What is orthodox quantum mechanics?

David Wallace

April 20, 2016

Abstract

What is called "orthodox" quantum mechanics, as presented in standard foundational discussions, relies on two substantive assumptions — the projection postulate and the eigenvalue-eigenvector link — that do not in fact play any part in practical applications of quantum mechanics. I argue for this conclusion on a number of grounds, but primarily on the grounds that the projection postulate fails correctly to account for repeated, continuous and unsharp measurements (all of which are standard in contemporary physics) and that the eigenvalue-eigenvector link implies that virtually all interesting properties are maximally indefinite pretty much always. I present an alternative way of conceptualising quantum mechanics that does a better job of representing quantum mechanics as it is actually used, and in particular that eliminates use of either the projection postulate or the eigenvalue-eigenvector link, and I reformulate the measurement problem within this new presentation of orthodoxy.

1 Introduction: the orthodox view of orthodoxy

"Orthodox" or "standard" quantum mechanics, as typically presented in text-book philosophy-of-physics discussions[1], consists of these components:

The structural core: This has three parts:

1. **States:** The possible states of a quantum system are represented by normalised vectors in some complex Hilbert space.

2. **Observables:** To any physical quantity used to describe the system (often called an 'observable') is associated a self-adjoint operator on that same Hilbert space.

3. **Dynamics:** The state of a quantum system evolves over time according to the *Schrödinger equation*:

$$\frac{\mathrm{d}}{\mathrm{d}t} \left| \psi(t) \right\rangle = -\frac{i}{\hbar} \widehat{H} \left| \psi(t) \right\rangle \tag{1}$$

where $\widehat{H}$ is the self-adjoint operator corresponding to the system's energy.

---

[1]See, e. g. , Albert (1992), Barrett (1999), Bub (1997), Penrose (1989, ch.5–6)

(This can be generalised in certain respects, in particular by allowing quantum states to be mixed rather than pure, and in fact I think this generalisation does a better job of capturing real-world quantum mechanics than the pure-state version (cf Wallace 2013) but for simplicity I use the pure-state version in this paper.)

The Born (probability) rule : Suppose some quantity O has associated operator $\widehat{O}$, which can be written

$$\widehat{O} = \sum_i o_i \widehat{\Pi}(i) \tag{2}$$

where the o_i are the distinct eigenvalues of the operator and $\widehat{\Pi}(i)$ projects onto the subspace of states with eigenvalue o_i. (Recall that any self-adjoint operator can be so written — this is the 'spectral resolution' of the operator.) Then if O is measured on a quantum system with state $|\psi\rangle$, then:

1. The only possible outcomes of the measurement are the eigenvalues o_i of the operator;

2. The probability of the measurement giving result o_i is

$$\Pr(O = o_i) = \langle\psi|\widehat{P}(i)|\psi\rangle. \tag{3}$$

The projection postulate (aka the collapse law): Suppose some quantity O, as above, is measured on a quantum system in state $|\psi\rangle$. Then the measurement induces a stochastic transition on the state, so that:

1. Immediately after the measurement, the system is in one of the states

$$|\psi_i\rangle = \frac{\widehat{\Pi}(i)|\psi\rangle}{\|\widehat{\Pi}(i)|\psi\rangle\|}. \tag{4}$$

2. The probability that the system transitions into state $|\psi_i\rangle$ is given by
$$\Pr(|\psi\rangle \to |\psi_i\rangle) = \langle\psi|\widehat{P}(i)|\psi\rangle. \tag{5}$$

(The projection law thus restricts the generality of the Schrödinger equation: systems evolve under it *only when a measurement is not taking place*.)

The eigenvector-eigenvalue link (E-E link): Given an quantity O as above:

1. A system in state $|\psi\rangle$ possesses a definite value of O if and only if $|\psi\rangle$ is an eigenstate of $\widehat{O}$, $\widehat{O}|\psi\rangle = o_i|\psi\rangle$.

2. In this case, the definite value is the associated eigenvalue o_i.

Given one additional assumption — that if a measurement of O returns value o_i, the measured system actually has value o_i of O — the Born rule can be derived from the projection postulate and the eigenvalue-eigenvector link. For if O is measured on a system in state $|\psi\rangle$, by the projection postulate it will transition into an eigenstate of $\widehat{O}$, with the probability of transitioning given by (5); after the collapse, it will have a definite value of O by the eigenvalue-eigenvector link; if measurement simply reports that definite value, the Born rule follows.

In any case, it is standard in foundations of quantum mechanics to treat both the projection postulate and the E-E link as core components of orthodox QM. Interpretations of QM like Everett's and Bohm's, for instance, are specifically described as 'no-collapse' interpretations in view of the fact that they drop the collapse law from the postulates of QM; discussions of the ontology of the GRW collapse theory Albert and Loewer (1996) talk of the need to abandon the E-E link; attempts at interpretation-neutral discussions of the ontology of QM (e.g., Skow 2010; Darby 2010; Bokulich 2014; Wolff 2015; Wilson 2016) typically take the E-E link as a starting point.

Furthermore, typical statements of the quantum measurement problem typically take the E-E link, and/or the projection postulate, as central. The measurement problem is the problem of *macroscopic indefiniteness*, of quantum states that describe macroscopic systems in states that are indefinite with regard to ordinary properties such as the location of pointers or the heartbeats of cats. Or — if macroscopic indefiniteness is to be removed via the projection postulate — it is the problem of *dynamical ill-definedness*, of the lack of any well-defined recipe as to when collapse occurs (it is easy to show that collapse cannot be a consequence or special case of the Schrödinger equation applied to a complex measuring system).

The purpose of this paper, by contrast, is to argue that orthodox quantum mechanics in fact consists only of the structural core and the Born rule. The projection postulate, and the eigenvector-eigenvalue link, are at best parts of a proposed interpretation of QM that goes beyond orthodoxy, at worst unmotivated distractions. As such, in formulating (as opposed to solving) the quantum measurement problem, we should begin with just the structural core and the Born rule. We might introduce one or both as part of a *solution* to the measurement problem, but we confuse the dialectic by taking them as initial common ground.

To be clear what I mean: I will not (here) argue that the *best* or *right* way to interpret QM, or to solve the quantum measurement problem, involves abandoning collapse and/or the E-E link. I will argue that QM as actually practiced by physicists — and what does "orthodox" QM mean, if not that? — already proceeds without either.

In section 2 I provide some evidence that physicists in practice do not seem to make use of a collapse rule. I strengthen this in sections 3–4 by arguing that the collapse rule is incapable of handling two standard kinds of experimental setup: those involving repeated measurements, and those involving continuous observation. In section 6 I point out that the rapid spreading of wavepackets

3

under the Schrödinger equation means that the E-E link makes the ridiculous claim that essentially any system, including macroscopic systems, is maximally indefinite in position, and hence that the E-E link does not have the resources to say when systems are actually localised; in section 7 I deploy a result of Hegerfeldt to show that this generalises to pretty much any observable.

