works in practice: it explains why interference does not, in practice spoil the probabilistic interpretation at the macro level. But because decoherence is an emergent, high-level, approximately-defined, dynamical process, there is no hope of incorporating it into any modification of quantum theory at the fundamental level. Decoherence does, however, act as a significant constraint on such modifications—a constraint which, in the case of relativistic quantum field theory, is likely to be exceedingly hard to satisfy.

Decoherence could, however, play a role in a solution to the measurement problem which leaves the equations of quantum theory alone and treats the objective macroscopic reality we see around us as itself an emergent phenomenon. Such a strategy is committed to the claim that, at the fundamental level, the quantum state continues to describe the physical state of the world: it is, therefore, ultimately Everett's strategy. Decoherence finds its natural role in the measurement problem as the process which explains why quantum mechanics, interpreted as Everett advocates, can be fundamentally deterministic and non-classical, but emergently classical. It does not, however, in any way blunt the metaphysically shocking aspect of Everett's proposal: no one quasi-classical branch is singled out as real; all are equally part of the underlying quantum reality.

Acknowledgements

This paper has drawn on many discussions with Simon Saunders. An earlier version was presented at the fourteenth Seven Pines symposium, generously supported by Lee Gohlike.

References

- Anandan, J. and H. R. Brown (1995). On the reality of space-time geometry and the wavefunction. *Foundations of Physics* 25, 349–360.
- Anderson, P. W. (1972). More is different: Broken symmetry and the nature of the hierarchical structure of science. *Science* 177, 393.
- Anglin, J. R. and W. H. Zurek (1996). Decoherence of quantum fields: Pointer states and predictability. *Physical Review D* 53, 7327–7335.
- Ballentine, L. E. (1990). *Quantum Mechanics*. Englewood Cliffs: Prentice Hall.
- Bassi, A. and G. Ghirardi (2003). Dynamical reduction models. *Physics Reports* 379, 257.
- Bell, J. and M. Nauenberg (1966). The moral aspect of quantum mechanics. In A. De Shalit, H. Feshbach, and L. Van Hove (Eds.), *Preludes in Theoretical Physics*, pp. 279–86. Amsterdam.: North Holland. Reprinted in Bell, J.S.,

- Speakable and Unspeakable in Quantum Mechanics* (Cambridge University Press, 1987), pp. 22–28.
- 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.
- Butterfield, J. (2011a). Emergence, reduction and supervenience: a varied landscape. Forthcoming in *Foundations of Physics*; available online at <http://philsci-archive.pitt.edu/5549/>.
- Butterfield, J. (2011b). Less is different: Emergence and reduction reconciled. Forthcoming in *Foundations of Physics*; available online at <http://philsci-archive.pitt.edu/8355/>.
- Cohen-Tannoudji, C., B. Diu, and F. Laloë (1977). *Quantum Mechanics*, Volume 1. New York: Wiley-Interscience.
- 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.
- DeWitt, B. and N. Graham (Eds.) (1973). *The many-worlds interpretation of quantum mechanics*. Princeton: Princeton University Press.
- Dieks, D. and P. E. Vermaas (Eds.) (1998). *The Modal Interpretation of Quantum Mechanics*, Dordrecht. Kluwer Academic Publishers.
- 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.
- 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, R. Tumulka, and N. Zanghi (2005). Bell-type quantum field theories. *Journal of Physics A* 38, R1.
- Everett, H. I. (1957). Relative State Formulation of Quantum Mechanics. *Review of Modern Physics* 29, 454–462. Reprinted in DeWitt and Graham (1973).
- 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 (2000). Quantum theory needs no “interpretation”. *Physics Today* 53(3), 70–71.

- 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.
- Griffiths, R. (1993). Consistent interpretation of quantum mechanics using quantum trajectories. *Physical Review Letters* 70, 2201–2204.
- Halliwel, J. J. (1998). Decoherent histories and hydrodynamic equations. *Physical Review D* 35, 105015.
- Harrigan, N. and R. W. Spekkens (2010). Einstein, incompleteness, and the epistemic view of states. *Foundations of Physics* 40, 125.
- 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 Appearance of a Classical World in Quantum Theory* (2nd Edition ed.). Berlin: Springer.
- Kent, A. (1990). Against Many-Worlds Interpretations. *International Journal of Theoretical Physics* A5, 1764. Expanded and updated version at <http://www.arxiv.org/abs/gr-qc/9703089>.
- Kochen, S. and E. Specker (1967). The problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics* 17, 59–87.
- Landau, L. and E. Lifshitz (1977). *Quantum Mechanics (Non-Relativistic Theory)* (3rd ed.). Butterworth Heinemann.
- Misra, B. and E. C. G. Sudarshan (1977). The Zeno’s paradox in quantum theory. *Journal of Mathematical Physics* 18, 756.
- Penrose, R. (1989). *The Emperor’s New Mind: concerning computers, brains and the laws of physics*. Oxford: Oxford University Press.
- 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.
- Rae, A. I. (1992). *Quantum Mechanics* (3rd ed.). Bristol: Institute of Physics Publishing.
- Redhead, M. (1987). *Incompleteness, Nonlocality and Realism: A Prolegomenon to the Philosophy of Quantum Mechanics*. Oxford: Oxford University Press.
- Sakurai, J. J. (1994). *Modern Quantum Mechanics*. Reading, MA: Addison-Wesley.

- Saunders, S., J. Barrett, A. Kent, and D. Wallace (Eds.) (2010). *Many Worlds? Everett, Quantum Theory, and Reality*, Oxford. Oxford University Press.
- 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 (2007). A minimalist pilot-wave model for quantum electrodynamics. *Proceedings of the Royal Society of London A* 463, 3115–3129.
- 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>.
- 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. (2007). Quantum probability from subjective likelihood: Improving on Deutsch’s proof of the probability rule. *Studies in the History and Philosophy of Modern Physics* 38, 311–332.
- 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.
- Wallace, D. (2010a). 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. (2010b). How to prove the Born rule. In S. Saunders, J. Barrett, A. Kent, and D. Wallace (Eds.), *Many Worlds? Everett, Quantum Theory and Reality*. Oxford: Oxford University Press. Available online at <http://arxiv.org/abs/0906.2718>.
- Wallace, D. (2011). *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford: Oxford University Press.
- Zeh, H. D. (1993). There are no quantum jumps, nor are there particles! *Physics Letters A* 172, 189.
- Zurek, W. (1998). Decoherence, chaos, quantum-classical correspondence, and the algorithmic arrow of time. *Physica Scripta* T76, 186–198.
- 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. (1994). Preferred observables, predictability, classicality, and the environment-induced decoherence. In J. J. Halliwell, J. Perez-Mercader, and W. H. Zurek (Eds.), *The Physical Origins of Time Asymmetry*, pp. 175–212. Cambridge: Cambridge University Press.