

Probabilities in Realist Views of Quantum Mechanics*

Christopher G. Timpson[†]

*Brasenose College, Oxford, OX1 4AJ, UK;
Faculty of Philosophy, The University of Oxford.*

March 1, 2024

1 Introduction

“...because the modals are among the most fundamental of all our concepts, their analysis can consist, as some grammarians and philosophers have insisted, only in showing their relations to other concepts and not in revealing some more basic concepts as their constituents.”
(White, 1975, p.3)

It is a commonplace, particularly since Carnap (e.g., Carnap (1955)), to note that there is more than one notion of probability. For Carnap, the varieties, of course, were *inductive* and *statistical* probability; the first being to do with the probability that theories (hypotheses, propositions) should be true; the latter arising in probability claims one might find made within theories, e.g., ‘For a gas at a fixed temperature, the probability that a particle will have an energy E decreases exponentially with E ’, or ‘The probability that a C^{14} atom will decay in 5730 years is $1/2$ ’. Carnap thought of statistical probability essentially in frequentist terms.

Later authors, such as Mellor (1969, 1971) and Lewis (1980), framed the dichotomy slightly differently; and we shall follow them. For them, the central distinction is between personal probability and objective chance. Personal probabilities are agents’ degrees of belief or credences that various events should occur or hypotheses be true; objective chances are mind-independent facts about the world knowledge of which would rationally constrain an agent’s degrees of belief to take on particular values. The objective chances specify how likely an individual event is to occur given the physical circumstances which obtain; it is a notion which applies to the single case. How one is to understand these objective chances—perhaps even whether they exist at all—is a controversial question, various aspects of which other authors in this volume have been and will be touching on; but it is important for our purposes to note at least this:

First, there is a good deal more consensus that objective chances exist than consensus on how they are to be understood. But when seeking to understand

*This paper is the original, longer, version of my chapter of the same title which appears in Beisbart and Hartmann (eds.) *Probabilities in Physics* OUP (2011).

[†]`christopher.timpson@bnc.ox.ac.uk`

fundamental concepts, we need to be careful not to set our goals unachievably high. Sometimes, the most one can hope for is to reveal interconnections between various different concepts one is interested in, and to gain understanding that way; for there is not some more basic set of concepts that these others can be explained in terms of. Plausibly this is so in the case of probability. Second, one crucial aspect of the notion of probability is that probabilities need to be capable of guiding our actions; they must interface with the rational assessment of our expectations¹. Third, these two strands of thought can be brought together in a proposition which we have already enunciated. Clearly we do not get terribly far with understanding the objective-chance kind of probability when we say things like ‘the objective chances specify how likely an event is to occur’, or ‘the objective chance tells us how much the circumstances favour one event’s occurrence over another’s’. But we do get a bit further in understanding when we draw the link with rational expectation: if I know that the objective chance of something’s happening is x , then I ought to have degree of belief x that it will occur; this is one of the most important consequences of the claim that the objective chance is thus-and-so. Here the circle of concepts has been extended a little wider². This link between rational expectation and objective chance is a very familiar idea; so familiar, in fact, as often to pass unremarked beneath our notice; but it is conceptually very important, for all that. Famously, it was dubbed the *Principal Principle* by Lewis (1980)³, judged deserving of that august label as in his view, the principle captures everything that we in fact know about the concept of objective chance (or very nearly everything): chances (if they exist) conditionally constrain rational degrees of belief.

One might be tempted to ask why—*why* do chances do this; or how do they manage it? But in one (one) important sense this question is misguided: the principle is a conceptual or analytic truth; it is constitutive of the concepts involved that they are related in this way. There’s no *why* to it. If nothing did this, there’d be no chances; while degrees of belief not suitably responsive to knowledge of chances would simply fail to be rational. There isn’t a question of first identifying the chances and then having to determine whether or not they actually constrain degrees of belief in the way they should. These aren’t two separate tasks. Finding the chances is automatically (of necessity) finding that which constrains the degrees of belief in the familiar way.

In one sense, then, the question may be misguided, but not perhaps in all. Nearby there clearly remain interesting questions which one can quite well ask with purpose, most pertinently, perhaps: what are the physical features which

¹It is for this reason, for example, that one needs the notion of probability to apply to the single case: I need to decide what to do now, in the very next instance; to do that I need to know what the chances of the various particular possibilities I’m faced with are. It’s no good being told what the distribution of eventualities in the long run is or in an ensemble of cases; I need to know what pertains to my individual case; the one I face.

²But not enormously so: we don’t have an exceptional grip on the pertinent notion of rationality which does not itself involve probabilistic notions at some point. Why is x the rational degree of belief to have in whether A will occur? Because the circumstances favour A ’s happening to just that extent. Nevertheless it may be helpful to get a handle on what it is for one event’s happening to be favoured over another’s in terms of that fact making it rational to expect one over the other.

³A little more precisely, the principle says that conditional on knowledge of the chances and any *admissible* background information, my degree of belief will be as the chances have it. Admissible information is any information which does not already encode the fact of what the result would actually be.

give rise to the chances? It is primarily this kind of question with which we will be concerned in the following; the question of the link between objective chance and degrees of belief will become particularly important once more, though, in Section 5 when we turn to consider the Everett interpretation; and it will also loom large when we come to draw our final conclusions.

Our concern is with probability in realist approaches to quantum mechanics. I have in mind those approaches which are realist in the sense that they take the quantum formalism seriously as describing the world at a fundamental level, seamlessly and as a whole. (Though it may perhaps be felt that various additions to the bare quantum formalism are required to be made too, to manage the job.) That is, the formalism (perhaps with these additions) is at all times to apply to everything in the same way; there are to be no divisions into domains where the formalism applies and those where it does not; and the dynamics is to proceed uniformly throughout. The properties that the fundamental constituents of the world possess are to be read off fairly directly from the complete physical state postulated by the theory. A key role in providing these realist descriptions will be played by the quantum state. The world, or Universe, is given to us only once, so these states will apply to individual systems—the world as a whole and its subsystems—not to ensembles. Moreover, we shall primarily be concerned with those views in which the quantum state itself stands directly for a part of the ontology: there is to be a field-like entity in the world to which the quantum state corresponds. If $|\Psi\rangle$ is the (pure) universal quantum state, call the physical item it corresponds to the Ψ -*field*. At any time, there are unequivocal, mind-independent facts about what the quantum state of a given system is: the state is objective.

Examples are familiar: we shall be concerned with three: the GRW theory (as a stand-in for more realistic dynamical-collapse theories); the de Broglie-Bohm theory (in its own right and in view of the insight it gives into general deterministic hidden-variable theories); and the Everett interpretation (which boldly goes it alone with the universal quantum state and the unitary dynamics). The largest class missing here, which I shall not be discussing, are modal interpretations⁴. We will also restrict attention to the case of many-body non-relativistic quantum mechanics; for quite enough issues of interest arise here.

The two questions which will organise our discussion are these:

1. In what manner does probability enter the theory? and
2. What is the status of the standard probabilistic rule—the Born rule—in the theory?

We will take each of our examples in turn. But first, however, we need some further preliminaries.

2 States and Probabilities

We think of quantum mechanics as being in some important and historically novel sense a probabilistic theory. In standard presentations the quantum state

⁴Although note that the de Broglie-Bohm theory may be seen as a particular kind of modal interpretation, in which the preferred variable is taken to be position (configuration) Bub (1997).

is usually introduced and motivated by means of a probabilistic interpretation⁵. Thus we often start off in fairly operationalist terms: the quantum state of a given system tells us the probabilities for the outcomes of various measurement procedures on that system, where we make no effort to analyse what's involved in a measurement, beyond labelling each possible outcome with a particular operator. In the general case, for a measurement \mathcal{M} , each possible outcome i (think of a dial on a black-box measuring device with a pointer pointing in some direction to indicate the result) is associated with some positive operator E_i on the Hilbert space \mathcal{H} of the system⁶, where $\sum_i E_i = \mathbf{1}$; they sum to the identity. (Such operators are often known as *effects*. In the most familiar case, they will be projectors P_i onto subspaces associated with particular eigenvalues of observable quantities of interest.) The quantum state is then introduced as a positive normalised linear functional of operators on \mathcal{H} ; mathematically, this will be given by a density operator ρ on \mathcal{H} , a positive operator of unit trace. The point is that such a normalised linear functional will assign numbers to the operators E_i which will satisfy the axioms of the probability calculus: they can be interpreted as probabilities that a given outcome should occur. The rule is that the probability $p(i)$ of obtaining outcome i on measurement of \mathcal{M} is:

$$p(i) = \text{Tr}(\rho E_i). \quad (1)$$

This is the *Born rule*. Observe that given that ρ is normalised and the E_i are positive and sum to the identity, each of the $p(i)$ will be a real number between 0 and 1 where $\sum_i p(i) = 1$; moreover due to the linearity of the functional $\text{Tr}(\rho \cdot)$, for two distinct outcomes i and j of \mathcal{M} , the probability that one or other will occur, $p(i \vee j)$, will be $\text{Tr}(\rho(E_i + E_j)) = p(i) + p(j)$, as required.⁷

In the case in which the system is in a pure state $\rho = |\psi\rangle\langle\psi|$ and one is measuring some maximal (non-degenerate) observable with eigenvectors $|\phi_i\rangle$ (so that $P_i = |\phi_i\rangle\langle\phi_i|$), (1) takes the very familiar form:

$$p(i) = \text{Tr}(|\psi\rangle\langle\psi| |\phi_i\rangle\langle\phi_i|) = \langle\psi|\phi_i\rangle\langle\phi_i|\psi\rangle = |\langle\psi|\phi_i\rangle|^2. \quad (2)$$

Now this is all very well and good. These algorithms (1) and (2) serve to connect the quantum formalism to empirical data and predictions in the lab, at least in a minimal (albeit essential) kind of way. But from the perspective of a realist view of quantum mechanics, it is very far from the *fundamental* role of the quantum state that it provide probabilities for measurement outcomes. Rather,

⁵It was Born (1926) of course, who first posited the probabilistic interpretation of the quantum state, though to begin with he misassociated probabilities with the amplitude rather than the modulus squared of the state.

⁶A linear operator A on a Hilbert space \mathcal{H} is positive iff $\forall|\psi\rangle \in \mathcal{H}, \langle\psi|A|\psi\rangle \geq 0$. The eigenvalues of such an operator will all be greater than or equal to zero. Positive operators on a complex Hilbert space will be Hermitian.

⁷The import of Gleason's remarkable theorem (Gleason, 1957) is that (1) is the *only* expression which will provide a normalised real function on outcomes of projective measurements (i.e., where $E_i = P_i$) which is additive for orthogonal projectors, at least for dimensions greater than two. We note, of course, that this additivity requirement is stronger than just that probabilities of outcomes for a *given* measurement be additive (be probabilities). Controversially, it connects different measurement processes too, as a given projector will belong to more than one orthogonal set summing to the identity. It is a requirement of *non-contextuality* of probabilities, cf., famously, Bell (1966). For extension of the theorem to the more general case of positive operators, where it applies also in the $d = 2$ case, see Busch (2003); Caves et al. (2004).

the state is to play a role in describing what is fundamentally *in* the world and how it behaves (the occurrent categorical features of the world). The notion of measurement and of probabilities for measurement outcomes will have no part to play at all in characterising the state at the basic level. Measurement processes won't be picked out in the fundamental story about how the world is; rather they will just be one particular kind of dynamical interaction amongst very many others, of no especial interest or importance save anthropocentrically. *Prima facie* it is not at all obvious how the fundamental descriptive role of the state should relate to, or allow for, a Born rule-style probabilistic interpretation. If Born rule probabilities are to arise from the state, they must do so as a corollary of the analysis of those specialised procedures we call measurement interactions within the fundamental terms of the theory, which terms will not themselves make reference to measurements, outcomes, observers or probabilities.

Nevertheless, however Born rule probabilities are to arise, arise they must, at least to a good approximation. The empirical good-standing of quantum mechanics is based on the predictions that issue from the Born rule. If we can't recover Born rule probabilities within our fundamental dynamical story then the theory we are considering must be deemed empirically inadequate. It is a *sine qua non* that a role for probability be found within our realist views and that the probabilities match (to reasonable degree) those given by the Born rule, for all that the state in such theories is not to receive an interpretation fundamentally in terms of probabilities for measurement outcomes.

The set of density operators—potential quantum states—is convex. That is to say, if ρ_1 and ρ_2 are density operators and $0 \leq \lambda \leq 1$, then $\rho = \lambda\rho_1 + (1-\lambda)\rho_2$ is also a member of the set. The pure states $\rho = |\psi\rangle\langle\psi|$, $\forall |\psi\rangle \in \mathcal{H}$, make up the extremal elements of the set: these cannot be reached by convex combinations of other elements of the set; they also comprise the boundary of the set⁸. Density operators which may be arrived at by taking convex combinations of pure states, and thus lie inside the boundary, are called *mixed states*. The more mixed a state, as measured for example by $-\text{Tr}\rho \log \rho$, the von Neumann entropy (cf. Wehrl (1978)), or by $\text{Tr}(\rho^2)$ (Fano, 1957), the length of the state in the Hilbert-Schmidt norm on operators, the less able one is to predict the result of measurements on systems in that state (cf. Timpson (2003)). The more mixed a state, the more spread-out the probabilities for measurement outcomes it provides are.

Convexity is a very natural structure to the set of states when one commences with a probabilistic interpretation of the quantum state. If a preparation procedure itself involves probabilities, so that it may produce a system in state ρ_1 with a probability λ or in state ρ_2 otherwise, then the probability of getting outcome i on a subsequent measurement of \mathcal{M} on the system will simply be the sum of the probabilities which would follow from ρ_1 and ρ_2 individually, weighted by the probability that they were in fact produced:

$$p(i) = \lambda\text{Tr}(\rho_1 E_i) + (1 - \lambda)\text{Tr}(\rho_2 E_i). \quad (3)$$

Eqn. (3), by linearity, is equivalent to:

$$p(i) = \text{Tr}((\lambda\rho_1 + (1 - \lambda)\rho_2)E_i); \quad (4)$$

⁸This is a distinctive property of the quantum case: in general, convex sets can have boundary elements which need not also be extremal.

that is, considering the probabilistic preparation procedure as a whole, we can take it simply to produce a system in state $\rho = \lambda\rho_1 + (1 - \lambda)\rho_2$, as this state captures all the statistics that we will expect to see on measurement.