For the remainder of the paper, I explore what 'orthodox quantum mechanics' is, shorn of the E-E link and the projection postulate. (This part of the paper draws on some parts of Wallace (2016), albeit deployed in a rather different context). I consider (in section 8) a view which treats preparation and measurement as primitive, but ultimately reject it (in section 9) on the grounds both that it too struggles with continuous and repeated measurements, and that it cannot handle applications of QM where results from QM are integrated into larger pieces of historical science. With this as a starting point, I finally suggest (section 10) that orthodoxy should be understood as an inchoate attitude to the quantum state, where its dynamics are always unitary but where it is interpreted either as physically representational or as probabilistic, according to context. In the concluding section I reflect on the right formulation of the measurement problem given this conception of what 'orthodox QM' actually is.

2   Against collapse: indirect evidence

The projection postulate appears in Dirac (1930) and von Neumann (1955), the first two codifications of the axioms of QM. It continues to be widespread, though not universal, in first courses on QM to this day: an unscientific perusal of my shelf reveals that collapse is included in about half of the books there that present QM from scratch. (The Born rule, of course, appears in all of them.) But for all this (I will argue) it plays no real role in applications of quantum mechanics in physics. It is rather hard to prove a negative, but here I give some suggestive reasons to think that physical practice abjures collapse.

Firstly, collapse is conspicuously absent from *second* courses in QM, and in particular in courses on relativistic QM. This ought to strike a student as peculiar (it certainly struck the author, as an undergraduate, as peculiar): collapse, as formally defined in QM, is a *global* phenomenon, applying to the whole quantum state and so affecting, simultaneously, systems spatially far from one another. In relativity, this notion of simultaneity is frame-dependent (or simply meaningless, depending how you think about conventionality of simultaneity, but in any case problematic). One would expect, if collapse is really part of orthodox QM, that the first chapter of any relativistic QM textbook would start with a careful discussion of exactly how the collapse postulate is to be applied in the relativistic context. I have not once seen any such textbook so much as consider the question.

Again, to be clear: the point is not that collapse is *unsatisfactory* in the relativistic regime. Of course it is; the tension between relativity and QM has been known at least since the EPR paper (Einstein, Podolsky, and Rosen 1935). But relativistic QM textbooks contain, not an unsatisfactory collapse rule, but

no collapse rule at all. One concludes that the theory must be applicable without any mention of collapse. And indeed it is: the name of the game in relativistic QM is to calculate probablility distributions over physical quantities — most often, over the various energies, momenta and particle numbers of the decay products of some scattering experiment — and for this, only the Born rule is required; collapse plays no part.

Secondly, the theoretical physics community has been worrying for forty years about the so-called "black hole information loss paradox" originally identified by Hawking (1976). (See, e. g. , Page (1994, Belot, Earman, and Ruetsche (1999) and references therein, though the debate continues in lively fashion to this day.) At its heart, the paradox is simply that black hole decay is *non-unitary* and as such can't be described within the Schrödinger-equation framework. But state-vector collapse is also non-unitary! So if the collapse law is part of orthodox QM, quantum-mechanical dynamics were never unitary in the first place: they were an alternating series of unitary and non-unitary processes. So why be so desparate to preserve unitary in the exotic regime of black hole decay, when it is ubiquitous in far more mundane cases? One has the clear impression that (at least this part of) the theoretical physics community does *not* in fact think that dynamics is non-unitary in any other contexts in physics, rendering black hole decay uniquely problematic. Tempting though it might be for this advocate of the Everett interpretation to claim that the community has adopted the many-worlds theory *en masse*, a more mundane account is simply that (what they regard as) *orthodox* QM does not include the collapse postulate.[2]

Thirdly, modern quantum field theory largely abandons Hamiltonian methods in favour of the path-integral approach. But in that approach it is not even clear how collapse is to be defined (and, again, textbook presentations never seem to mention the issue), and yet the theory still seems to produce empirically successful predictions.

Finally, and as an admittedly crude indicator, searching the archives of *Physical Review* for projection postulate, wave-function collapse and the like 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. (For comparison, searches for terms like state vector or Hilbert space or Schrodinger equation typically turn up several tens of thousands of references.)

3   Against collapse: inadequacy for repeated measurements

The case of repeated measurements — when some quantity is measured on a quantum system and then, a short while later, measured again — has actually

---

[2]Of course, plenty of people working on black hole decay *are* fairly explicit advocates of the Everett interpretation, and I have argued elsewhere that quantum cosmology generally is *tacitly* committed to the Everett interpretation, but it's clear that the majority of the community embrace Mermin's "shut up and calculate" approach (Mermin 2004).

been used, since Dirac, as an argument *for* state-vector collapse. The argument goes like this: repeated measurements must give identical results; so if a measurement of O gives outcome o_i, then a subsequent measurement of O immediately afterwards must also give outcome o_i, The only way this is compatible with the Born rule is if the state of the system immediately before this second measurement is an eigenstate of $\widehat{O}$ with eigenvalue o_i — so to get repeated measurements right, wavefunction collapse is a requirement.

... which would be all very well, if repeated measurements *did* give identical results. But:

- photon detectors typically absorb photons: immediately after a measurement on a photon, the photon no longer exists;

- The Stern-Gerlach apparatus detects an atom's spin by slamming it very hard into a screen; this process is in no way guaranteed to preserve that atom's spin.

- More generally, measuring something by slamming it very hard into something else is probably the single most commonly used tool in the experimental physicist's toolbox.

In fact, 'non-disturbing' measurements, in which repeated measurements indeed give the same results, are decidedly uncommon in quantum mechanics and require some skill to set up (see Home and Whitaker (1997) for discussion). So a collapse rule explicitly designed to ensure that repeated measurements give the same results is in flat conflict with a lot of observed physics.

By contrast, quantum mechanics *without* collapse has no trouble with repeated measurements — non-disturbing or otherwise. The familiar trick, following von Neumann's original prescription, is to include the measurement device in the physical analysis. Suppose for simplicity that $\widehat{O}$ is non degenerate,

$$\widehat{O} = \sum_i o_i \left| o_i \right\rangle \left\langle o_i \right|, \tag{6}$$

and suppose that the measurement device has some observable M corresponding to the possible measurement outcomes. In von Neumann's original version, M is the position of the centre of mass of some pointer; here for convenience I take $\widehat{M}$ too as being discrete and nondegenerate,

$$\widehat{M} = \sum_i m_i \left| m_i \right\rangle \left\langle m_i \right|. \tag{7}$$

Then the measurement interaction is assumed to have form

$$\left| o_i \right\rangle \otimes \left| m_0 \right\rangle \rightarrow \left| \varphi_i \right\rangle \otimes \left| m_i \right\rangle. \tag{8}$$

Applying this measurement process to a system initially in state

$$\left| \psi \right\rangle = \sum_i \lambda_i \left| o_i \right\rangle \tag{9}$$

and a measurement device initially in state $|m_0\rangle$ gives the outcome

$$|\psi\rangle \otimes |m_0\rangle \to \sum_i \lambda_i |\varphi_i\rangle \otimes |m_i\rangle. \qquad (10)$$

Applying the Born rule to a measurement *of M* now tells us that the probability of getting m_i is $|\lambda_i|^2$ — exactly what the Born rule requires for a measurement of O on the original system, and it is for exactly this reason that this process indeed qualifies as a measurement.