But while convexity is a natural structure given the probabilistic interpretation, it is less obviously so in a realist interpretation of quantum mechanics. The realist may simply have to take this state-space structure as given (it's the structure of the empirically adequate theory). They may nevertheless make use of the notion of a mixed state of course. For a realist interpretation there will always be an unequivocal fact about what the state of a given system is, yet one might still be ignorant of what that state is; and under that constraint, for the purposes of prediction (so long as one has suitably recovered the Born rule in one's theory), the density operator given by the appropriate convex combination of the possible states is the right one to work with; while recognising that there is a fact about what the true state of the system in question actually is (which fact this density operator does not represent). This conclusion is reinforced when we recognise that there is no way to identify perfectly the quantum state of an individual system⁹. The best one can do is make measurements on very many identically prepared systems and thereby infer the density operator giving the expected statistics for measurements on systems prepared in that way (this is sometimes called *quantum state tomography*). If the preparation procedure involves probabilities, as above, then the mixed state $\rho = \lambda\rho_1 + (1 - \lambda)\rho_2$ is the one which would be inferred and would be the appropriate one to assign to each system being produced by the preparation procedure. By contrast, if one could determine the states of individual systems, then the appropriate states to assign would be ρ_1 or ρ_2 depending on which of these each individual system produced was actually in.

We have been working here primarily with an *ignorance* construal of mixed states. If the actual (realist) quantum state of a system is pure, then a mixed state will only be assigned due to ignorance of what that pure state is. But as is very well known, there is another way in which mixed states arise: in the presence of entanglement, the states (*reduced states*) assigned to subsystems of the entangled system (by in each case tracing out all the other subsystems but one) will be mixed, even if the overall state for the system is pure. For these mixed states, there is no further fact about what true state the subsystem might actually be in: the specified mixed state is it. This is the distinction between *properly* and *improperly* mixed states due to d'Espagnat (1976). The former arise due to ignorance of the actual state (which may itself be pure or mixed) while the latter arise due to entanglement¹⁰.

It's an important feature of the set of density operators that there are very many different ways in which a given density operator may be expressed as a convex combination of others (see Hadjisavvas (1981); Hughston et al. (1993)

⁹See Busch (1997) for a very nice general discussion.

¹⁰Here's the argument why we need to distinguish the two. Consider some N -party entangled state. By definition, this state cannot be expressed as an N -party product state, nor a convex combination of such states (a separable state). The reduced state for each subsystem individually will be some mixed state. Assume that each such state may be given an ignorance interpretation: there is some underlying less-mixed (possibly pure) state that each subsystem actually has. The true state of the N -party system would then simply be the tensor product of each of these true states for subsystems, or a convex combination of these if there were further correlations between them. But then the total state would not be entangled. Thus reduced states of entangled systems cannot be given an ignorance interpretation. QED.

for the extent of this freedom). For improperly mixed states, there is no correct convex decomposition: any one will do, just as one might choose any system of coordinates one wishes to describe a situation: there is no fact of the matter about which is the right one. For properly mixed states, though, when one is an objectivist, there certainly will be a correct decomposition: it will be the one giving the chances with which the actual possibilities might have been produced¹¹. Moreover there will be a fact about which of these possibilities is the actual state.

The upshot of all this is that there are two distinct, but connected, ways in which probabilities may feature in a realist view of quantum mechanics. First and foremost we need to recover the correct probabilities (chances) for the outcomes of measurements. But we may also consider probability distributions over what the actual state might be. In what follows we shall concentrate on the former and will make the usual assumption that the total state of the Universe is pure. Generic pure states of complex systems will be massively entangled.

3 Dynamical Collapse: The GRW Theory

Suppose we just start with the quantum state of the Universe, $|\Psi\rangle$, and assume it to evolve unitarily, in accord with the Schrödinger equation. The fundamental ontology of this theory is pretty sparse: we have some spatio-temporal structure within which lives the Ψ -field. Everything else, including the macroscopic features of our experience, will be supposed to supervene on how this physical entity is arranged over time in its spatial arena. Commonly, for N -body non-relativistic quantum mechanics (our simplified model), the arena within which the Ψ -field evolves is taken to be rather high dimensional: a $3N$ -dimensional space, as the position basis representation of the quantum state $|\Psi(t)\rangle$ is the wavefunction $\Psi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N, t)$ (suppressing indices for internal degrees of freedom such as spin) which is mathematically defined on a $3N$ -dimensional space at a given time. In this picture, the three-dimensional space of our experience will need to be one of the things which emerges from the underlying ontology. But one needn't take this line: it is also possible to think of the ontology as involving a non-separable field living in usual four-dimensional spacetime (see Wallace and Timpson (2010) for discussion). In either case, though, we face a problem: it's a familiar one about measurement¹².

Let's think about how an idealised measurement process might be modelled in this theory. Take $\{|\phi(\mathbf{x})\rangle\}$ to be a set of states each localised around position \mathbf{x} ; these might be narrow Gaussians, in the position basis. Then some macroscopic object having (approximately) well localised position, such as a lab bench or the pointer on some measuring device, might be modelled by some product of such states for each of its constituent parts. Very roughly, then, the state $|\chi(\mathbf{x})\rangle_M$ (' M ' for *Macro*) of some macroscopic object localised around position

¹¹For an argument that the proper/improper distinction still gets a grip in no-collapse quantum mechanics, see Timpson and Brown (2005).

¹²It is sometimes felt that there is another problem too: that even in a dynamical collapse theory, the Ψ -field on its own is insufficient to ground a reasonable ontology (e.g., Maudlin (2007); Allori et al. (2008)). I remain unconvinced, but in any case, as our main target is probability, not controversies over quantum ontology, this point need not detain us.

\mathbf{x} will be something like:

$$|\chi(\mathbf{x})\rangle_M = |\phi(\mathbf{x})\rangle_1 |\phi(\mathbf{x})\rangle_2 \dots |\phi(\mathbf{x})\rangle_K, \quad (5)$$

where K is the number of microscopic systems comprising the macroscopic object¹³. In terms of the fundamental ontology, this determinate position for a macroscopic object corresponds to (supervenes on) some hump of intensity of the Ψ -field. The mathematical representation of this hump will, in the position basis, be a narrow $3K$ -dimensional Gaussian centred on $(\mathbf{x}_1 = \mathbf{x}, \mathbf{x}_2 = \mathbf{x}, \dots, \mathbf{x}_K = \mathbf{x})$.

Now suppose we wish to measure some observable quantity which is associated with a (non-degenerate) set of eigenstates $|\xi_i\rangle_S$ (' S ' for *System*). A decent measuring apparatus for that quantity will have to have a set of indicators which will reliably tell us if the system being measured is in one of those eigenstates; and which one if so. We can take different positions of a pointer to play this role; and there will be a corresponding set of states $|\chi(\mathbf{x}(i))\rangle_M$ for when the pointer is in one of those positions $\mathbf{x}(i)$. Take $|\chi(\mathbf{x}(0))\rangle_M$ to be the 'apparatus ready' state.

We will wish to arrange that the natural dynamics for the system-apparatus interaction will take

$$|\xi_i\rangle_S |\chi(\mathbf{x}(0))\rangle_M$$

to

$$|\xi_i\rangle_S |\chi(\mathbf{x}(i))\rangle_M.$$

(This is the case for non-disturbing and for non-destructive measurements; incorporating these more general scenarios does not change matters substantively). Then the position of the pointer will allow one to infer what the eigenstate was.

So far so good. If the system comes into the apparatus in an eigenstate, then the pointer of the apparatus will respond, moving from the 'ready' position to another macroscopically distinct position, thereby indicating the eigenstate. In terms of the fundamental ontology, the part of the Ψ -field constituting the macroscopic pointer will now hump in a different place, correlated with the part of the Ψ -field constituting the system. But of course, what happens if the system to be measured comes in in some superposition $\sum_i \alpha_i |\xi_i\rangle_S$? The linear (since unitary) Schrödinger dynamics will take us to an entangled state:

$$|\Psi\rangle = \sum_i \alpha_i |\xi_i\rangle_S |\chi(\mathbf{x}(i))\rangle_M |\mathcal{E}\rangle_E, \quad (6)$$

(where $|\mathcal{E}\rangle_E$ is the state of the rest of the Universe, the environmental degrees of freedom not yet involved in the measuring process). And this is problematic. For we now have a state in which the pointer is in macroscopically distinct positions at the same time (the Ψ -field has humps in various different locations, each constituting a pointer pointing in some different direction); and this is very far from constituting a unique result of the measurement; and it is not what we in fact observe. Something seems to have gone very wrong: thus the problem of measurement.

¹³Here, as elsewhere, I have ignored questions of symmetrisation for simplicity; these don't affect the argument. In such a more realistic treatment (5) would involve entanglement.

As Bell remarks in his characteristically lucid discussion of the GRW theory (Bell, 1987a), this is just the problem Schrödinger had with his cat (Schrödinger, 1935):

He [Schrödinger] thought that she could not be both dead and alive. But the wavefunction showed no such commitment, superposing the possibilities. Either the wavefunction, as given by the Schrödinger equation, is not everything, or it is not right. (Bell, 1987a, p.201)

Call this Bell's Dilemma. The general point of a dynamical collapse theory is to explore the latter horn of the dilemma: to seek to modify the fundamental dynamics so that embarrassing macroscopic superpositions like (6) do not arise. But this is a delicate business; at the same time as killing off embarrassing superpositions, one needs to keep the virtuous ones; the ones we know do obtain. We don't want to ruin everything by changing the dynamics, just deal with this problem about macrosuperposition. The theory of Ghirardi et al. (1986) was the first neat and convincing suggestion that this might plausibly be done¹⁴.

Clearly we need to move away subtly from the unitary dynamics. A very general framework which will encompass both unitary and non-unitary dynamics, but still within the broad structure of quantum theory, is provided by the notion of a *quantum operation* (see Nielsen and Chuang (2000, chpt. 8) for a clear and self-contained discussion). A quantum operation \mathcal{O} is a trace non-increasing completely positive map. A positive map is a linear operator which maps positive operators to positive operators; complete positivity requires that the action of the map on a subsystem leaves a positive operator on a larger system containing the subsystem positive too. The point is that such maps will (even in the presence of entanglement) map density operators to density operators, up to normalisation.

The general form for a quantum operation is given by $\mathcal{O}_i(\rho) = \sum_k O_{ik}\rho O_{ik}^\dagger$, where the O_{ik} are linear operators on the Hilbert space of the system for which $\sum_k O_{ik}^\dagger O_{ik} \leq \mathbf{1}$. In fact this sum will itself constitute a positive operator; and if we choose carefully, we can arrange it so that $\sum_{ik} O_{ik}^\dagger O_{ik} = \mathbf{1}$, i.e., each of the $\sum_k O_{ik}^\dagger O_{ik} = E_i$ is an effect operator. Following the action of a quantum operation on ρ , we won't yet have another allowed quantum state until we re-normalise (it is this which introduces non-linearity into the state-change rule unless the map was trace-preserving). Thus the state following the physical process represented by the operation \mathcal{O}_i will be:

$$\rho^i = \frac{\mathcal{O}_i(\rho)}{\text{Tr}(\mathcal{O}_i(\rho))} = \frac{\sum_k O_{ik}\rho O_{ik}^\dagger}{\text{Tr}(\rho E_i)}, \quad (7)$$

since $\text{Tr}(\sum_k O_{ik}\rho O_{ik}^\dagger) = \text{Tr}(\rho \sum_k O_{ik}^\dagger O_{ik}) = \text{Tr}(\rho E_i)$, due to the cyclicity of the trace. Now follows an important point. In the quantum operation formalism, it is specified that the probability that a process represented by an operation \mathcal{O}_i should occur is given by $\text{Tr}(\rho E_i)$, where E_i is the effect given by summing the operation elements O_{ik} , as above. From the perspective of a probabilistic

¹⁴Theories along the same lines fixing some of the flaws in the original GRW proposal soon followed, for example the *Continuous Spontaneous Localisation* theories developed by Pearle (1989) and others. See Bassi and Ghirardi (2003) for a comprehensive review of the general programme. Major difficulties still loom in making these kinds of theories compatible with relativity, but see Fleming (1988); Myrvold (2002, 2003); Tumulka (2006) for some proposals.

interpretation of the quantum state, this makes a lot of sense. If a fixed range of processes \mathcal{O}_i might occur, each with probability $p(i)$, but we don't know which, then the subsequent properly mixed state ρ' will be $\sum_i p(i)\rho^i$. If $p(i)$ is given by $\text{Tr}(\rho E_i)$, then that sum becomes $\sum_i \mathcal{O}_i(\rho)$, which will be of unit trace; and ρ' too will then be of unit trace, i.e., an allowed density operator. However, along the lines we noted before, this reasoning is far less compelling in a realist setting where the fundamental role of the state is not made out in terms of probabilities: if quantum states are supposed primarily to describe how things are, then why should my ignorance about what actual state a system is in be itself captured mathematically by something which is a candidate quantum state? In response perhaps the best one can do is, once more, simply to point to the fact that these are the given structures within an empirically successful theory.

At any rate, this is the general setting. What GRW proposed was that there is a fixed probability per unit time that each individual microsystem will be subject to a spontaneous localisation process; at all other times, meanwhile, systems evolve according to the Schrödinger dynamics. If the probability per unit time for a localisation to occur (a 'hit' as they are sometimes known) is chosen judiciously—i.e., it is suitably low—then isolated quantum systems will tend to evolve unitarily and there is only an appreciable chance of a localisation to occur for (very) large agglomerations of systems.