There is no requirement here that $|\varphi_i\rangle = |o_i\rangle$ or even that the distinct $|\varphi_i\rangle$ are orthogonal — indeed, the measurement process could perfectly well dump the measured system in some fixed post-measurement state $|\varphi_0\rangle$ (as in the case of photon absorption) in which case the measurement process is

$$|\psi\rangle \otimes |m_0\rangle \to |\varphi_0\rangle \otimes \sum_i \lambda_i |m_i\rangle. \qquad (11)$$

The 'non-disturbing' measurements are then the ones where indeed $|\varphi_i\rangle = |o_i\rangle$. In these cases, but only these, if we bring in a second copy of the measurement device and repeat the measurement interaction, we get

$$|\psi\rangle \otimes |m_0\rangle \otimes |m_0\rangle \to \left(\sum_i \lambda_i |o_i\rangle \otimes |m_i\rangle \right) \otimes |m_0\rangle \to \sum_i \lambda_i |o_i\rangle \otimes |m_i\rangle \otimes |m_i\rangle. \qquad (12)$$

Applying the Born rule in this case to a joint measurement of $M \times M$, we find that indeed, when measurements are non-disturbing the probability is 100% that two successive measurements give the same result.

If there is a lesson to learn from repeated measurements it is that the Born rule, by itself, does not define *transition* probabilities, but only probabilities at an instant (an issue I return to in section 10). But "wave functions collapse on measurement" does not solve this problem satisfactorily, and indeed gives flatly incorrect results.

4 Against collapse: inadequacy for continuous measurement

Continuous measurements — where a system is constantly observed to see if, or how quickly, it undergoes some change — are commonplace in physics. For instance, radioactive decay measurements — where a Geiger counter is placed near some radioactive substance, and the rate of decay is recorded — are among the most straightforward demonstrations of quantum mechanics' probabilistic nature. Yet they fit strikingly badly into the wavefunction-collapse framework.

It is not that the *physics* of (for instance) radioactive decay is problematic, at least phenomenologically. (Actually calculating decay rates *ab initio* is another matter: the nucleus is a complex and strongly bound system, and hard to treat

analytically.) The idea is that — if the decay rate is $1/\tau$ — then an undecayed particle, in state $|\text{undecayed}\rangle$, evolves over some short time interval δt like

$$|\text{undecayed}\rangle \to \left(1 - \frac{\delta t^2}{2\tau}\right)|\text{undecayed}\rangle + \sqrt{\frac{1}{\tau}}\delta t\,|\text{decay products}\rangle. \qquad (13)$$

Meanwhile, the decay-product state's own evolution over time, which can be represented as

$$|\text{decay products}\rangle \to \widehat{U}(t)\,|\text{decay products}\rangle \equiv |\text{decay products};t\rangle \qquad (14)$$

explores a very large region of Hilbert space and, in particular, satisfies

$$\langle\text{decay products}|\text{decay products};t\rangle \simeq 0 \qquad (15)$$

for $t_0 < t < T$, where T is the (extremely large) Poincaré recurrence time for the system and $t_0 \ll \tau$ (i.e., the rate at which the radioactive products evolve away from their original state is much quicker than the particle's decay rate). Under these assumptions we can deduce

$$\widehat{U}(t)\,|\text{undecayed}\rangle = \mathrm{e}^{-t/2\tau}\,|\text{undecayed}\rangle + \int_0^t \mathrm{d}\xi\,\frac{\mathrm{e}^{-\xi/2\tau}}{\sqrt{\tau}}\,|\text{decay products};(t-\xi)\rangle$$
$$(16)$$

at least for $T \gg t \gg \tau$.

Applying the Born rule to this system gives exactly the results we would expect: at time t, the probability of the system being undecayed is $|\mathrm{e}^{-t/2\tau}|^2 = \mathrm{e}^{-t/\tau}$. And no assumption of wavefunction collapse is required to derive this probability. But suppose we make that assumption anyway: when, in that case, is the wavefunction supposed to collapse?

One possibility would be to model the continuous process of measurement by a frequent but discrete series — applying the projection postulate every Δ seconds — and then taking $\Delta \to 0$. As long as Δ is long enough — technically speaking: as long as $\Delta \gg t_0$ — this iterated collapse will leave the probabilities unaffected. But it is the content of the quantum Zeno paradox (Misra and Sudarshan 1977) that as $\Delta \to 0$, the evolution of the system is entirely halted: in this limit, the state of the particle remains $|\text{undecayed}\rangle$ forever.

Misra and Sudarshan did assume (at least for the purposes of their paper) that observation required collapse, and so that continuous observation required continuous collapse; hence "paradox". But it is the collapse postulate, not anything about continuous observation *per se*, that delivers this impossible result. Modelling of measurement as a physical process, as per the previous section, reveals (cf Home and Whitaker 1997) that:

- The Zeno 'paradox' is a real (and empirically confirmed) physical effect: if a discrete measurement process is carried out repeatedly (and, crucially, if the time taken to carry out every individual measurement is short compared to the timescales on which the measured system evolves), then the rate of evolution of the system really is reduced by the measurements, and tends to zero as the frequency of repetition tends to infinity.

8

- A continuous observation can also be modelled as a physical process, and in this case the relevant variable is the response speed of the measurement device compared to the timescale on which the measured system evolves. Again, when the former is much faster than the latter, evolution is heavily suppressed. But an observation can be 'continuous' even while its response time is relatively slow: in the case of a Geiger counter, the response time is so slow compared to the evolution timescales of decay that Zeno slowing is negligible. (The relevant system timescale is not the decay rate, but the evolution time of the decay products, i.e. t_0.)

(This is not to say that the Zeno effect is *entirely* non-paradoxical, even when understood without the distorting reference to collapse. Paradoxical (though non-contradictory) consequences arise when the measurement process involves energy exchange between measurement device and system only when the system is in a state distinct from its original state, so that the presence of the measurement device appears to halt the system's evolution even though the two are not interacting. This is related to the phenomenon of interaction-free measurement, as seen in the Elitzur-Vaidman bomb problem Elitzur and Vaidman (1993). For further discussion see Home and Whitaker, *ibid*, or (for an unapologetically pro-Everettian perspective!) Wallace (2012, pp.390–3).)