A GRW localisation process takes the following form. Take the function $\phi(\mathbf{x}_i - \mathbf{x})$ to be a Gaussian function

$$\phi(\mathbf{x}_i - \mathbf{x}) = R \exp\left(\frac{-(\mathbf{x}_i - \mathbf{x})^2}{2a^2}\right), \quad (8)$$

which is centred on \mathbf{x} with a standard deviation (width) of a , which we make very small. R is a suitably chosen normalisation factor (of which more in a moment). Now we can define (helping ourselves harmlessly to the calculational tool of position eigenstates $|\mathbf{x}_i\rangle$) a Gaussian quantum operation element—a localisation operator—as:

$$\hat{\phi}(\mathbf{x}_i - \mathbf{x}) = \int d^3x_i \phi(\mathbf{x}_i - \mathbf{x}) |\mathbf{x}_i\rangle\langle\mathbf{x}_i|. \quad (9)$$

Notice that this operator is self-adjoint with positive eigenvalues, i.e., a positive operator. It operates on the subspace of the total Hilbert space which corresponds to system i (in what follows, take its action on the rest of the Hilbert space to be the identity). Take $|\psi\rangle$ to be an N -body quantum state. The (unnormalised) result of the action of the localisation on the state will be:

$$|\psi\rangle\langle\psi| \mapsto \hat{\phi}(\mathbf{x}_i - \mathbf{x}) |\psi\rangle\langle\psi| \hat{\phi}(\mathbf{x}_i - \mathbf{x}). \quad (10)$$

Now if $\phi(\mathbf{x}_i - \mathbf{x})$ is a Gaussian function, then $\phi^2(\mathbf{x}_i - \mathbf{x})$ is too, though with a narrower width of $a/\sqrt{2}$. Then, correspondingly, $\hat{\phi}^2(\mathbf{x}_i - \mathbf{x})$ will also be a positive operator; and if we have chosen the normalisation factor R correctly, then:

$$\int d^3x \hat{\phi}^2(\mathbf{x}_i - \mathbf{x}) = \mathbf{1}, \quad (11)$$

that is, $\hat{\phi}^2(\mathbf{x}_i - \mathbf{x})$ is an effect operator (summing over the different possibilities—different centres of localisation \mathbf{x} —adds up to unity) corresponding to the localisation procedure (10), which is a trace-decreasing quantum operation. (In

this case the localisation quantum operation only involves a single operation element, i.e., $k = 1$.)

What all this means is that the probability density that system i will be hit by a localisation represented by $\widehat{\phi}(\mathbf{x}_i - \mathbf{x})$, i.e., by a localisation which is centred on the point \mathbf{x} , will be given (following the usual formula for quantum operations) by:

$$p_i(\mathbf{x}) = \text{Tr}(|\psi\rangle\langle\psi| \widehat{\phi}^2(\mathbf{x}_i - \mathbf{x})) = \langle\psi| \widehat{\phi}^2(\mathbf{x}_i - \mathbf{x}) |\psi\rangle. \quad (12)$$

In other words, when the component of $|\psi\rangle\langle\psi|$ corresponding to system i (i.e., the reduced state of i) has a large weight in a narrow region centred around \mathbf{x} , the probability of a localisation centred on \mathbf{x} will be high; while if the weight is low, the probability of a localisation occurring centred there is low too.

Following normalisation (as in (7)), localisation will take the starting pure state of the system to another pure state. It is easiest to see the effect of a localisation in the position basis. Assume (without loss of generality) that the localisation occurs to system 1 (we can always think of simply labelling whichever system got hit ‘system 1’). The result of (10) is simply to multiply the wavefunction of the system by the Gaussian (8):

$$\psi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N) \mapsto \phi(\mathbf{x}_1 - \mathbf{x}) \psi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N). \quad (13)$$

If we just had a single system ($N = 1$) then the result is easy to visualise: Following the localisation, the wavefunction of the system would now be sharply peaked around point \mathbf{x} in space, under the envelope provided by $\phi(\mathbf{x}_1 - \mathbf{x})$. If previously the wavefunction had had significant spatial extension, this would now be sharply curtailed under the exponentially decreasing limits of the Gaussian. With a many body system, the result is slightly less straightforward to visualise. Whilst the localisation only acts on the first system’s subspace, if the N -body state is entangled, this can still have non-trivial effects on the other systems (*vide* the EPR argument!). Taking an orthonormal set of functions $\eta_j(\mathbf{x}_1)$ for the first system, we can express the N -body wavefunction as:

$$\psi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N) = \sum_j \beta_j \eta_j(\mathbf{x}_1) \tilde{\psi}_j(\mathbf{x}_2, \mathbf{x}_3, \dots, \mathbf{x}_N), \quad (14)$$

where the $\tilde{\psi}_j(\mathbf{x}_2, \mathbf{x}_3, \dots, \mathbf{x}_N)$ are a set of wavefunctions for the remaining systems¹⁵. Following localisation of system 1 we have, of course:

$$\sum_j \beta_j \phi(\mathbf{x}_1 - \mathbf{x}) \eta_j(\mathbf{x}_1) \tilde{\psi}_j(\mathbf{x}_2, \mathbf{x}_3, \dots, \mathbf{x}_N).$$

For each term in the sum, the effect on the first system is just as before: the previous wavefunction is squashed under the narrowly concentrated envelope of the Gaussian. For the other systems, this means that the amplitude of the component $\tilde{\psi}_j$ in the total state will be affected, for each j ; leading overall to a change to the state for the other systems too. By design, the effect is most

¹⁵In fact, from the biorthogonal (Schmidt) decomposition theorem, $\{\eta_j(\mathbf{x}_1)\}$ and $\{\tilde{\psi}_j(\mathbf{x}_2, \mathbf{x}_3, \dots, \mathbf{x}_N)\}$ can be chosen in such a way that both constitute an orthonormal set for their respective spaces, though we shan’t need this.

marked when we have a superposition of macroscopically distinct positions. If (14) were a wavefunction like:

$$\sum_j \beta_j \phi(\mathbf{x}_1 - \mathbf{x}(j)) \phi(\mathbf{x}_2 - \mathbf{x}(j)) \dots \phi(\mathbf{x}_N - \mathbf{x}(j)), \quad (15)$$

i.e., a superposition of distinct N -body wavefunctions each well localised around some $\mathbf{x}(j)$ (compare our localised state (5) from earlier); and if the $\mathbf{x}(j)$ are macroscopically separated positions the distances between which are enormously larger than the Gaussian width a , then when system 1 is hit by a localisation, the only term in the superposition which will survive is one where $\mathbf{x}(j)$ is very close to the centre of the localisation \mathbf{x} . That is, all but one of the individually localised, but superposed, position states will be killed off, to leave only one component with non-vanishing weight¹⁶. And this—again by design—is of course the fate of our embarrassing post-measurement superpositions like (6). For a macroscopic object, one composed of very, very many microsystems, even though the probability per system per unit time of a localisation is very low, the sheer number of component systems mean it is practically certain that at least one component system within it will suffer a localisation; and that single event on its own is enough to kill off all but one element in a superposition of macroscopically distinct positions. The usual specification for the parameters in the GRW theory is that the probability per unit time per system for a localisation might be roughly 10^{-15}s^{-1} ; and the width of the Gaussian about 10^{-7}cm . An individual system would then have a lifetime of about 10^8 years; but if our pointer comprised (conservatively) of the order of 10^{20} systems, the expected lifetime before it was subject to a localisation would be 10^{-5}s . One would be free to tweak these parameters in response to experiment, of course.

It is crucial for the GRW theory to work that the probability distribution for the locations of the centres of hits be given by the expression (12); which one might notice is exactly the expression for the Born rule as applied to approximate position measurements under a probabilistic interpretation of the quantum state. With this assumption in place (and it is an assumption—nothing about the details of the localisation process forced us logically to make it), localisations will be drawn to spatial locations given high weight by the quantum state; thus in a state like (15), a localisation will be overwhelmingly likely to occur near *one* of the $\mathbf{x}(j)$ positions. Moreover—and this is essential—it guarantees that we recover the Born rule appropriately.

For a measurement procedure like the one above which led to the superposition (6) we want to end up saying that we will get the outcome corresponding to eigenvalue ξ_i (etc.) with a probability $|\alpha_i|^2$; the weight of the corresponding component of the quantum state. Now the probability that the i th component of the state will be left (that outcome obtains) given a localisation (call the system in the pointer it happens to ‘ l ’) will be given by summing all the probabilities that the centre for the localisation be sufficiently near $\mathbf{x}(i)$; for then and only then will the i th component be left:

$$p(i|\text{localisation}) = \int_C d^3x p_l(\mathbf{x}) = \int_C d^3x \langle \Psi | \hat{\phi}^2(\mathbf{x}_l - \mathbf{x}) | \Psi \rangle, \quad (16)$$

¹⁶Strictly speaking, because a Gaussian has tails stretching to infinity—admittedly exponentially decreasing ones—there will be tiny, tiny amounts of the other components of a superposition like (15) left after localisation; but they will be vanishing small.

where C is the region close to $\mathbf{x}(i)$ —close on the order of the width a . For state (6), this equates to:

$$\int_C d^3x \sum_i |\alpha_i|^2 \langle \phi(\mathbf{x}_l - \mathbf{x}(i)) | \hat{\phi}^2(\mathbf{x}_l - \mathbf{x}) | \phi(\mathbf{x}_l - \mathbf{x}(i)) \rangle, \quad (17)$$

where $|\phi(\mathbf{x}_l - \mathbf{x}(i))\rangle$ is a Gaussian state for l localised near $\mathbf{x}(i)$. But the inner product $\langle \phi(\mathbf{x}_l - \mathbf{x}(i)) | \hat{\phi}^2(\mathbf{x}_l - \mathbf{x}) | \phi(\mathbf{x}_l - \mathbf{x}(i)) \rangle$ will be approximately equal to zero unless \mathbf{x} is close to $\mathbf{x}(i)$, so we lose all but the i th term in (17); while correspondingly the integral of the inner product over C will be approximately equal to one, thus:

$$p(i|\text{localisation}) \approx |\alpha_i|^2, \quad (18)$$

as required. The reason for the approximation sign is the fact of the Gaussian's tails; but this is a good approximation. We have recovered the Born rule in the GRW theory.

Let us now take a step back from all these gory details and think a little more abstractly about how the various probabilities enter the theory. In essence we have two quite distinct fundamental probabilities postulated; first the probability per unit time for localisation of a system; and second, the probability distribution for the spatial location of the centres of localisation. Both of these are objective chances: knowledge of their values would constrain one's degrees of belief about, respectively, when a localisation might happen and where it will be centred. The first probability is just a postulated brute fact, whose value is not dependent on any particular state of the Universe; the second is postulated too, but here the particular values of the chances depend on facts about the quantum state at a given time. These are the fundamental probabilities; we also have derivative probabilities: the Born rule probabilities for measurements. Facts about the values of these probabilities are determined by the facts about the probabilities for the centres of localisation and (in fine detail) by the shape of the localisation function in the theory. These Born-rule probabilities are objective chances too: knowledge of the GRW theory and of the current quantum state would constrain one's degrees of belief in the outcomes of measurements. We noted, however, that it was simply an assumption in the theory that the probability distribution for centres of localisation should take the form it does; and it was only by making this assumption that the Born rule values were recovered.

The GRW theory is a fundamentally stochastic theory of the Universe as a whole: the basic probabilities it postulates are not explained in terms of anything else. It's just how the laws are that they involve this probabilistic evolution. As with any stochastic theory, what—from a God's eye view—the laws of the theory determine given an initial state, is a set of possible histories of the Universe and a probability measure over them. One history, having a certain weight, is realised; the probability per unit time for localisation affects when the choices between alternative histories are made; and the localisations choose which path through history-space is taken. There will be some histories where embarrassing macroscopic superpositions manage luckily to survive; but these of course will be of extraordinarily low weight.¹⁷

¹⁷For further discussion of aspects of probability in the GRW theory, see Frigg and Hoefer (2007), who in particular are concerned with the question of which of the standard philosoph-

4 de Broglie-Bohm Theory

The pilot-wave theory developed by Bohm (1952) (hereafter ‘the Bohm theory’) following ideas put forward by de Broglie (1927) is a non-local, contextual, deterministic hidden variable theory for quantum mechanics. It takes the other horn of Bell’s proposed dilemma above: the quantum state isn’t everything. In addition to the existence of the Ψ -field, it is postulated that systems possess definite positions $\mathbf{x}_1(t), \mathbf{x}_2(t), \dots, \mathbf{x}_N(t)$ at all times. All the other physical facts then supervene jointly on the Ψ -field and on the particle (sometimes ‘corpuscle’) positions. In particular, in this theory, it is not sufficient for an object like a pointer to be located in some position that there be a suitable hump in the Ψ -field, but the component systems have actually to *be there* too.

In this theory, the main role of the Ψ -field is to guide the motion of the N particles. The wavefunction $\Psi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N, t)$ always evolves unitarily in accord with the Schrödinger equation, but in addition there is a rule for the motion of systems: momenta are defined for each system according to $\mathbf{p}_i = \nabla_i S$, where $S(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N, t)$ is the phase of $\Psi(t)$. Hence a definite spacetime trajectory may be defined for each particle, where this trajectory for an individual system will depend on the many-body wavefunction (hence the explicit non-locality—action at a distance—in the theory).

In this theory it is to be the definite positions for particles which provide the single definite outcome for a measurement process. In a successful measurement, as we know, the wavefunction for the joint object-system and measuring apparatus will have spread out into a superposition of non-overlapping wavepackets on configuration space (cf. state (6)). The determinate values of the positions for the object-system and for the particles making up the pointer will pick out a definite point in configuration space; and the outcome that is observed, or obtains, is the one corresponding to the wavepacket whose support contains this point. However, all the terms in the superposition remain intact and continue to evolve unitarily.