5 Against collapse: inadequacy for unsharp measurement

The view that measurements are represented by collections of mutually orthogonal projectors is now thirty years out of date. Quantum measurement theory now regards the "projection-valued measurements" (PVMs) that can be so represented merely as a special case of a more general framework: "positive-operator-valued measurements" (POVMs).[3] In the POVM framework, measurements are represented by collections $\{\widehat{M}_i\}$ of self-adjoint operators that (i) are positive (that is, have no negative eigenvalues, or equivalently, satisfy $\langle\psi|\,\widehat{M}_i\,|\psi\rangle \geq 0$ for any state $|\psi\rangle$); (ii) sum to unity, $\sum_i \widehat{M}_i = \widehat{1}$.

For instance, consider measuring a particle's phase-space position: that is, consider simultaneously measuring its position and its momentum. Within the PVM framework, this is impossible: position and momentum do not commute. But in modern measurement theory, this simply tells us that we cannot make a simultaneous *sharp* measurement of position and momentum. We can measure both provided we are prepared to accept a little noise in the measurement process, and for macroscopically large systems the noise can be very small indeed — which is reassuring, since manifestly we *do* simultaneously measure the position and momentum of macroscopic bodies.

A phase-space POVM (in, for simplicity, one spatial dimension) can be defined by starting with some state $|\varphi\rangle$ that is a wavepacket approximately lo-

---

[3]For more detail on the physics of this section, see, e. g., Busch, Lahti, and Mittelstaedt (1996).

calised around position and momentum zero (say, a Gaussian), so that

$$|\varphi(p,q)\rangle = \exp\left(-i\widehat{X}p\right)\exp\left(+i\widehat{P}q\right)|\varphi\rangle \tag{17}$$

is the same state translated so as to be localised around position q and momentum p. Then the family of operators

$$\widehat{M}_{p,q} = \frac{1}{2\pi}|\varphi(p,q)\rangle\langle\varphi(p,q)| \tag{18}$$

is a POVM and can be used to represent the unsharp phase-space measurement. It will give probability distributions over position and momenta separately which are smearings-out of the sharp results obtained from the Born rule, with the level of smearing depending on the width of the wavepacket in position and momentum space and becoming negligible in both cases for macroscopically large systems.

Similarly, the POVM framework can handle fuzzy measurements of a single quantity, as might be appropriate when the measurement device is imperfect. Given an observable O corresponding to an operator $\widehat{O}$ with spectral resolution (2), suppose that $f^1, \ldots f^N$ are N functions from the spectrum of $\widehat{O}$ to the nonnegative reals, satisfying

$$\sum_{k=1}^{N} f^k(o_i) = 1 \tag{19}$$

for all i. Then the family of operators $f_k(\widehat{O})$ is a POVM. If $f_k(o_i) = \delta_i^k$, this just reduces to a sharp measurement of O, but more general measurements of O can be represented by more general choices of the f^k.

The POVM generalisation of traditional measurement theory is by now routine, and mathematically speaking is a straightforward generalisation of the Born rule. But it has no associated collapse law, and so it is opaque how to apply collapse in POVM contexts. In addition, POVMs are not associated with the spectral decompositions of the operators representing physical quantities, so to deduce what POVM is being applied, we need to model the measurement process as a unitary interaction with the measurement device, and then in due course apply the Born rule with respect to a macroscopic quantity pertinent to the measurement device.

The lessons of continuous, repeated and unsharp measurements are the same: in any measurement processes more complicated than a simple, non-repeated discrete measurement, reliably getting the physics right requires treating the system's behaviour unitarily, and if necessary physically modelling the measurement process. Collapse is *at best* an unreliable shorthand. And of course, it is *only* in "measurement processes more complicated than a simple, non-repeated discrete measurement" that the collapse rule could play any role in physical practice anyway. If we measure the system once and immediately discard it, the Born rule is all we need.

I conclude that the collapse postulate plays, and can play, no real part in actual applications of quantum mechanics.

6 Against the eigenvalue-eigenvector link: problems for position

Consider a mass-m point particle — either a fundamental particle, or, more typically, the centre-of-mass degree of freedom of some rigid body like a dust mote or a table. Restricting it, for simplicity, to one dimension, its most significant observables are position and momentum, corresponding to operators $\widehat{X}$ and $\widehat{P}$ respectively, obeying the commutation relation $\left[\widehat{X}, \widehat{P}\right] = i\hbar$.

It is a standard result of quantum mechanics[4] that:

- $\widehat{X}$ and $\widehat{P}$ have continuous spectra (reflecting the fact that these are not quantised quantities, that any position or momentum is a possible result of a measurement), and can be expressed as

$$\widehat{X} = \int_{-\infty}^{+\infty} \mathrm{d}x\, x\, |x\rangle \langle x| \quad \text{and} \quad \widehat{P} = \hbar \int_{-\infty}^{+\infty} \mathrm{d}k\, k\, |k\rangle \langle k| . \tag{20}$$

- Any quantum state can be expressed in the position basis as

$$|\psi\rangle = \int_{-\infty}^{+\infty} \mathrm{d}x\, |x\rangle \langle x|\psi\rangle \equiv \int_{-\infty}^{+\infty} \mathrm{d}x\, \psi(x) |x\rangle \tag{21}$$

  where $\psi(x) \equiv \langle x|\psi\rangle$ is the *position-space wavefunction* (or often just *wavefunction*) of the state.

- Similarly, any quantum state can be expressed in the momentum basis as

$$|\psi\rangle = \int_{-\infty}^{+\infty} \mathrm{d}k\, |k\rangle \langle k|\psi\rangle \equiv \int_{-\infty}^{+\infty} \mathrm{d}k\, \hat{\psi}(k) |k\rangle \tag{22}$$

  where $\psi(k) \equiv \langle k|\psi\rangle$ is the *momentum-space wavefunction* of the state.

- The position and momentum bases are related by

$$|k\rangle = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{+\infty} \mathrm{d}x\, \mathrm{e}^{ikx} |x\rangle \tag{23}$$

  from which it follows that the position and momentum representations are Fourier transforms of one another:

$$\hat{\psi}(k) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{+\infty} \mathrm{d}x\, \mathrm{e}^{-ikx} \psi(x). \tag{24}$$

Suppose we apply the E-E link to the position of the particle. For the particle to *definitely have position* x, it would need to be an eigenstate of $\widehat{X}$ — that

---

[4]And, like most 'standard results of quantum mechanics', there are some tacit additional mathematical assumptions required. See (Ruetsche 2011, ch.3) for details.

is, it would need to be in state $|x\rangle$. That isn't possible: because the spectrum of the position operator is continuous, the eigenstates of position are so-called 'improper eigenstates' — at least as QM is normally used, they do not represent an actually-attainable state of a quantum system. ('Legal' quantum states are normalised — $\langle\psi|\psi\rangle = 1$ — whereas the norm $\langle x|x\rangle$ is infinite, or at any rate undefined.)