While the Bohm theory is non-local in three-dimensional space, it is local in the $3N$ -dimensional space of particle configurations. If we introduce the notation $\mathbf{X}(t)$ for the point in configuration space picked out by the definite values $\mathbf{x}_1(t), \mathbf{x}_2(t), \dots, \mathbf{x}_N(t)$, then (assuming for simplicity particles of equal mass m) the velocity $\dot{\mathbf{X}}(t)$ of this point will be given by

$$\dot{\mathbf{X}}(t) = \frac{1}{m} \nabla S(\mathbf{X}(t)), \quad (19)$$

where ∇ is the $3N$ -dimensional gradient operator ($\partial/\partial\mathbf{x}_1, \partial/\partial\mathbf{x}_2, \dots, \partial/\partial\mathbf{x}_N$). In other words, the trajectory of the point $\mathbf{X}(t)$ is determined by the phase of the wavefunction at $\mathbf{X}(t)$ alone. This means that once the wavefunction has separated out into non-overlapping packets, within one of which $\mathbf{X}(t)$ is to be found, then from the point of view of the motion of the particles, all the other wavepackets can be ignored, apart from the one $\mathbf{X}(t)$ is sitting in: all the other bits of the Ψ -field will be irrelevant to the particle dynamics. We must then make a distinction between those cases where re-combination of the disjoint wavepackets is a live option, as in an interferometry experiment, for example;

ical analyses of the chance notion might fit best within the GRW setting. As indicated above, I’m perhaps less convinced of the call for the analysis of chance in that philosophical sense.

and where it is not: where dynamical decoherence induced by the environment ensures that the time before re-combination (re-coherence) could occur would be many times the life of the Universe. In the latter case we are justified in replacing, from the point of view of calculations, the true wavefunction of the Universe $\Psi(t)$ with a so-called ‘effective wavefunction’ $\psi(t)$; the component of the true wavefunction which is actually guiding the motion of the particles. In realistic measurement scenarios, we will always have decoherence and will thus move to the effective wavefunction post-measurement; but this is simply a calculational aid: the true, objective, wavefunction retains all its components. (If we liked, we could model the move to the effective wavefunction by a quantum operation, analogously to the GRW theory (in fact the operations involved would be almost identical); but it’s not worth it: this is not part of the fundamental dynamics; it’s not part of the dynamics at all.)

So far we have not seen anything of probability. Shortly we shall. But first it will be useful to see a little of how the Bohm theory was formally developed. Taking the single particle ($N = 1$) case to begin with, substitute the polar form of the wavefunction $\Psi(\mathbf{x}_1, t) = R(\mathbf{x}_1, t) \exp(-iS(\mathbf{x}_1, t)/\hbar)$ into the Schrödinger equation, giving rise to two equations:

$$\frac{\partial S}{\partial t} + \frac{(\nabla_1 S)^2}{2m} + V - \frac{\hbar^2}{2m} \frac{\nabla_1^2 R}{R} = 0; \quad (20)$$

and

$$\frac{\partial R^2}{\partial t} + \nabla_1 \cdot \left(R^2 \frac{\nabla_1 S}{m} \right) = 0, \quad (21)$$

where m is the mass of the particle and V is the classical potential. The analogy of (20) to the classical Hamilton-Jacobi equation¹⁸ motivates the definition of the momentum of the particle as $\mathbf{p}_1 = \nabla_1 S$. With this definition, (21) becomes:

$$\frac{\partial R^2}{\partial t} + \nabla_1 \cdot (R^2 \dot{\mathbf{x}}) = 0, \quad (22)$$

which is precisely the form of a local conservation equation for the quantity R^2 , i.e., the modulus squared of the wavefunction.

Now for probability. Since our theory is deterministic, probabilities will have to enter by way of a probability distribution over initial conditions. Thus we introduce a probability density $p(\mathbf{x}_1, t)$ giving the probability with which our particle will at time t be located at position \mathbf{x} (not merely be *found to be located there on measurement*). This probability density will satisfy a continuity equation under the physical dynamics: if the density in a volume decreases this will be in virtue of a flow through the surface of the region, thus:

$$\frac{\partial p(\mathbf{x}_1, t)}{\partial t} + \nabla_1 \cdot (p(\mathbf{x}_1, t) \dot{\mathbf{x}}_1) = 0. \quad (23)$$

Picturesquely, we could imagine a hypothetical infinite ensemble of particles (a kind of fluid) occupying space with a density given by $p(\mathbf{x}_1, t)$; the amount that the number of particles in a small region dV decreases in a time dt will be given by the density at the surface multiplied by the rate of flow (velocity) through the surface. Where we expect to find particles will depend on what

¹⁸See e.g. Landau and Lifshitz (1976, §47).

their trajectories are, as determined by the velocity field $\dot{\mathbf{x}}_1(t)$; the trajectories drag the probability density around, as it were. Comparing (22) and (23) it is extremely tempting immediately to stipulate that the probability density for positions should be given by the modulus squared of the wavefunction. If we make this stipulation then our theory will be exactly empirically equivalent to standard quantum theory: the probability that a particle will be found at \mathbf{x} will be given by the mod square of the amplitude at \mathbf{x} , as this is the probability that the particle is *at* \mathbf{x} . The preceding reasoning goes through in exactly the same way for the N -body case; we need merely make suitable replacements of \mathbf{X} for \mathbf{x}_1 throughout. Then notice that we will recover the exact Born rule probabilities for measurements for quantities other than position too. The probability that a measurement will have outcome i will be given by the mod squared weight of the associated wavepacket; as that's just the probability that \mathbf{X} will be within that wavepacket. Notice too, finally, that $p(\mathbf{X}, t) = |\Psi(\mathbf{X}, t)|^2$ will be a dynamically invariant equality, due to the continuity equations (22) and (23); if the probability density is given by the mod square of the amplitude at one time, then it will be for all other times, if it evolves solely under the physical dynamics¹⁹. The reason for the caveat is that once the results of a measurement are known, the probability density will change discontinuously—the chance that the systems involved are located in the support of any of the disjoint wavepackets but one will drop to zero. In this case, the new probability density will still be given by the modulus squared of a wavefunction—but now it will be the wavefunction for that packet of the Ψ -field which is currently driving the evolution of the systems' positions; following measurement, that is, it will be the newly arrived at effective wavefunction.

Now despite all these welcome consequences of the choice $p(\mathbf{X}, t) = |\Psi(\mathbf{X}, t)|^2$, there is nonetheless reason to pause. As many have noted, the dual role of $\Psi(\mathbf{X}, t)$ as both guiding systems and providing probabilities for their location are logically distinct. *Prima facie* there seems to be no reason at all why the probabilities *should* be given by the mod square of the wavefunction—the latter has, fundamentally, nothing to do with probabilities: it describes the guiding field. Suspicion increases when we recognise that the Bohm theory is empirically equivalent to quantum mechanics not only *if*, but *only if* the probabilities are given by $|\Psi(\mathbf{X}, t)|^2$; move away from this distribution and many strange things start to happen. Yet given that it is the particle positions and the guiding Ψ -field which are truly fundamental in the Bohm theory, why should we in fact happen to see only phenomena which are consistent with the normal quantum predictions? The theory itself would allow much else. This smacks of conspiracy; or of an *ad hoc* theory which is parasitic on normal quantum mechanics, rather than being a genuine alternative in its own right.

Bohm (1952) himself explicitly countenanced the possibility that situations could arise in which the probabilities would differ from $|\Psi(\mathbf{X}, t)|^2$ and thus we would expect predictions differing from those of quantum theory; in particular, violation of the position-momentum uncertainty principle would become possible: one could reliably prepare systems for which the spread of the statistics for measurements of canonically conjugate variables would be less than the familiar quantum bounds. In essence, one begins to see something of the actual locations

¹⁹In fact a stronger claim holds true also: $|\Psi|^2$ is the *unique* dynamically invariant probability distribution (Goldstein and Struyve, 2007).

of particles; their configuration is no longer entirely hidden within the bounds of the wavepacket in whose support they are contained. Thus Bohm remarked that

...if the theory is generalized...The probability density of particles will cease to equal $|\Psi|^2$. Thus experiments would become conceivable that distinguish between $|\Psi|^2$ and this probability; and in this way we could obtain an experimental proof that the normal interpretation, which gives $|\Psi|^2$ *only* a probability interpretation, must be inadequate. (Bohm, 1952, I §9)

However he then went on to suggest that arguments could be given that $p(\mathbf{X}, t)$ could be expected to tend to $|\Psi(\mathbf{X}, t)|^2$ as a kind of equilibrium distribution.

This thought has been pursued in detail by Valentini (1991a) who has demonstrated that $p(\mathbf{X}, t) = |\Psi(\mathbf{X}, t)|^2$, can indeed be derived as the ‘quantum equilibrium’ distribution towards which systems will tend, as the result of a *subquantum H-theorem*. In addition he proved that the impossibility of superluminal signalling via measurement and the uncertainty principle hold in general only in quantum equilibrium (Valentini, 1991b). Thus these crucial quantum mechanical features in fact arise as effective features of an underlying nonlocal and deterministic theory, rather than being postulated as fundamental parts of the theory. Moreover, in this respect the Bohm theory is characteristic of deterministic hidden variable theories in general: no-signalling will hold only in the equilibrium state (Valentini, 2002a), as will the impossibility of experimentally distinguishing non-orthogonal quantum states (Valentini, 2002b); while out of equilibrium the important assumption (as in any C^* -algebraic theory) that expectation values will be additive (probabilistic states be linear in their arguments) also fails (Valentini (2004); Timpson (2004, §9.2.2)). With so much characteristic quantum structure hanging on the probabilities taking the equilibrium value, we can see why the question of why the probabilities might naturally take that value is such an important one.

Here’s a sketch of Valentini’s subquantum H -theorem. The Kullback-Leibler relative entropy $H(p(x)||q(x))$ is a nice measure of the distinctness of two probability distributions (Kullback, 1959, chpts. 1 and 2):

$$H(p(x)||q(x)) = \int dx p(x) \log \frac{p(x)}{q(x)}; \quad (24)$$

it is equal to zero only when the two distributions are identical. Let us then define our subquantum H as

$$H(t) = \int d^{3N} p(\mathbf{X}, t) \log \frac{p(\mathbf{X}, t)}{|\Psi(\mathbf{X}, t)|^2}. \quad (25)$$

Now clearly, if $p(\mathbf{X})$ starts off different from $|\Psi(\mathbf{X})|^2$, then given their respective continuity equations (eqns. (23) and (22)) above, they will remain equally different over time; and H will not change. Similarly to the case in (Gibbsian) classical statistical mechanics, then, we need to introduce coarse-grainings of p and $|\Psi|^2$; and then again, much as in the classical case, it will follow that the corresponding coarse-grained \bar{H} will decrease over time:

$$\frac{d\bar{H}}{dt} \leq 0; \quad (26)$$

at the minimum of \overline{H} , the coarse grained $\overline{p(\mathbf{X}, t)} = \overline{|\Psi(\mathbf{X}, t)|^2}$. The intuitive reason this happens, as Valentini remarks, is that given their respective continuity equations, p and $|\Psi|^2$ behave like two fluids, but being stirred together by the same stick—the configuration point $\mathbf{X}(t)$. When the trajectory is sufficiently complicated, as it will be in realistic dynamical systems, the two pretend fluids will be thoroughly mixed together and distinguishable only on a very fine-grained level. For some instructive results from simulations, which also give some indication of times for approach to equilibrium, see Valentini and Westman (2005).

Once we have the Born-rule probabilities for the total system, Born-rule probabilities for sub-systems follow immediately. The probabilities for measurements on system i will just be given by the reduced state of system i , taking the trace over all the other systems. When, following a measurement or preparation procedure, i is in an effective pure state, the probabilities for its location will then simply be given by the mod square of the corresponding wavefunction.

The subquantum H -theorem is a very nice result. Conceptually, however, it shares some of the standard difficulties that attend the story of the approach to equilibrium in classical statistical mechanics. What justifies coarse-graining? Well one can appeal to clumsiness of our measuring apparatuses; their inability to be responsive to features on a very small scale. This kind of response often fails to carry conviction in the classical statistical mechanics case (why should the increase of entropy of the Universe care about my inability to determine positions very finely?), but perhaps it is a little more plausible in the quantum case. In the Bohm theory any account of measurement will need to involve the wavefunction too; and any kind of realistic measurement (as in the schemas given above) will involve wavepackets of finite extent in configuration space: you get result i of measurement (which need not be a *position* measurement, recall) when \mathbf{X} is found in the support of the i th wavepacket; finer grain than that is not relevant to the outcome of the measurement—and it was for these outcomes that we wanted Born rule probabilities. Admittedly, the detailed picture of measurement becomes somewhat less clear out of equilibrium, but what is clear is that in Bohmian theory of measurement there is a further dynamical player—the Ψ -field—which will naturally structure into discrete lumps during measurement. The prospects for justifying coarse-graining in this direction deserve further analysis.

The second main, familiar, problem is that it can't be true that just *any* old initial probability distribution will tend towards the quantum equilibrium: the theory is deterministic and time-symmetric, after all. For every solution to the equations which would have a non-equilibrium distribution heading towards equilibrium, the time-reverse solution in which the distribution starts closer to equilibrium and heads away from it also exists. What Valentini suggests is crucial, then, is that there be no fine-grained microstructure in the initial distribution and wavefunction; i.e., that at $t = 0$ the coarse grained and fine grained p (respectively, $|\Psi|^2$) are equal (this is explicit in his proof); it is this which is violated in the time-reversed states.

Even with this concession that the result is not entirely universal, holding for *any* set of initial conditions of the Universe, we still seem to have made progress; the explanatory burden regarding 'Why $|\Psi|^2$?' has been considerably lessened. Rather than having to insist that the Universe had to start with this particular single distribution to explain the observed phenomena, we can say

that it could have started in any one of a number of distributions and we would still get what we see today. We can't say that *any* initial condition would do; but we're doing far better than if we had to insist on one and one alone.

It seems, then, that a plausible story can be told in the Bohmian context of how Born-rule probabilities arise dynamically from a preceding non-equilibrium distribution. Let us now compare with probabilities in the GRW theory. In both theories, Born-rule probabilities are derivative rather than fundamental, although in different ways in the two cases: in GRW they derive from the localisation probabilities; in Bohm they arise as an equilibrium probability distribution. In both cases, the fundamental role of probability is to do with positions: a distribution over particle locations in the Bohm theory, a distribution over localisation centres in GRW. From a God's eye view, the Bohm theory will again define a set of possible histories, which in some respects will look very similar to the GRW space of histories. A given initial quantum state of the Universe not a total energy eigenstate will typically evolve into a branching structure; in the Bohm theory a path through this branching will be picked out by the particles' continuous trajectory. In GRW, the path through the branching will be picked out by the discrete localisation events. The other difference is that branches of the state not chosen in GRW typically get killed off, while in Bohm, they survive, merely having no further effect on particle trajectories.