So: no system has a *perfectly* definite position. This is not in itself problematic. It is a standard result of functional analysis that functions $f(\widehat{X})$ may be defined by

$$f(\widehat{X}) = \int_{-\infty}^{+\infty} \mathrm{d}x\, f(x)\, |x\rangle\langle x|. \tag{25}$$

In particular, if Σ is some compact (i.e., closed and bounded) subset of the real numbers, and if Λ_Σ is defined by

$$\begin{aligned} \Lambda_\Sigma(x) \quad &= 1 \quad \text{if } x \in \Sigma \\ &= 0 \quad \text{otherwise} \end{aligned} \tag{26}$$

$$\tag{27}$$

— that is, if Λ_Σ represents the property of being in Σ — then

$$\Lambda_\Sigma(\widehat{X}) = \int_\Sigma |x\rangle\langle x| \tag{28}$$

projects, according to the E-E link, onto all and only those states which are definitely located in Σ. In the position representation, this is all and only states whose wavefunction vanishes outside Σ. This suggests that if we want to represent a reasonably-well-localised particle, we should choose one with a wavefunction confined to some reasonably small Σ. (And similarly *mutatis mutandis* if we want to consider systems localised in *momentum*.)

The first thing to say about this is that it is not how physicists *in fact* represent localised particles. The standard strategy in physics is to represent a particle localised at some point x_0 by a Gaussian, i.e. a state with wavefunction

$$\psi(x) = \mathcal{N}\exp(-(x - x_0)^2/2L^2). \tag{29}$$

$\psi(x)$ is very small when $|x - x_0|/L \gg 1$, so a state like this, if its position is measured, is nearly certain to be found within a few multiples of L from x_0 — L is the "effective width" of the state, the size of the region in which it is 'effectively localised' in physics parlance. But $\psi(x) \neq 0$ for every value of x there is — so according to the E-E link, the particle is completely delocalised, no matter how small L might be.

Perhaps not too much should be made of this. Physicists use Gaussians because they are mathematically very convenient, rather than from some deep commitment to what 'true' localisation is like. Perhaps we should think of the Gaussian as just a very convenient approximation to a 'really' localised state, with the latter having a wavefunction with genuinely compact support.

But suppose that a quantum system, at some initial time, *does* have such a wavefunction — say, ψ_0, which is localised inside some compact region Σ. If we represent that same system instead in the momentum-space representation — that is, with its momentum-space wavefunction, which is the Fourier transform $\hat{\psi}$ of ψ — then we find that $\hat{\psi}$ does not itself have compact support. (This is a consequence of the classical Paley-Wiener theorem,[5] which says *inter alia* that the Fourier transform of a compactly supported L^2 function is holomorphic.) Via the $E-E$ link, this tells us that any particle whose *position* is not completely indefinite has a completely indefinite *momentum.*

That might itself be worrying: we might have hoped, given the uncertainty principle, that a particle could definitely have both (a) a position within some region of width L and (b) a momentum within some region of width $\hbar/L$, but we won't get that from the $E-E$ link. But worse is to come: for consider the *time evolution* of this system. We might expect that a particle whose momentum is completely indefinite will spread out instantaneously over all of space — and indeed, this is exactly what happens. Even 'confining' the system inside some potential well will not prevent its spreading out, for a quantum system will 'tunnel' through any potential barrier unless it is infinite, and so unphysical.

So: no body can be localised to any degree at all for more than an instant, at least on the E-E definition of 'localised'. (I have argued for this on intuitive physical grounds but will supply a mathematical proof in the next section). No assumption about the 'microscopic' nature of the body in question has been made: the argument applies as readily to chairs, tables and planets as to electrons or atoms, and so chairs, tables and planets, according to the E-E link, have at almost all times a completely indefinite location. If we assume the projection postulate, of course, a system *is* localised immediately after a position measurement — but the operative word is 'immediately'. An arbitrarily short time after the measurement, delocalisation is complete.

I conclude that the E-E link is of no use in understanding what it really is for a physical system to be localised to any degree.[6] A fortiori, it cannot be being used in physics to do useful work in our understanding of localisation. As we shall see, this is not a feature unique to spatial localisation.

7 Against the eigenvector-eigenvalue link: problems for basically any quantity

Hegerfeldt's theorem (Hegerfeldt 1998a; Hegerfeldt 1998b) is as follows:

> **Hegerfeldt's theorem:** Suppose that the spectrum of the Hamiltonian of some quantum system is bounded below (something that

---

[5] See, e. g. , Rudin (1991, pp.196-202).

[6] The line of argument here has some resemblance to that used by Albert and Loewer (1996) to argue that the E-E link should be rejected in the GRW theory in place of a "fuzzy link". But Albert and Loewer attributed the problem to the Gaussian collapse function used in the GRW theory, whereas as we have seen, the problem arises even in the absence of any collapse event, as a consequence of ordinary Schrödinger dynamics.

holds of essentially any physically reasonable Hamiltonian) and let $|\psi(t)\rangle$ be some dynamical history of that system (i. e., some solution to the Schrödinger equation). Then if $\widehat{\Pi}$ is any projection operator[7], *exactly one* of the following holds:

1. $\langle\psi(t)|\widehat{\Pi}|\psi(t)\rangle \neq 0$ for all times t except for some nowhere dense set of times of measure zero (i. e., for all times except some set of isolated instants);

2. $\langle\psi(t)|\widehat{\Pi}|\psi(t)\rangle = 0$ for all times t.

Hegerfeldt proved the theorem as part of an investigation into localisation in relativistic quantum mechanics[8] but in fact it causes severe difficulties for the E-E link in general. For consider again our operator

$$\widehat{O} = \sum_i o_i \widehat{\Pi}_i \tag{30}$$

(which includes, as a special case, the sort of discretisations of position we considered previously.) The observable O corresponding to $\widehat{O}$ will according to the E-E link, definitely *not* have value o_i with respect to state $|\psi(t)\rangle$ iff $\langle\psi(t)|\widehat{\Pi}_i|\psi(t)\rangle = 0$. So Hegerfeldt's theorem can be rephrased as

> **Hegerfeldt's theorem (indefiniteness form):** Given a system evolving unitarily over some interval of time under a Hamiltonian whose spectrum is bounded below, a given property is either (a) definitely *not* possessed at every time in that interval, or (b) *not* definitely not possessed at almost every time in that interval.

Put another way: suppose there is some property that, at some time in the indefinitely distant future, the system might have some probability to be found to possess. Then, according to the E-E link, it is *immediately* — that is, within an arbitrarily short window of time — indefinite whether the system has that property.