What we tend to think of as the major difference between probability in GRW and in Bohm is that in the former the probability measure over histories is given by a stochastic dynamics, while in the latter it is given by a distribution over the initial particle configurations. We tend to find the latter more puzzling than the former. But how much of a difference really is this? Granted, on one hand we have a significant difference: in the GRW theory, the complete physical state at one particular time (plus the dynamics) does not determine the complete physical state at all others; while in the Bohm theory it does. But from the point of view of the probabilities? In the Bohm theory we have one brute objective chance: the chance that the initial configuration of particle positions should be as they were. In GRW we have many, many chance events: the brute timings of localisations and the state-dependent chances for localisation centre. Is it any more (or less) puzzling that there should be chance in the one theory than the other?

No doubt our disquiet regarding a Bohmian probability distribution over initial conditions of the Universe stems in part from the thought that we could only make sense of such a distribution if we had an infinite ensemble of Universes to play with; and the distribution was telling us something about the number of Universes within the ensemble having a given configuration. And this is a hopeless conception (not only for its fictiveness). But invoking the ensemble in the first place only stems from thinking in frequentist terms; from thinking that probability has to be made out in that way. And that was never going to work in any case, for well known reasons. Yes, we must drop the conception of an ensemble of Universes, but drop frequentist pre-conceptions too (as we must anyway) and that proves no problem for the intelligibility of our probability distribution over initial conditions. A second—and perhaps related—source of disquiet arises in the thought that the probability distribution in question must really arise as a distribution over the possible outcomes of some chancey process. But the initial state of the Universe is just that—initial, the first one—it cannot be the result of some process taking place in time; and invoking instead

a chancey process occurring *outside* time seems just to descend into gibberish. Thus we seem to reach a paradox: if the first event was chancey then it wasn't the first event. But this paradox may be circumvented when we note that not all chances need be chances for the outcome of some chancey process. It is quite intelligible to ask about the chance of whether something simply was so.

5 Everett

The Everett approach (Everett, 1957) rejects Bell's dilemma. We need take neither horn to deal with the trouble about macrosuperposition, for really there is no trouble at all. The Everett view would insist that one should read (6) literally: If one component of such a superposition would, on its own, count as representing a well-localised pointer, then the superposition as a whole *does* represent differently localised pointers. If a suitably Gaussian hump of Ψ -field would constitute a well-localised pointer (cat, table, whatever), then a number of distinct such humps constitute a number of distinct, differently localised pointers (cats, etc.). There is indeed a plurality of outcomes of a measurement. The reason we don't observe it is that we too are part of the total physical system. Our observing goings-on would be just another entangling operation, just like the initial measurement; once we have interacted with the measuring device then we too form part of the plurality; we branch, as it were; and we don't see what's going on in the other branches. Correlated with each of the different well-localised pointers will now be differently disposed observers, with different internal configurations (brain states and all the rest of it), corresponding to having observed different things. So the existence of macrosuperposition is perfectly consistent with our experience, if we take care to think about what would physically be involved in observing things; or so the claim.

In modern versions of the theory (Saunders, 1995, 1996b, 1998; Wallace, 2002, 2003a, 2010a) the crucial role of decoherence in all this is emphasised. Decoherence theory (Joos et al., 2003; Zurek, 2003) tells us that a natural consequence of quantum dynamics is the emergence at a coarse-grained level of a branching structure of effectively autonomous components of the quantum state which individually will evolve according to quasi-classical laws. One path through this branching structure will represent a quasi-classical history. Observers able to monitor and interact with their surroundings emerge as inhabitants within such independent quasi-classical histories; corresponding to particular complex arrangements of fundamental particles; evolved to take advantage of the existence of enduring records that decoherence allows (Gell-Mann and Hartle, 1990; Saunders, 1993). Often branching is described in terms of a branching into different concrete worlds. The independence of the evolution of the components is crucial: it is this which justifies 'separate world' talk and which ensures that realistic observers won't detect that in fact they are part of a global superposition.

This is all well and good, but we clearly now face a major problem with probability. Following a measurement procedure we have a branching of the total state into components, each of which represents one of the outcomes. But all of these in fact obtain. We don't get one single outcome of measurement, we get them *all*. And what room could there be for a notion of probability here? If we don't get one outcome of some chance event occurring at the expense of

others, then how is probability to get any grip at all? This has been recognised as a fundamental—perhaps crippling—problem for the theory for a long time.

This is the pure form of the problem: that probability is rendered unintelligible if all of the outcomes the probabilities are supposed to pertain to in fact occur. There are other forms the worry has taken too, though, for instance transmuting into the claim that it is impossible to see how the Born rule probabilities in particular could be recovered within the theory.

Everett himself saw no problem; for him it was good enough that one could define an appropriate normed additive measure over branches:

The situation here is fully analogous to that of classical statistical mechanics, where one puts a measure on trajectories of systems in the phase space by placing a measure on phase space itself, and then making assertions which hold for “almost all” trajectories. (Everett, 1957, §5)

He also suggested that the Born rule measure could be derived. Requiring a positive function m only of the coefficients α_i in a post-branching superposition, he suggested the measure should be independent of the phase²⁰. With the assumption too of additivity, he showed that $m(\alpha_i) = |\alpha_i|^2$; i.e., we have the Born rule measure.

But to what purpose? The whole point is that the situation in Everett is precisely *not* like that in classical statistical mechanics: in statistical mechanics, one trajectory is realised; in Everett, they all are. Of course one can define a measure over the branches and of course the Born rule will give an example of such a measure, but the question is, what has that measure to do with probability? The challenge is that we have overwhelming reason to say: nothing at all!

The point is often emphasised by highlighting the existence of branches with deviant statistics. Imagine a long run of repeated spin- z measurements on spin-1/2 systems identically prepared in a spin- x eigenstate. The Born-rule probabilities for the outcomes are 50/50 spin-up/spin-down. But in the Everett interpretation, all the outcomes are realised, which means that as well as branches in which the observed relative frequencies of outcomes are close to the Born-rule values, there will also exist many branches with radically deviant statistics; including branches in which the outcomes were all spin-up and branches in which they were all spin-down, even in the infinite limit. In a usual probabilistic setting we will say: such statistics are unlikely, very unlikely, or they are certain not to occur: the possibility of their occurrence may be discounted. But not so in Everett, it seems, for there, by contrast, they are *certain to occur*. Granted, one can prove law of large number-type theorems that say that branches with deviant statistics will get low or vanishing weight as measured by the Born-rule measure²¹, but this tells us nothing until we can connect that measure with *probability*. How can low or vanishing Born-rule weight tell us that a possible sequence of results can be discounted when we know that it is

²⁰His argument here is obscure: it seems to advert to the impossibility of determining the phase of a quantum state—but this is only the overall phase; and relative phases are still significant and can be determined.

²¹Everett (1957, §5) appealed to this sort of idea. Various formal proofs of quantum versions of laws of large numbers were given by Finkelstein (1963); Hartle (1968); Ochs (1977); Fahri et al. (1989). See Caves and Schack (2005) for discussion of these proofs.

in fact certain to occur? Why should we expect to see Born-rule statistics, when branches with non-Born rule statistics seem to be of exactly the same status?

A related challenge is due to Graham (1973):

It is extremely difficult to see what significance [Everett's] measure can have when its implications are completely contradicted by a simple count of the worlds involved, worlds that Everett's own work assures us must all be on the same footing. (Graham, 1973, p.236)

Faced with a binary measurement we will have a branching into two distinct components; two different worlds. Why shouldn't the probabilities be 50/50, then, irrespective of the Born-rule value? Why shouldn't the amount that we expect to see given statistics just be given by counting the number of worlds in which these statistics will be instantiated? And those numbers will be given by a standard permutation argument and have nothing to do with Born-rule weights. The number of branches having a relative frequency f of 'up' outcomes in a sequence of N independent repeated trials of the spin measurement will simply be

$$\frac{N!}{(N(1-f))!(Nf)!}$$

which has nothing to do with amplitudes of the state at all.

A number of authors (Deutsch, 1985; Lockwood, 1989; Albert and Loewer, 1988) sought to respond to this counting argument by introducing a further plurality of elements into the Everettian setting, whether of further worlds, or points of view, or minds. The idea, roughly, is that by adding these extra elements, one can make up the numbers to the Born-rule values. These are theories of many-many worlds, or of many minds; the basic ontology is now to contain much more than the Ψ -field. In outline we can think of the branching structure of the state indicating the *types* of world that are available at a given time, while the further elements indicate *how many* of each type is instantiated. Thus following a binary branching where the branches have weights 2/3, 1/3 respectively, say, one might insist that one branch receives twice as many of these further elements (worlds or minds) as the other one; so now a straight count of the possibilities will give us back the Born-rule values. Clearly there will need to be a continuous infinity of these elements to give us a real-valued measure.

This was not a salutary development. First, looked at individually, each one of these elements corresponds effectively to a hidden variable. It picks out one of the histories as the preferred, or instantiated one. But then the elements are introduced *en masse*, with numbers in accord with Born-rule weights. We have an uncountable swarm of hidden variables. But what's the point of this? If we are going to introduce hidden variables at all then we might as well introduce one alone: introducing more isn't going to make things better; and the extra ones are superfluous. With the introduction of the singular hidden variable, the probability problem is already immediately transformed: we no longer have the distinctive Everettian problem, as it's no longer the case that all of the branches have the same status; we are back essentially to statistical mechanics-type probabilities. The second point is that introducing the swarm only backs the probability problem up one level²². Why should the proportion of the swarm

²²The case is exactly like introducing an ensemble of Universes when we were puzzling over the Bohmian distribution over initial conditions.

each branch has associated with it determine what I should expect to see on measurement? This would only seem to make sense if there were some separately given chance that I should myself be found in one of the worlds associated with a given swarm-member. But where does this probability come from and what sense could be made of it? And if one needs this further probability anyway, what use the swarm?

Thankfully, Saunders (1996a, 1998) cut neatly through this tangle, thereby instigating the first crucial movement towards the *decision-theoretic turn* which has transformed the probability debate in Everett.

The first point is to note a tacit premiss in the counting argument; it is simply being assumed that each of the worlds is equally likely. In the counting argument,

[the] key objection is that a simple count of worlds contradicts the quantum mechanical probabilities. But a simple count of worlds cannot have anything to do with probability *unless we presuppose that the worlds are all equiprobable*. Whence, now, this *a priori* equiprobability? (Saunders, 1996a, p.136)

Thus it is assumed in this objection that the notion of probability does make sense in an Everettian setting but an indifference-type argument is given to supply the values of the probabilities. But one can learn by experiment that the values of probabilities a principle of indifference would supply should be corrected; principles of indifference are always trumped by the data: I can learn that a die is biased. Similarly (granted that the notion of probability makes sense) we can learn that the Born-rule probabilities are in fact the right ones in a quantum mechanical setting, rather than weighting each branch equally. Once we allow that probability makes sense, we can simply note that the Born-rule probabilities are the empirically correct ones. In the theory they can simply be postulated as correct. It's as meaningful to say that the outcomes should have Born-rule weights as to say that they should have equal weights. If one is truly puzzled by probability in Everett, the counting argument already concedes too much, for it concedes that the chance of outcomes makes sense, quibbling only with the values given. From that position it is easy to get to the Born rule. Simply postulating that the probabilities are Born-rule probabilities would be no worse than stipulations of probability in GRW, Bohm or classical statistical mechanics.

The counting argument presupposed that probability makes sense in Everett. The second point, then, is to argue that space for probability does indeed exist, even if all the outcomes of a measurement (a branching) occur. The crucial move of Saunders (1998) was to point out that from the point of view of a decision-making agent faced with a branching event, the situation will be first-person identical to facing a genuinely indeterministic branching in which only one event occurs. This is a *subjective uncertainty* reading of a branching event. From my point of view, coming up to a branching, I will be uncertain about what I will see post measurement; I am uncertain about which of the outcomes I will have to deal with. Even though all of the outcomes will obtain, I can expect to see only one; for there is no such thing as the first-person expectation of a conjunctive experience of all the outcomes; and I shall surely see *something*. With respect to decision-making and dealing with the future I will be in exactly the same state of mind as I would be if I were facing a genuinely uncertain

event. Thus just as when faced with a genuinely indeterministic event, I can have degrees of belief about what will obtain; about what I will see after the event. And now we have room for probability.²³

At this point one might feel enough has been done. We have personal probabilities and the Born-rule weights could simply be postulated as the objective chances which conditionally constrain personal probabilities. The physical facts which determined the chances would be amplitudes of the Ψ -field. Remarkably, however, it seems that one can go further. As first argued by Deutsch (1999), it appears that it can be *proved* that one's personal probabilities ought to accord with the Born rule; proved, that is, that the Born rule values *must be* the unique objective chances. This is an extraordinary result if it can be maintained. Previously we have had to stipulate what the objective chances are in a theory. Now they are forced on us; and their ineluctability as rational constrainers of degrees of belief laid bare.

5.1 The Deutsch-Wallace Argument

Deutsch's strategy makes essential use of a decision-theoretic framework for dealing with probability. The most basic way of thinking about degrees of belief (cf. Ramsey (1926)), the entry-level technique for operationalising or defining this term of art, is by means of (idealised) betting behaviour. The degree of my belief that an outcome should occur, my personal probability for the event, can be determined by how I would be prepared to bet, were I a betting man. There is an equation between what I take fair odds to be and what I deem the probabilities. The latter can be inferred from the former. Evens means that I judge the outcome as likely to occur as not, i.e., it has probability of 1/2; two-to-one on means I think there's a 2/3 probability that the outcome will obtain; and so on. Fair odds are the odds at which I think neither I nor the bookie can be expected to lose; should I be a betting man, therefore, I would be indifferent between taking the bet and keeping my stake; the expected net value of the bet (according to me) is zero.

A more sophisticated development of the same idea (allowing, for example, that I might care about more than just money...) is provided by decision theory (e.g. Savage (1954)). In decision theory, we conceive of an agent confronted with a number of choices to make, or acts to perform, each of which has a range of possible consequences (even nothing of interest happening would count as a consequence of one sort). While not invoking any numerical values, we can still ask which of these choices the agent would prefer over others, in light of what the respective consequences might be. We can thereby hope to come to an ordering—a preference ordering—amongst the possible choices. If these preferences satisfy a suitable set of axioms of rationality (e.g., a ground-level one being that preferences should be transitive), and if the structure of choices is rich enough, then it is possible to prove a representation theorem which says that one choice is preferred to another if and only if a certain number attached to the first is greater than a certain number attached to the second, where these

²³An alternative view on how uncertainty may enter in a deterministic Everettian setting is given by Vaidman (1998, 2002).

numbers are given, for a choice labelled c , by an expression:

$$EU(c) = \sum_i p_c(i)U_c(i), \quad (27)$$

where the $p_c(i)$ will be numbers satisfying the probability axioms and are therefore interpreted as the subjective probabilities which the agent thinks the consequences i of the choice would have, and U_c is a real-valued function taken to represent the utilities attached by the agent to the consequences of the choice. $EU(c)$ is termed the expected utility of the choice c . Importantly, with such a representation theorem, the (qualitative) preference ordering fixes the (quantitative) probabilities effectively uniquely.

Deutsch (1999) transferred this *modus operandi* to the Everettian setting and considered an agent facing choices amongst various different measurements that might be performed, where the outcomes of the measurements each have some utility attached. More specifically, the scenario is one of an agent considering a number of *quantum games* that could be played, where the game consists of the preparation of a system in some state, a measurement on that system; and the payment of a reward (positive or negative) to the agent, consequent on the outcome observed. The agent's preference order amongst the games can be expressed by a value function (expected utility) \mathcal{V} on games; and from decision theory axioms, the usual probability representation theorem follows: the value \mathcal{V} assigned to a given game can be expressed (uniquely) as the weighted sum of the personal probabilities for outcomes of the measurement involved, multiplied by the rewards assigned to the outcomes. Then comes the punch line. These personal probabilities, however, have (provably) to be the Born rule values. That is, the value of a quantum game to a rational agent is given by the sum of the Born-rule weights multiplied by the rewards for outcomes. Result.

Deutsch's original argument did not immediately persuade, however (see Barnum et al. (2000)), not least because it was unclear that the Everett interpretation was assumed in his proof. The next important step was then taken by Wallace (2003b) who clarified the logic behind Deutsch's argument, highlighted a crucial suppressed premiss and provided a simplified version of the proof. Wallace (2007) clarified the argument yet further and provided an improved alternative version of the proof from weaker assumptions²⁴.

Following Wallace's (2003) clarification of Deutsch's original argument, we represent a quantum game by an ordered triple $G = \langle |\psi\rangle, \hat{X}, \mathcal{P} \rangle$, where $|\psi\rangle$ is the state the system is prepared in, \hat{X} is a self-adjoint operator representing the observable to be measured (assume a discrete spectrum $\sigma(\hat{X})$) and \mathcal{P} is a function from the spectrum of \hat{X} into the reals, representing the different rewards or *pay-offs* the agent will receive on the occurrence of a particular outcome of the measurement. Different outcomes might receive the same pay-

²⁴Closely related proofs may be found in Saunders (2005) (which also includes helpful discussion of the general philosophical issues about probability in quantum mechanics) and Zurek (2005). Greaves (2004) provides an interesting alternative perspective which rejects the subjective uncertainty starting point for probability in Everett, but argues that the decision-theoretic argument for Born rule probabilities may nonetheless go through. See Wallace (2006) for discussion of this alternative and more on the philosophy behind the general argument. Greaves (2009) and Greaves and Myrvold (2010) take a different tack and resolve the question of how, in a theory in which all possible outcomes of branching occur, specific probabilistic claims such as the Everett interpretation makes could ever be supported.

off, so a given G will factor the Hilbert space of the system into orthogonal subspaces each associated with a different value of the pay-off.

Our agent is then to consider their preference amongst quantum games by establishing a suitable ordering amongst the various triples G . It is crucial to recognise that the relation between the representations G of games and physical processes in the lab is many-many. There are various different ways in which one can construct a measuring apparatus corresponding to a given operator; and there is more than one triple G which will correspond to a given physical process (Wallace, 2003b). For instance, it is straightforward that a device apt to measure an observable \hat{X} will equally measure $f(\hat{X})$ where f is a function that maps the spectrum of \hat{X} to another set of real numbers. There is no physical difference between measuring the two: this is simply a matter of re-labelling what eigenvalues we should take a given measurement indicator to indicate. (Look back at the process leading to (6); the only important thing there were the *eigenvectors*, not the eigenvalues.) Therefore, a physical process which would instantiate $\langle |\psi\rangle, \hat{X}, \mathcal{P} \cdot f \rangle$ would also instantiate $\langle |\psi\rangle, f(\hat{X}), \mathcal{P} \rangle$; the very same pay-offs would be had for the same physical outcomes; the only difference is an arbitrary labelling. (This is called *pay-off equivalence*.)

Another important example arises when we note that it is an arbitrary matter where one takes state preparation to end and measurement to begin. For example, consider a process starting with an electron in an ‘up’ eigenstate of spin in the x -direction which is then rotated $\pi/2$ about the y -axis into an ‘up’ z -spin eigenstate by the application of a suitable homogeneous magnetic field, before finally entering a Stern-Gerlach apparatus oriented in the z -direction. Is this a measurement of the x -component of spin of an x -spin eigenstate, or a measurement of the z component of spin of a z -spin eigenstate? Answer: whichever you please. It’s a matter of convention. More generally, a process which instantiates a preparation $U|\psi\rangle$ (where U is a unitary operator on the Hilbert space of the system) followed by a measurement corresponding to $U\hat{X}U^\dagger$ will also count as instantiating a preparation of $|\psi\rangle$ followed by a measurement corresponding to \hat{X} . In other words, both $\langle |\psi\rangle, \hat{X}, \mathcal{P} \rangle$ and $\langle U|\psi\rangle, U\hat{X}U^\dagger, \mathcal{P} \rangle$ would be instantiated by the very same process. (This is called *measurement equivalence*; Wallace (2003b, pp.422-423) proves a more general case.)

Given this freedom to attach different representations G of games to the physical processes an agent faces, we can usefully define an equivalence relation \simeq on triples G :

- G and G' are equivalent, $G \simeq G'$, iff some particular physical process would count as an instantiation of G and also of G' .

Wallace (2003b) then goes on to show via application of pay-off equivalence and measurement equivalence that two triples G and G' are equivalent in this sense if and only if they share the same set of pay-offs (the range of the pay-off function is the same in each case) and the subspaces (generally distinct) which each associates with a given pay-off have identical weights as determined by the Born rule. This is *general equivalence*. If P_c^G is the projector onto the subspace which G associates with pay-off value c , then the weight $W_G(c)$ which G assigns to that pay-off is:

$$W_G(c) = \langle \psi | P_c^G | \psi \rangle. \quad (28)$$

General equivalence then states that:

$$G \simeq G' \text{ iff } W_G = W_{G'}. \quad (29)$$

General equivalence is a claim that links the possibility that two mathematical triples G and G' would be instantiated by one and the same process to the mathematical weights they respectively assign to subspaces associated with particular groupings of measurement outcomes. So far this says nothing at all about what rational agents ought to believe; about how they ought to arrange their preferences. That comes next.

Our agent's preference order is to be on triples G . If a given physical process would equally well instantiate G and G' then we must be indifferent between the two, for there is only one physical situation to consider: there can be no grounds for a difference. That is, if $G \simeq G'$ then $\mathcal{V}(G) = \mathcal{V}(G')$; their values to the agent must be the same (Wallace (2003b) calls this *physicality*). But then note that G is equivalent to G' iff their pay-offs receive equal Born-rule weight (from (29)). Putting these together we have the result that we must be indifferent between G and G' if their pay-offs receive the same Born-rule weight. The crucial move has now been made, connecting equal weight with indifference. From here it is a matter of turning the mathematical handles (though see Wallace (2003b, §5) for the details); but we have seen the first step in establishing that the Born-rule weights—the projections of the state onto the subspaces associated with given pay-offs—are all that are relevant in fixing the preference order.

In this argument it is crucial that the preference order is on the triples G . If I am indifferent (etc.) between G and G' then I am indifferent between *them*, irrespective of what physical process they might be instantiated by on a given occasion. This is the suppressed premiss Wallace identifies; without it the proof breaks down. For example:

Consider operators A , B and C on the Hilbert space of the system, where A and B commute, A and C commute, but B and C don't. Take B and C to be non-degenerate, with a complete set of eigenvectors $\{|\eta_i\rangle\}$ for B and $\{|\xi_i\rangle\}$ for C , where elements from one set are non-orthogonal to those from the other. A will be degenerate; its eigen-subspaces corresponding to groupings of the one-dimensional eigen-subspaces of B and C . It is possible to measure A either by performing a measurement correlating eigenstates $|\eta_i\rangle$ of B with the apparatus pointer states and then coarse-graining the outcome (call this process 1), or by correlating eigenstates $|\xi_i\rangle$ of C with the apparatus pointer states and coarse-graining the outcome (call this process 2). We can now consider three triples:

$$G = \langle |\psi\rangle, A, \mathcal{P} \rangle, G' = \langle |\psi\rangle, B, \mathcal{P}' \rangle, \text{ and } G'' = \langle |\psi\rangle, C, \mathcal{P}'' \rangle.$$

Assume (without loss of generality) that the pay-off function for G , namely \mathcal{P} , is just the identity function: the pay-off associated with an outcome of the measurement is just the corresponding eigenvalue of A . Then take \mathcal{P}' and \mathcal{P}'' to be corresponding coarse-grained pay-off functions for their respective measurements: for any outcome corresponding to an eigenstate which falls in a given eigen-subspace of A , the corresponding A eigenvalue pay-off is received. Now all of G , G' and G'' will be equivalent to one another. G and G' are clearly equivalent: process 1 is a single process which when followed by the pay-off would instantiate either. Similarly G and G'' are clearly equivalent: process 2

followed by the pay-off would be a single process which would instantiate either. But also G' and G'' are equivalent to one another: either of processes 1 or 2 followed by the pay-off would count as an instantiation of either²⁵. We can infer that the three are all equivalent by appeal also to the transitivity of equivalence relations: if G and G' are obviously equivalent (reflecting on process 1) and G and G'' are obviously equivalent (reflecting on process 2), then it follows that G' and G'' are also equivalent and moreover I should be indifferent between them. But if one thought that it made a difference which way A was measured, then this reasoning collapses. If I thought that it was important which of processes 1 or 2 was used to play the quantum game then on the basis of the palpable equivalences $G \simeq G'$ and $G \simeq G''$ I might not automatically infer that $G' \simeq G''$; in fact the equivalence relation would begin to break down. Similarly, while looking at process 1 I might think: yes, I must be indifferent between G and G' ; and while looking at process 2 I might think: yes I must be indifferent between G and G'' ; but if I'm not indifferent between the processes too, then I can't thereby infer that I should be indifferent between G' and G'' . And where would that leave G ? A given item cannot have more than one place in a preference order.

What justifies the setting of the proof in terms of preferences on triples G , then, Wallace argues, is the natural (previously unarticulated) principle of *measurement neutrality*. Measurement neutrality is the claim that an agent is indifferent between different ways in which a measurement may be physically realised, thus indifferent between different processes which might be used to instantiate a game represented by a particular triple G . As we have seen, if this weren't so, then the agent's preference ordering could not be represented on objects $(|\psi\rangle, \hat{X}, \mathcal{P})$; we'd have to take into account the detailed manner the measurement was physically realised; and the proof would not get going.

Now, measurement neutrality is in fact a standard assumption in the quantum theory of measurement; we typically do assume that any of the different ways one might build a device which would count as measuring the observable quantity represented by a given operator are equally good. On reflection, however, measurement neutrality may begin to look rather too much like an assumption of non-contextuality of probabilities (cf. fn.7) which *is* a contested assumption. The specific reason non-contextuality is contested is the thought that there might be more to the physical situation of measurement which would be relevant to the probabilities one might expect than just the quantum state and the specification of an operator associated with the measurement. This will be painfully obvious when one has a hidden variable which helps determine the outcome observed, for example. But arguably these reasons for concern melt away in the Everettian context (cf. Wallace (2003b, §7), Wallace (2007, §7), Saunders (2005, §§5.4–5)). Clearly we don't have any hidden variables to worry about; and unlike in an operationalist theory where measurement is an unanalysed notion which might hide any number of decision-theoretically relevant

²⁵True, if the pay-off functions for G' and G'' were *fine-grained* then there'd be no process which would equally well instantiate either; but that's not the case we're considering here. Equally, there's no process which counts both as a *measurement* for B and for C , but we're interested in the instantiation of *games as a whole*, not measurements. And these will be identified by what physically gets put in and what physically comes out. What come in are systems in various quantum states and what come out are the concrete goods corresponding to pay-offs. Under these conditions of identification, either of process 1 plus pay-off or process 2 plus pay-off do indeed count as co-instantiating G' and G'' .

sins, measurement is a transparent, dynamically analysed process in Everett. Then we should look at that process.

To recap, we want a measurement process to tell us reliably when a system measured was in some eigenstate. To do that we need the process to establish correlations between input eigenstates and pointer states of an apparatus, where the pointer states are suitably macroscopic and stable. The process will naturally involve decoherence and a relatively enduring branching structure will be established. Should I care about different ways of performing a measurement corresponding to the operator A ? Well, what's to care about? First, the particular eigenvalues of A obviously aren't important and aren't involved. If an system in an eigenstate of A is the input and a clear indication of its having had that state is the output, then that's all that's relevant. Second, the exact nature of the measuring device and the pointer states isn't important. If I did care about that then it could in any case just be incorporated into the pay-off function for an outcome and would not affect my personal probabilities. Third, the details of the dynamics do not matter. So long as they are such as to establish the required correlations then that's all that's relevant. Notice in particular that given that I am indifferent regarding the specifics of the pointer states, so too will I be indifferent regarding the exact mapping of input eigenstates to pointer states. Measurement neutrality thus looks a defensible proposition.