Put yet another way: anything that might at some future point be indefinite will be indefinite immediately. This seems to render the E-E link fairly useless as a description of ontology. We might have imagined that systems begin having some definite value of a given quantity, then gradually evolve so as to be indefinite across several values of that quantity, and in due course become completely indefinite with respect to that quantity (perhaps until some wavefunction collapse restores definiteness). But dynamically, that can't happen: indefiniteness is immediate if it is going to happen at all.

(I should, in fairness, acknowledge one context in which we seem to be able to get some content out of the E-E link even given Hegerfeldt's theorem. In the specific case of angular momentum (including both orbital angular momentum of some bound system, and the intrinsic spin of a particle), we could imagine

---

[7]In fact, it suffices for $\widehat{\Pi}$ to be a positive operator.
[8]See Halvorson and Clifton (2002) for discussion of its significance in this context.

the angular momentum *precessing*, so that the state at time t is an eigenstate of angular momentum with respect to some angle $\Omega(t)$. The system would then at all times have a definite angular momentum even though it would only be definite for an instant with respect to angular momentum in a given direction. But this relies on special features of the angular-momentum case, in particular does not generalise to position and momentum, and looks highly likely to be unstable once angular momentum couples to other degrees of freedom.)

To push the consequences of Hegerfeldt's theorem further (and also provide a rigorous justification of the claims of the previous section), suppose $\langle\psi(t)|\,\widehat{\Pi}\,|\psi(t)\rangle = 0$ for all t, and consider

$$\mathcal{S} = \mathrm{Span}\{|\psi(t)\rangle\}. \tag{31}$$

$\mathcal{S}$ is time invariant, and so must be spanned by (possibly improper) eigenstates of the Hamiltonian. And of course any element of $\mathcal{S}$ is an eigenstate of $\widehat{\Pi}$ with eigenvalue 0. So we can conclude that:

> **Complete indefiniteness corollary:** Unless some operator has some (possibly improper) eigenstates in common with the Hamiltonian, its associated quantity is completely indefinite at almost every time.
>
> (Readers uncomfortable with my casual use of improper eigenstates can just rephrase the requirement as "$\widehat{O}$ has an eigensubspace invariant under the Hamiltonian".)

As an application of this result, suppose that the Hamiltonian is non-degenerate (that is: has no two eigenstates with the same eigenvalue). Then a necessary and sufficient condition for a quantity not to be almost always completely indefinite is that it is a function of the Hamiltonian.

As another, consider some collection of scalar particles interacting via some potential:

$$\widehat{H} = \sum_i \frac{\widehat{P}_i^2}{2m_i} + V(\widehat{X}_1, \dots \widehat{X}_n), \tag{32}$$

for some smooth function V. In every case I know, and in particular in the case of free particles, the eigenfunctions of this Hamiltonian have only isolated zeroes. (This is easily provable in the case where V is a polynomial or other holomorphic function, so that the eigenfunctions themselves are holomorphic; it also follows in one dimension from the uniqueness theorem for solutions of ordinary differential equations.) But in that case, no eigenfunction has compact support, so no projector onto localised states is time-invariant. It follows that every particle has a completely indefinite position almost always.

The underlying problem here is a radical mismatch between the E-E link and the way quantum mechanics *actually* handles the idea of a system's becoming more spread out (speaking loosely) with respect to a given quantity. QM handles the latter through probabilities: the likelihood of a particle localised at x being

found very far from x is initially negligibly small and only gradually increases — and, depending on the dynamics, may never increase beyond negligible levels. But the E-E link is all-or-nothing: as soon as the system has any probability, even $10^{-10^{20}}$, of being found in some region, it is completely indefinite whether it is in that region.

I conclude that statements about a system's properties that rely on the E-E link convey essentially no information about a system between measurements (and, at the instant of measurement, everything empirically salient is coded in the Born rule in any case). As such, the E-E link cannot plausibly play a role in orthodox QM.

8 Quantum mechanics in practice: the lab view

Neither the collapse rule nor the E-E link can be part of orthodox, i. e. actually-used-in-practice, QM. It isn't that they are ultimately unsatisfactory on philosophical grounds, but rather that they are not even *prima facie* satisfactory, and lead to nonsense (all systems maximally indefinite all the time) or violation of empirical predictions (all measurements non-disturbing; continuous measurement impossible to define without Zeno freezing).

So what *do* physicists do, if they don't do "orthodox" QM as it is usually understood? In the rest of the paper I shall consider two paradigms to describe orthodox QM. The first — the "lab view" — is not uncommonly found in more careful foundational discussions of QM in the physics literature (especially in quantum information) but is not ultimately satisfactory to account for physical practice; the second — the "decoherent view" — does, I think, provide an adequate fit to physical practice.

In the Lab View (my presentation here is modelled on Peres (1993)), any application of QM should be understood as applying to some experimental setup, and 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 outcome of the measurement is given by the Born Rule.

The Lab View itself does not force a unique interpretation of the underlying physics, but it is often presented in parallel with a particular interpretation, each of which is sometimes claimed as 'orthodoxy'. Particularly prominent examples include:

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.

But most relevant for our purposes is:

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, via the E-E link. But at the final moment of measurement at the end of the experiment, the Projection Postulate is applied, jumping the system into an eigenstate of the quantity being measured.

Here we might seem to find a rehabilitation of orthodoxy. But consider: (i) as we have seen, the E-E link in practice tells us nothing about the physical state of the system between preparation and measurement, for it is almost certain that the system is maximally indefinite with respect to any quantity of interest pretty much throughout its evolution; (ii) collapse, occurring as it does at the very end of the physical process described by the Lab View, can do no actual work in physical predictions beyond what we already get from the Born rule.

17

(As a terminological aside, although contemporary physics often uses "Copenhagen interpretation" to refer to measurement-induced collapse, the historical views developed under that name are closer to complementarity and to straight instrumentalism. See Cushing (1994) and Saunders (2005) for further discussion.)

But in any case, the Lab View is itself insufficient to do justice to actual applications of QM, once they transcend the prepare-evolve-measure framework we have considered so far.

9   Limitations of the Lab View

We have already considered situations that go beyond what the Lab View, strictly speaking, can handle: those when the measurement process is not the end of our interaction with the system, where measurements are repeated or continuous. Furthermore, and even outside these cases, the Lab View's stipulation that measurement is primitive is itself in conflict with physical practice: measurement devices are physical systems, made from atoms and designed and built on the assumption that their 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?

The method used is in each case the same (and we have already seen it play out in our discussion of the projection postualte):

1. Insofar as the physics of the measurement process are relevant, we 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. We avoid infinite 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*.

The 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)). But it is really a special case of a more general requirement, for modern applications of quantum mechanics go beyond cases where measurement is repeated or continuous and embrace cases where we

cannot really avoid interpreting the QM probabilities as entirely separate from a formal 'measurement' process.

This is particularly clear in cosmology, where it has long been suggested that Lab View quantum mechanics is unsuitable 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.

Nor is the 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 pre-historic 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.[9]

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? Again, only in the most Procrustean sense.) 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:

---

[9]This is 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 world-view.

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.

10   Quantum mechanics in practice: the decoherent view

Once again: the point is not that the orthodoxy of the Lab View is conceptually inadequate, and so we must seek an unorthodox alternative; it is that physicists manifestly *are* doing quantum mechanics in regimes beyond the reach of the Lab View, so they must *already* have a method for applying it that goes beyond the Lab View.

In fact, the method is fairly obvious. 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 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 of course, this cannot straightforwardly 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[10] 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 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, 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

---

[10]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.

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.

To put the position intentionally crudely: orthodox QM, I am suggesting, consists of shifting between two different ways of understanding the quantum state according to context: interpreting quantum mechanics realistically in contexts where interference matters, and probabilistically in contexts where it does not. *Obviously* this is conceptually unsatisfactory (at least on any remotely realist construal of QM) — it is more a description of a *practice* than it is a stable *interpretation*. But why should that be surprising? Philosophers have spent decades complaining that physicists' approach to QM is philosophically unsatisfactory, after all. In a way, philosophers' version of 'orthodoxy' does physicists too much credit in providing a self-consistent realist account of QM that just lacks a satisfactory account of exactly when collapse happens, even as it does them too little credit in failing to recognise the unsuitability of the orthodox version of orthodoxy to physical practice.

And in fact, physics has made considerable progress in clarifying just when we can, and cannot, get away with a probabilistic interpretation of the quantum state, and in particular in helping us understand why we can reliably get away with it in macroscopic contexts. The *decoherence theory* developed by, *inter alia*, Joos and Zeh (Joos and Zeh 1985; Zeh 1993), Zurek (1991, 1998), Gell-Mann and Hartle (1993), Omnes (1988, 1992), and Griffiths (1984, 1993) in the 1980s and 1990s is concerned precisely with when the quantum state can be treated as probabilistic, understood either (in the environment-induced decoherence framework of Joos, Zeh and Zurek) because interference is suppressed with respect to some basis, or directly (in the consistent-histories framework of Griffiths, Omnes, and Gell-Mann and Hartle) by finding a consistent rule to assign probabilities to histories.[11] Hence my name for this position: the "decoherent view".

And where is collapse in all this? Well, the condition of decoherence can be reinterpreted as a condition for when we can impose an explicit collapse rule without empirically contradicting quantum theory. But that 'condition' is precisely the condition in which we can get away with treating the quantum state probabilistically, and from that perspective, "collapse" is just probabilistic conditionalisation. Of course, we can continue to think of the quantum state non-probabilistically, and use decoherence as a condition for when a *physical* collapse process can be introduced, but now we are well outside the assess-

---

[11]Appreciating that this is the task being performed by decoherence in contemporary physical practice also goes some way to explaining why the physics community has regarded decoherence as a major step towards understanding the interpretation of QM, something not generally shared by philosophers.(Barrett (1999, p.230) is typical: "That decoherence destroys simple interference effects does not solve the measurement problem since it does not explain the determinateness of our measurement records . . . In order to observe a single determinate record there must somewhere be a single determinate record.")

ment of orthodoxy, and well along the path towards a proposed *solution* of the measurement problem.

11 Two applications of the decoherent view

To illustrate the efficacy of the decoherent view in doing justice to physical practice, I consider two examples, from radically different sectors of physics: Stern-Gerlach type experiments, and the emergence of structure in the early Universe from primordial fluctuations of the inflaton field.

The Stern-Gerlach-type experiments I have in mind proceed as follows:

1. A beam of silver atoms emerges from a furnace.

2. That beam is split by a magnetic field and the beam corresponding to spin down in the z direction is discarded.

3. The beam is subjected to a series of interference experiments.

4. The spin of the atoms in the beam in (say) the x direction is measured by once again splitting the beam and measuring what fraction of atoms are in each beam.

Initially, the spin degrees of freedom of a silver atom is in a mixed state

$$\rho_1 = \frac{1}{2}(|+_z\rangle \langle +_z| + |-_z\rangle \langle -_z|). \tag{33}$$

(The justification of this state comes from quantum statistical mechanics and lies outside the scope of this paper.) After the magnetic field is applied, the particles spin and position degrees of freedom become entangled, having state:

$$\rho_2 = \frac{1}{2}(|+_z\rangle \langle +_z| \otimes |\varphi_+(t)\rangle \langle \varphi_+(t)| + |+_z\rangle \langle +_z| \otimes |\varphi_-(t)\rangle \langle \varphi_-(t)|) \tag{34}$$

where $|\varphi_+(t)\rangle$ and $|\varphi_-(t)\rangle$ are wavepacket states of negligible overlap. The mixed state cannot be used for interference experiments, so we can get away with treating it probabilistically. We now discard the $-$ part of the beam, and continue to operate only on the $+$ part; conditional on the silver atom still being in the apparatus, its spin state must be $+_z$, and so we update the state by conditionalising, to the pure state

$$|\psi_3\rangle = |+_z\rangle \otimes |\varphi_+(t)\rangle \tag{35}$$

(A more realistic treatment might allow for slight overlap of wavepackets, so that there is still some admixture of $|-_z\rangle$.)

Now we do a series of interference experiments with the system. At this point, treating it probabilistically will get us into trouble, so we avoid doing so: we continue to evolve the state unitarily and abjure probabilistic conditionalising.

Finally, we split the beam again, so it has form

$$|\psi_4\rangle = \alpha_+ \,|+_z\rangle \otimes |\chi_+(t)\rangle + \alpha_- \,|-_z\rangle \otimes |\chi_-(t)\rangle \qquad (36)$$

(with the values of $\alpha_\pm$ depending on the details of the interference processes, and with $|\chi_+(t)\rangle$ and $|\chi_-(t)\rangle$ again having negligible overlap. We once again treat this probabilistically (since we are going to do no further interference experiments, and indeed are about to entangle the system with a macroscopically large measurement device) and interpret $|\alpha_\pm|^2$ as the probability that the particle's spin is in fact $\pm$.

As for primordial structure formation, it works as follows (here I follow Weinberg (2008, pp.470–474), and excerpt a more detailed discussion in Wallace (2016)). A quantum field theory (the inflaton field) is coupled to spacetime geometry in a perturbative fashion, and allowed to evolve in time. In the very early Universe the system is inherently quantum-mechanical and the mod-squared amplitudes of the various modes of the field cannot consistently be interpreted probabilistically. But as the universe expands, the various mechanisms of decoherence come into play and — still very early in the Universe's history — we reach the point at which a probabilistic interpretation is consistent. At that point, we interpret those mod-squared amplitudes as probabilities of the actual modes of the inflaton field having various values; this determines a probability distribution over various possible inhomogeneities in the density of the early Universe, and that distribution is fed into cosmological simulations of structure formation. There is no measurement here, and no natural point for a collapse — only a quantum state which, in due course, we can get away with treating probabilistically.