Let us reiterate the conclusion: the result of the Deutsch-Wallace argument is that in Everettian quantum theory, an agent's degrees of belief are rationally constrained to take on the values given by the Born rule; thus if objective chances are those things knowledge of which would rationally constrain one's degrees of belief, then the objective chances are the square amplitudes of the wavefunction. Moreover, it has been *proven* that these physical quantities play the objective chance role.

5.2 Alternative proof

Perhaps a quicker way to personal probabilities from preference orderings (Savage, 1954) is first to define preferences amongst consequences from preferences amongst choices (acts). If the ensuing qualitative ordering on consequences then satisfies suitable axioms there should again be a representation theorem which provides unique numerical probabilities for consequences given choices. Again, the individual's preference ordering determines particular personal probabilities.

This is the route Wallace (2007) takes in his improved version of the proof. Consider ordered pairs $\langle E, M \rangle$ where E is the outcome (which may be coarse-grained) of a measurement process M . We will consider preferences amongst the consequences $E|M$: that E should happen given that M was performed. Introduce the ordering relation \succeq amongst consequences: $E|M \succeq F|N$ iff the outcome F given N is not preferred over the outcome E given N . We may also read this as a qualitative personal probability: I will order as $E|M \succeq F|N$ iff I think that it's as least as likely that E will happen given M is performed as that F will happen given that N is performed. I'm then indifferent between two consequences, $E|M \simeq F|N$ iff $E|M \succeq F|N$ and $F|N \succeq E|M$ (N.B., \simeq now has a different definition from the one given in the preceding section.)

Wallace (2007) then argues for *Equivalence*:

- If E and F are events and the Born-rule weight given to E on measurement

of M is equal to the Born-rule weight of F on N , then $E|M \simeq F|N$.

From there, with the addition of a weak and universally agreed assumption about the possibility of preparing and measuring states²⁶ and three intuitive decision-theory axioms²⁷ the *quantum representation theorem* follows:

- The ordering \succeq is uniquely represented by the Born-rule probability measure on outcomes: the quantitative personal probabilities $p(E|M)$ must be given by the Born-rule weights $W_M(E)$.

Equivalence follows from measurement neutrality, but the direct argument proceeds by physically modelling the pay-offs an agent is concerned with and then considering an *erasure* procedure.

Take the simple case of a spin measurement on a spin-1/2 system and consider two games. In the first the agent gets a reward if the outcome is spin-up in the z -direction, in the other if the outcome is spin-down in the z -direction. As this is a decision problem, the agent is only concerned about getting the reward; by definition this is all they put value on. Let $|\text{reward}\rangle, |\text{no reward}\rangle$ be the states corresponding to the having and not having of the reward, respectively. Suppose the initial state $|\psi\rangle$ is $|\uparrow_x\rangle = 1/\sqrt{2}(|\uparrow_z\rangle + |\downarrow_z\rangle)$ then we have two options for the subsequent state, depending on what game was chosen:

$$\text{Game 1 : } |\psi_1\rangle = \frac{1}{\sqrt{2}}(|\uparrow_z\rangle|\text{reward}\rangle + |\downarrow_z\rangle|\text{no reward}\rangle); \quad (30)$$

$$\text{Game 2 : } |\psi_2\rangle = \frac{1}{\sqrt{2}}(|\uparrow_z\rangle|\text{no reward}\rangle + |\downarrow_z\rangle|\text{reward}\rangle). \quad (31)$$

Now given that the agent only cares about having the reward, they will be indifferent if, following the measurement, we were to erase the spin outcome, for example by re-setting it to a standard state $|0\rangle$ with the help of an auxiliary system which was then discarded. Then the states following this erasure process would be:

$$|\psi'_1\rangle = |\psi'_2\rangle = \frac{1}{\sqrt{2}}|0\rangle(|\text{reward}\rangle + |\text{no reward}\rangle). \quad (32)$$

Since one and the same state is arrived at, the agent must be indifferent between Game 1 followed by erasure and Game 2 followed by erasure. As they are indifferent to erasure too, it follows that they are indifferent between Games 1 and 2; indifferent, that is, between these two games having equal Born-rule weights for pay-offs. Notice that if the pay-off branches weren't equally weighted, then one would not arrive at the same state following erasure and so the argument for indifference would fail.

This sketch is then generalised to argue for indifference between cases of equal weights other than of the value of $1/2$ (consider, e.g., spin measurements on a spin-1 system with a state $\sqrt{w}|\uparrow_z\rangle + \sqrt{w}|\downarrow_z\rangle + \sqrt{1-2w}|0_z\rangle$, followed by erasure; with Games 1 and 2 defined as before) and to allow for indifference between

²⁶That for any n there is some system with an n dimensional Hilbert space which may be prepared in any state and at least one non-degenerate measurement be made on it.

²⁷In addition to transitivity (that if x is preferred to y and y to z , then x is to z) only *separation* and *dominance* are added. Separation is the assumption that there is some outcome which is not impossible; and dominance is the idea that an event doesn't become less likely if more outcomes are added to it; in fact it will become more likely unless the added outcomes are not null; that is, certain not to happen themselves.

equally weighted pay-offs arising under different measurements (imagine erasing details of which procedure was followed). A fully formalised and explicit version of the proof is provided in Wallace (2010b).

As Wallace remarks, there is a notable premiss at play in the argument from erasure: *branching indifference*. It could well be that a realistic erasure process might give rise to a branching into a number of distinct decohered branches in each of which I get a particular value of the pay-off. If I am to be indifferent to erasure then I need to be indifferent to the possibility of such a branching. Similarly, Equivalence tells us that it is only the weight of a pay-off that is important, in particular, it is not to the point how many future copies of myself might receive the reward. Thus *branching indifference* is the claim that

- An agent should be indifferent regarding the occurrence of a process whose effect is only to cause the world to branch, with no pay-offs (positive or negative) ensuing for their post-branching descendents.

This idea was already in play with measurement neutrality; and it seems a natural consequence of construing branching as involving subjective uncertainty from the perspective of the agent. I face a branching, with a pay-off associated with each branch. I'm uncertain which pay-off I'll get. Now suppose one of the branches would in fact branch further, before I get the cash (or the slap, whichever it is). I would be just as uncertain as in the absence of this sub-branching about whether I would get the cash: my personal probability for getting the reward would be just as it was before; I'm indifferent between the sub-branches as in each I get the same pay-off and it's the pay-off which I care about. This kind of consideration is further supported (Wallace (2007, §9), Wallace (2010b, §5), Saunders (2005, §5.4)) by noting that branching is a ubiquitous process which happens all the time and it would be wildly beyond the cognitive capacities of any realistic agent to keep track of the detail of these branchings; moreover, branching is only an approximately defined notion, so there is no fact of the matter about what the number of branches actually is. We can't but be (rationally) indifferent to these processes.

6 Assessment: Too good to be true?

In the Deutsch-Wallace result, the philosophical stakes are high:

Subjective and objective probability emerge at the end of the day as seamlessly interjoined: nothing like this was ever achieved in classical physics. Philosophically it is unprecedented; it will be of interest to philosophers even if quantum mechanics turns out to be false, and the Everett interpretation consigned to physical irrelevance; for the philosophical difficulty with probability has always been to find *any* conception of what chances are, in physical terms, that makes sense of the role that they play in our rational lives. (Saunders, 2005, p.236)

In Everett the chances are grounded on the amplitudes of the Ψ -field and the Deutsch-Wallace proof shows why personal probabilities have to take those decreed values; why the mod-squared values are the chances. It is interesting to reflect on why the proof does not transfer to other objectivist views of quantum

mechanics (cf. Wallace (2003b, §7), Wallace (2007, §6)); why we can't get the philosophical benefits there too. The short answer is that the extra components one has to add in these theories spoil the proof. If there's a hidden variable—as in the Bohm theory—then I can't be indifferent to everything except branch weights when weighing up what I expect to see; as the value of the hidden variable will be crucial. Even if the theory decrees that the value of the variable is epistemically inaccessible in detail (as in the Bohm theory, in equilibrium, at least) so its value could not form part of my rational deliberations, still the probability distribution over the possible values is a crucial component which cannot be ignored. That distribution could always trump my equal-weight indifference. Similarly in a collapse theory, the probabilities with which collapses would occur would trump my equal-weight judgements of indifference. It is only in the Everett interpretation, where nothing is added and all the different branches are ontologically on a par—where all the possible outcomes obtain—that the only differentiation between possibilities to which I must not be indifferent are their Born-rule weights. Thus perhaps what seemed to be the Everett interpretation's greatest weakness *vis-a-vis* probability turns out to be its greatest strength.

For some, the Deutsch-Wallace result seems to good to be true. (See for example objections by Lewis (2005), Hemmo and Pitowsky (2007) and Rae (2010).) For the state-of-the art, see the discussions in Saunders et al. (2010), particularly the chapters of Albert and Price; Wallace (2010b) therein anatomises a large range of objections—potential counter-example rules for probability in Everett which would be alternatives to the Born-rule way of choosing—and he identifies exactly which of his axioms each of these conflicts with.

In the broadest terms, concerns perhaps spring from two general directions. On the one hand, one might be concerned that it just seems too much to claim that the Born rule values can be derived *purely* from axioms of rationality; there must be something else going in; and the philosophical punch of the proof may thereby be lessened. On the other, there is a (less articulate) suspicion that probability simply can't be got to work in Everett, so there must be something wrong with the proof somewhere.

Regarding the first, the question of course is what the name of the game is. What kinds of axioms or assumptions are legitimate if one is to show that rational degrees of belief are indeed conditionally constrained? Some of the motivation behind branching indifference, for instance, might look rather more of a practical nature than of a rational one. However it may help here to reflect that rationality is not a pure disembodied schema. Absent perhaps theological discussions, there is no such thing as pure disembodied rationality. Rational assessment is always assessment of (certain of) the activities of embodied creatures like ourselves as more or less rational. The question is not to what extent do our activities partake in a crystallised domain of pure rationality, but rather, given what our aims, interests and abilities are, or sensibly could be, are we acting appropriately in our attempts to achieve those aims; for example by avoiding probably faulty ways of proceeding? If rational assessment is assessment of hum-drum creatures like ourselves, then it's entirely appropriate that the formalisation of the modes of assessment should mention limitations which are essential to the kinds of decision making activities we can engage in. Thus it seems there is scope for principles which involve some suitable aspects of practical considerations; we should not be chary of branching indifference on those

grounds, for example.

Regarding the second concern, the question is where to pick one's battles. If one is prey to a general sense that probability won't work in Everett then I suggest it is a mistake to try and quibble with the formal details of the Deutsch-Wallace proof, or to proffer alternative rules for probability that do not, on the surface at least, seem to be ruled out simply by considerations of rationality. For recall that the Everettian can simply re-trench. Suppose it were shown that the Deutsch-Wallace proof does not ultimately provide what it seems to; that the proof does not demonstrate that rational degrees of belief are conditionally constrained by mod-squared amplitudes. Then the Everettian could just return to accepting the Born rule as a bare posit; and they'd be no worse off than any other physical theory involving probabilities; no worse off, certainly than GRW or the Bohm theory.

No, the place to dig one's heels in—if, that is, one were so inclined—is right at the very beginning, before the wheels of the proof even begin to turn. The decision-theoretic apparatus only gets a grip once we allow that an agent can form a preference ordering in the face of various choices they are offered. Each of these choices involves a branching in the future. Is it so obvious that somebody will be able to form such an ordering in the first place? Faced with the prospect of all of the outcomes of a measurement obtaining, decision making may just conceptually break down. Each choice offered may itself be so imponderable or unintelligible that one couldn't even begin to entertain which one might prefer. The rational agent may simply be left dumbstruck: there is no way to proceed. Of course, if the subjective uncertainty reading of branching can be maintained, then this problem is automatically answered: there is no difference from the point of view of the agent between an Everettian case and a normal case of indeterministic uncertainty. But then the question becomes: are we so sure that the proposal of subjective uncertainty is correct? It looks highly plausible; but perhaps we need to double-check.²⁸

The tough-minded Everettian's response to this kind of ground-floor challenge might be blunt. We are hung up on the question of whether there is space for probability when all the outcomes of a measurement obtain. The sceptic denies that the concept of probability, even personal probability, makes sense in such a scenario. We are not doubting the adequacy or intelligibility of the ontological story about branching, however. But if that's so, then it's perfectly possible that, for all we know, we *do* live in an Everettian universe; and for the entire history of humankind we have been subject to innumerable branchings. If the concept of probability does not get a grip in such a world, then so much the worse for it. We had the wrong concept; we must adopt an error theory about it. In such a world, claims about probability would come out uniformly false. Yet we still managed to get by and make decisions. In truth, the role in our intellectual lives that we thought probability played was being played all along by *robability* (call it); where robability may be made out in the kind of Everettian decision-theoretic way we have been dealing with. And perhaps it would play its role in such a world. But for all the merits of this tough-minded argument, I can't help feeling that robability may end up being just a bit too cheap.

²⁸See Lewis (2007); Saunders and Wallace (2008) for recent discussion.

7 Conclusion

Whether or not the Deutsch-Wallace result is too good to be true it is certainly, as Saunders notes, singular. Let us close with a reverse question: how badly off are GRW and the Bohm theory that they do not have an analogous result? That their chances have (as things stand) to be postulated rather than proven? Is the provision of a proof that the quantities identified as fixing the chances in a theory satisfy the condition of constraining degrees of belief a norm to which any adequate probabilistic theory should aim? It's not clear to me that this is so; which is not to say that it's not an interesting bonus if you happen to be able to get it. Lewis (1994) insisted:

Don't call any alleged feature of reality "chance" unless you've already shown that you have something, knowledge of which could constrain rational credence. (Lewis, 1994, p.484)

But his concern was not ours; he was concerned to assess candidates for the philosophical analysis of the chance notion; opposing a best-systems analysis invoking probabilistic laws to a conception of primitive modal relations between universals, for example. Our question, by contrast, has been: what are the particular circumstances, or which are the particular physical properties, which determine what the probabilities are? These are very different kinds of questions. Lewis' imperative does not hit us.