Thes examples probably strike the reader as uncomfortably opportunistic, even ad hoc. Indeed, they should so strike the reader. The ad hoc, opportunistic approach that physics takes to the interpretation of the quantum state, and the lack, in physical practice, of a clear and unequivocal understanding of the state — *this* is the quantum measurement problem, once the distractions of collapse and the E-E link are removed.

12 Conclusion: the measurement problem from the perspective of contemporary physics

Quantum mechanics, as actually practiced in mainstream physics, makes no use of the eigenstate-eigenvector link, nor of the collapse postulate. Its dynamics are unitary; the unitarily evolving quantum state is interpreted inchoately, as describing physical goings on in regimes where interference is important and as describing probabilities in regimes where it can be neglected. On pain of failure to account for interference, we cannot (it seems) consistently treat the state as probabilistic; on pain of failure to account for the probability rule, and more generally of failing to make contact with observation, we cannot (it seems) consistently treat the state as representational. The "measurement problem"

from this perspective, is the task of taking this inchoate practice and showing how it can be justified given, as starting point, a well-defined physical theory — where what counts as a "well-defined physical theory" will depend on one's general stance on scientific realism and the philosophy of science. Perhaps we can do so by showing how an ultimately physical superposition nonetheless appears emergently probabilistic (Everett's strategy); perhaps we can do so by showing that interference can after all be made sense of on probabilistic grounds (the quantum Bayesian strategy, and the ψ-epistemic one); perhaps we can do so by adding additional representational structure (the Bohmian strategy) or by changing the dynamics to introduce a stochastic element (the dynamical-collapse strategy) or by adopting a conception of scientific theories that diverges from standard realism (the complementarity, quantum-logic, and instrumentalist strategies; perhaps the quantum-Bayesian strategy too).

From this perspective, the distinction between 'pure interpretations' that leave the formalism of QM alone, and modificatory strategies that modify or supplement it, is clear. The (real) Copenhagen interpretation, quantum Bayesianism, and the Everett interpretation (whatever their strengths or weaknesses otherwise) fall into the former category, as would a (hypothethical) ψ-epistemic interpretation: their dynamics is unitary throughout, their formalism unsupplemented by hidden variables. Dynamical-collapse and hidden-variable theories are in the latter category, being committed to adding additional variables and/or to modifying the Schrödinger equation.

From this perspective, too, the "orthodox interpretation" — that is, the theory obtained by adding the E-E link and the projection postulate to unitary quantum mechanics, and deriving the Born rule from them — is just one more modificatory strategy, and a strikingly implausible and unattractive one to boot. Perhaps some better attempt to solve the measurement problem will incorporate one or both, perhaps in modified and improved form — but it is time to retire the theory that is based on them as a starting point for discussions of the measurement problem.

13 Acknowledgements

References

Albert, D. Z. (1992). *Quantum mechanics and experience.* Cambridge, Massachussets: Harvard 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.

Barrett, J. A. (1999). *The quantum mechanics of minds and worlds*. Oxford: Oxford 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.

Belot, G., J. Earman, and L. Ruetsche (1999). The Hawking information loss paradox: the anatomy of a controversy. *British Journal for the Philosophy of Science 50*, 189–229.

Bokulich, A. (2014). Metaphysical indeterminacy, properties, and quantum theory. *Res Philosophica 91*, 449–475.

Bub, J. (1997). *Interpreting the Quantum World*. Cambridge: Cambridge University Press.

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.

Busch, P., P. J. Lahti, and P. Mittelstaedt (1996). *The Quantum Theory of Measurement* (2nd revised ed.). Berlin: Springer-Verlag.

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.

Cushing, J. T. (1994). *Quantum Mechanics: Historical Contingency and the Copenhagen Hegemony*. Chicago: University of Chicago Press.

Darby, G. (2010). Quantum mechanics and metaphysical indeterminacy. *Australasian Journal of Philosophy 88*, 227–245.

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.

Einstein, A., B. Podolsky, and N. Rosen (1935). Can quantum-mechanical description of reality be considered complete? *Physical Review 47*, 777–780.

Elitzur, A. C. and L. Vaidman (1993). Quantum mechanical interaction-free measurements. *Foundations of Physics 23*, 987–997.

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 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 (1993). Classical equations for quantum systems. *Physical Review D 47*, 3345–3382.

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.

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.

Halvorson, H. and R. Clifton (2002). No Place for Particles in Relativistic Quantum Theories? *Philosophy of Science 69*, 1–28.

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

Hegerfeldt, G. A. (1998a). Causality, Particle Localization and Positivity of the Energy. In A. Böhm (Ed.), *Irreversibility and Causality*, pp. 238–245. New York: Springer.

Hegerfeldt, G. A. (1998b). Instantaneous Spreading and Einstein Causality. *Annalen der Physik 7*, 716–725.

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. and H. Zeh (1985). The emergence of classical particles through interaction with the environment. *Zeitschrift fur Physik B59*, 223–243.

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

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

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

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.

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.

Page, D. (1994). Black hole information. In R. Mann and R. McLenaghan (Eds.), *Proceedings of the 5th Canadian Conference on General Relativity and Relativistic Astrophysics*, pp. 1–41. Singapore: World Scientific.

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.

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.

Rudin, W. (1991). *Functional Analysis* (2nd ed.). McGraw-Hill.

Ruetsche, L. (2011). *Interpreting Quantum Theories*. Oxford: Oxford University Press.

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

Skow, B. (2010). Deep metaphysical indeterminacy. *Philosophical Quarterly 58*, 851–858.

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

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

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

Wallace, D. (2016). Interpreting the quantum mechanics of cosmology. Forthcoming in volume on philosophy of cosmology, Ijjas and Loewer (ed.).

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

Wilson, J. (2016). Quantum metaphysical indeterminacy. Talk to the Jowett Society, Oxford, February 26 2016.

Wolff, J. (2015). Spin as a determinable. *Topoi 34*, 379–386.

Zeh, H. D. (1993). There are no quantum jumps, nor are there particles! *Physics Letters A172*, 189.

Zurek, W. H. (1991). Decoherence and the transition from quantum to classical. *Physics Today 43*, 36–44. Revised version available online at http://arxiv.org/abs/quant-ph/0306072.

Zurek, W. H. (1998). Decoherence, einselection, and the quantum origins of the classical: the rough guide. *Philosophical Transactions of the Royal Society of London A356*, 1793–1820. Available online at http://arxiv.org./abs/quant-ph/98050.