Acknowledgements

I'd like to thank the editors for their invitation to contribute to this volume and their help in the process. Hilary Greaves and David Wallace were kind enough to comment on the manuscript, for which I'm grateful. Finally thanks are also due to David Wallace for conversations on probability stretching back over quite a few years.

References

- Albert, D. and Loewer, B. (1988). Interpreting the many worlds interpretation. *Synthese*, 77:195–213.
- Allori, V., Goldstein, S., Tumulka, R., and Zanghi, N. (2008). On the common structure of Bohmian mechanics and the Ghirardi-Rimini-Weber theory. *British Journal for the Philosophy of Science*, 59:353–389.
- Barnum, H., Caves, C., Finkelstein, J., Fuchs, C., and Schack, R. (2000). Quantum probability from decision theory? *Proceedings of the Royal Society of London A*, 456:1175–1182.
- Bassi, A. and Ghirardi, G. C. (2003). Dynamical reduction models. *Physics Reports*, 379:257–426.
- Bell, J. S. (1966). On the problem of hidden variables in quantum mechanics. *Reviews of Modern Physics*, 38:447–52. repr. in Bell (1987b, Chpt.1).
- Bell, J. S. (1987a). Are there quantum jumps? In *Schrödinger: Centenary of a Polymath*. Cambridge University Press. repr. in Bell (1987b) pp.201–212.
- Bell, J. S. (1987b). *Speakable and Unsayable in Quantum Mechanics*. Cambridge University Press. Second Edition 2004.
- Bohm, D. (1952). A suggested interpretation of the quantum theory in terms of hidden variables, I and II. *Physical Review*, 85:166–179;180–193.
- Born, M. (1926). On the quantum mechanics of collisions. *Zeitschrift für Physik*, 37:863–867. trans. in Wheeler and Zurek (1983) pp.52–55.
- Bub, J. (1997). *Interpreting the Quantum World*. Cambridge University Press, first paperback (1999) edition.
- Busch, P. (1997). Is the quantum state (an) observable? In Cohen, R. S., Horne, M., and Stachel, J., editors, *Potentiality, Entanglement and Passion-at-a-Distance*, pages 61–70. Kluwer Academic Publishers. arXiv:quant-ph/9604014.
- Busch, P. (2003). Quantum states and generalized observables: A simple proof of Gleason’s theorem. *Physical Review Letters*, 91(12):120403.
- Carnap, R. (1955). *Statistical and inductive probability*. Published by The Galois Institute of Mathematics and Art, 1955. Repr. in Brody, B.A editor *Readings in the Philosophy of Science* Prentice-Hall 1970, pp.440–450.
- Caves, C. M., Fuchs, C. A., Manne, K. K., and Renes, J. M. (2004). Gleason-type derivations of the quantum probability rule for generalized measurements. *Foundations of Physics*, 34(2):193–209.
- Caves, C. M. and Schack, R. (2005). Properties of the frequency operator do not imply the quantum probability postulate. *Annals of Physics*, 315(1):123–146.
- de Broglie, L. (1927). The new dynamics of quanta. Trans. in *Quantum Mechanics at the Crossroads: Reconsidering the 1927 Solvay Conference*, Bacciagaluppi G and Valentini A editors, Cambridge University Press 2009.
- d’Espagnat, B. (1976). *Conceptual Foundations of Quantum Mechanics*. Addison-Wesley, second edition.

- Deutsch, D. (1985). Quantum theory as a universal physical theory. *International Journal of Theoretical Physics*, 24:1–41.
- Deutsch, D. (1999). Quantum theory of probability and decisions. *Proceedings of the Royal Society of London A*, 455:3129–3127. arXiv:quant-ph/09906015.
- Everett, III, H. (1957). “Relative state” formulation of quantum mechanics. *Reviews of Modern Physics*, 29:454–62.
- Fahri, E., Goldstone, J., and Gutman, S. (1989). How probability arises in quantum mechanics. *Annals of Physics*, 192:368–382.
- Fano, U. (1957). Description of states in quantum mechanics by density operator techniques. *Reviews of Modern Physics*, 29(1):74–93.
- Finkelstein, D. (1963). The logic of quantum physics. *Transactions of the New York Academy of Sciences*, 25:621–637.
- Fleming, G. N. (1988). Hyperplane-dependent quantized fields and Lorentz invariance. In *Philosophical Foundations of Quantum Field Theory*, pages 93–115. Oxford University Press.
- Frigg, R. and Hoefer, C. (2007). Probability in GRW theory. *Studies in History and Philosophy of Modern Physics*, 38:371–389.
- Gell-Mann, M. and Hartle, J. (1990). Quantum mechanics in the light of quantum cosmology. In Zurek, W., editor, *Complexity, Entropy and the Physics of Information*, pages 425–458. Addison-Wesley.
- Ghirardi, G. C., Rimini, A., and Weber, T. (1986). Unified dynamics for microscopic and macroscopic systems. *Physical Review D*, 34:470–491.
- Gleason, A. M. (1957). Measures on the closed subspaces of a Hilbert space. *Journal of Mathematics and Mechanics*, 6:885–894.
- Goldstein, S. and Struyve, W. (2007). On the uniqueness of quantum equilibrium in Bohmian mechanics. *Journal of Statistical Physics*, 128:1197–1209.
- Graham, N. (1973). The measurement of relative frequency. In DeWitt, B. and Graham, N., editors, *The Many-Worlds Interpretation of Quantum Mechanics*, pages 229–53. Princeton University Press.
- Greaves, H. (2004). Understanding Deutsch’s probability in a deterministic multiverse. *Studies in History and Philosophy of Modern Physics*, 35:423–456.
- Greaves, H. (2009). The Everettian epistemic problem. *Studies in History and Philosophy of Modern Physics*.
- Greaves, H. and Myrvold, W. (2010). Everett and evidence. in Saunders et al. (2010).
- Hadjisavvas, N. (1981). Properties of mixtures of non-orthogonal states. *Letters in Mathematical Physics*, 5:327–332.
- Hartle, J. B. (1968). Quantum mechanics of individual systems. *American Journal of Physics*, 36:704–712.
- Hemmo, M. and Pitowsky, I. (2007). Quantum probability and many worlds. *Studies in History and Philosophy of Modern Physics*, 38:333–350.

- Hughston, L. P., Jozsa, R., and Wootters, W. K. (1993). A complete classification of quantum ensembles having a given density matrix. *Physics Letters A*, 183:14–18.
- Joos, E., Zeh, H. D., and Keifer, C. (2003). *Decoherence and the appearance of a classical world in quantum theory*. Springer, second edition.
- Kullback, S. (1959). *Information Theory and Statistics*. Dover, Dover 1968 edition.
- Landau, L. and Lifshitz, E. (1976). *Mechanics*. Butterworth-Heinemann, third edition.
- Lewis, D. (1980). A subjectivist’s guide to objective chance. In Jeffrey, R. C., editor, *Studies in Inductive Logic and Probability*, volume II. University of California Press. repr. in David Lewis *Philosophical Papers, Volume II*, Oxford University Press (1986), pp.83–131.
- Lewis, D. (1994). Humean supervenience debugged. *Mind*, 103:473–489.
- Lewis, P. J. (2005). Probability in Everettian quantum mechanics. Pitt Phil Sci Archive 00002716.
- Lewis, P. J. (2007). Uncertainty and probability for branching selves. *Studies in History and Philosophy of Modern Physics*, 38:1–14.
- Lockwood, M. (1989). *Mind, Brain and Quantum*. Blackwell.
- Maudlin, T. (2007). Completeness, supervenience and ontology. *Journal of Physics A*, 40:3151–3171.
- Mellor, D. H. (1969). Chance. *Proceedings of the Aristotelian Society Supplementary*, 43:11–34.
- Mellor, D. H. (1971). *The Matter of Chance*. Cambridge University Press.
- Myrvold, W. (2002). On peaceful coexistence: Is the collapse postulate incompatible with relativity? *Studies in the History and Philosophy of Modern Physics*, 33:435466.
- Myrvold, W. (2003). Relativistic quantum becoming. *British Journal for the Philosophy of Science*, 54:475500.
- Nielsen, M. A. and Chuang, I. (2000). *Quantum Computation and Quantum Information*. Cambridge University Press.
- Ochs, W. (1977). On the strong law of large numbers in quantum probability theory. *Journal of Philosophical Logic*, 6:473–480.
- Pearle, P. (1989). Combining stochastic dynamical state-vector reduction with spontaneous localization. *Physical Review A*, 39:2277–2289.
- Rae, A. I. M. (2010). *Studies in History and Philosophy of Modern Physics*.
- Ramsey, F. P. (1926). Truth and probability. In Kyburg, H. and Smokler, H., editors, *Studies in Subjective Probability*, pages 63–92. Wiley.
- Saunders, S. (1993). Decoherence, relative states, and evolutionary adaptation. *Foundations of Physics*, 23:1553–1585.
- Saunders, S. (1995). Time, quantum mechanics, and decoherence. *Synthese*, 102:235–266.

- Saunders, S. (1996a). Relativism. In Clifton, R., editor, *Perspectives on Quantum Reality*, pages 125–142. Kluwer Academic Publishers.
- Saunders, S. (1996b). Time, quantum mechanics, and tense. *Synthese*, 107:19–53.
- Saunders, S. (1998). Time, quantum mechanics, and probability. *Synthese*, 114:373–404.
- Saunders, S. (2005). What is probability? In Elitzur, A., Dolev, S., and Kolenda, N., editors, *Quo Vadis Quantum Mechanics?*, pages 209–38. Springer Verlag.
- Saunders, S., Barrett, J., Kent, A., and Wallace, D. (2010). *Many Worlds? Everett, Realism and Quantum Mechanics*. Oxford University Press. Forthcoming.
- Saunders, S. and Wallace, D. (2008). Branching and uncertainty. *British Journal for the Philosophy of Science*, 29:293–305.
- Savage, L. (1954). *Foundations of Statistics*. Wiley.
- Schrödinger, E. (1935). The present situation in quantum mechanics. *Naturwissenschaften*, 23:807–812;823–828;844–849. repr. in Wheeler and Zurek (1983, pp.152-167).
- Timpson, C. G. (2003). On a supposed conceptual inadequacy of the Shannon information in quantum mechanics. *Studies in History and Philosophy of Modern Physics*, 34(3):441–68.
- Timpson, C. G. (2004). *Quantum Information Theory and the Foundations of Quantum Mechanics*. PhD thesis, University of Oxford. arXiv:quant-ph/0412063.
- Timpson, C. G. and Brown, H. R. (2005). Proper and improper separability. *International Journal of Quantum Information*, 3(4):679–690.
- Tumulka, R. (2006). A relativistic version of the Ghirardi-Rimini-Weber model. *Journal of Statistical Physics*, 125(4):821–840.
- Vaidman, L. (1998). On schizophrenic experiences of the neutron or why we should believe in the many worlds interpretation of quantum mechanics. *International Studies in the Philosophy of Science*, 12:245–261.
- Vaidman, L. (2002). Many-worlds interpretation of quantum mechanics. *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Zalta EN editor, URL = <http://plato.stanford.edu/archives/fall2008/entries/qm-manyworlds/>.
- Valentini, A. (1991a). Signal-locality, uncertainty, and the sub-quantum H -theorem. I. *Physics Letters A*, 156(1,2):5–11.
- Valentini, A. (1991b). Signal-locality, uncertainty, and the sub-quantum H -theorem. II. *Physics Letters A*, 158(1,2):1–8.
- Valentini, A. (2002a). Signal-locality in hidden-variable theories. *Physics Letters A*, 297:273.
- Valentini, A. (2002b). Subquantum information and computation. *Pramana Journal of Physics*, 59:31. arXiv:quant-ph/0203049.
- Valentini, A. (2004). Universal signature of non-quantum systems. *Physics Letters A*, 332:187–193.

- Valentini, A. and Westman, H. (2005). Dynamical origin of quantum probabilities. *Proceedings of the Royal Society of London A*, 461:253–272.
- Wallace, D. (2002). Worlds in the Everett interpretation. *Studies in History and Philosophy of Modern Physics*, 33:637–661.
- Wallace, D. (2003a). Everett and structure. *Studies in History and Philosophy of Modern Physics*, 34:87–105.
- Wallace, D. (2003b). Everettian rationality: Defending Deutsch’s approach to probability in the Everett interpretation. *Studies in History and Philosophy of Modern Physics*, 34(3):415–440.
- Wallace, D. (2006). Epistemology quantised: Circumstances in which we should come to believe in the Everett interpretation. *British Journal for the Philosophy of Science*, 57(4):655–689.
- Wallace, D. (2007). Quantum probability from subjective likelihood: Improving on Deutsch’s proof of the probability rule. *Studies in History and Philosophy of Modern Physics*, 38:311–332.
- Wallace, D. (2010a). *The Emergent Multiverse: Quantum Mechanics according to the Everett Interpretation*. OUP. Forthcoming.
- Wallace, D. (2010b). How to prove the Born rule. In Saunders et al. (2010).
- Wallace, D. and Timpson, C. G. (2010). Quantum mechanics on spacetime I: Spacetime state realism. Forthcoming in the *British Journal for the Philosophy of Science*.
- Wehrl, A. (1978). General properties of entropy. *Rev. Mod. Phys.*, 50(2):221–260.
- Wheeler, J. A. and Zurek, W. H., editors (1983). *Quantum Theory and Measurement*. Princeton University Press.
- White, A. R. (1975). *Modal Thinking*. Basil Blackwell.
- Zurek, W. H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75:715–765.
- Zurek, W. H. (2005). Probabilities from entanglement, Born’s rule $p_k = |\psi_k|^2$ from envariance. *Physical Review A*, 71:052105.