To appear in the British Journal for the Philosophy of Science 2010

Evidence and Uncertainty Everett's Multiverse

Paul Tappenden[1]

12 December 2009

> How does it come about then, that great scientists such as Einstein, Schrödinger and De Broglie are nevertheless dissatisfied with the situation? Of course, all these objections are levelled not against the correctness of the formulae, but against their interpretation. …. The lesson to be learned from what I have told of the origin of quantum mechanics is that probable refinements of mathematical methods will not suffice to produce a satisfactory theory, but that somewhere in our doctrine is hidden a concept, unjustified by experience, which we must eliminate to open up the road.[2]

> It is truly surprising how little difference all this makes. Most physicists use quantum mechanics every day in their working lives without needing to worry about the fundamental problem of its interpretation.[3]

ABSTRACT

I endorse the view that it may be of no relevance to the acceptability of the Everett interpretation of quantum mechanics as a physical theory whether or not an informed observer can be uncertain about the outcome of a quantum measurement prior to its having occurred. However, I suggest that the very possibility of post-measurement, pre-observation uncertainty has an essential role to play in both confirmation theory and decision theory in a branching universe. This is supported by arguments which do not appeal to Van Fraassen's Reflection Principle.

[1] paulpagetappenden@gmail.com

[2] Born ([1954], pp. 8, 11).

1 Dendritic unitary evolution

The world is more like a tree than a worm (world-tube). So says the 'relative state' formulation of quantum mechanics introduced by Hugh Everett III [1957] and often called the many-worlds interpretation. Much work has been done on the idea but there is still an ongoing dispute as to whether it is even intelligible. Two issues dominate the debate on intelligibility: how quantum probabilities are to be understood in a branching universe where all physical 'possibilities' actually occur and how the furniture of our everyday world can be constituted by the fundamental ontology of Everett's theory, that of the quantum-mechanical wave-function[4]. My aim here is to contribute to the work on probability, though I shall suggest that an idea arising in the course of the discussion points the way towards what needs to be done on the ontological problem.

Consider an experiment with a Stern-Gerlach apparatus to make a measurement of the spin of a particle of a given type relative to a chosen spatial direction, z. To simplify, there are conventionally two possible outcomes for this measurement: the result UP (apparatus indicating z-spin up) and the result DOWN. Depending on the details of the experimental setup quantum mechanics assigns an 'amplitude' to each of these outcomes which is a complex number. According to the Born rule the square of the modulus of this number yields the probability of the outcome with that amplitude actually occurring. The squared moduli of the amplitudes for the outcomes UP and DOWN thus sum exactly to unity in our simplified example which ignores bizarre outcomes of very low amplitude.

Mainstream quantum mechanics has taken the Born rule to entail the so-called collapse of the wave-function. The Stern-Gerlach apparatus stochastically evolves into either the UP state, with probability pU, or the DOWN state with probability pD, where pU + pD = 1. Everett proposed replacing the concept of stochastic evolution with that of dendritic evolution. The details of this process, if it indeed occurs, have turned out to be intricate and involve a phenomenon known as decoherence[5] but I shall begin with a provisional simplification about which I shall say more later: the apparatus and its environment can be thought of as bifurcating into a branch where the apparatus has evolved into the UP state and a branch where the apparatus has evolved into the DOWN state. Each of these branches has a quantum-mechanical amplitude and the squared moduli of these amplitudes, a quantity often referred to as branch weight, are wU and wD respectively, where wU + wD = 1. Given that branch weight is the square of a real number it can be conveniently thought of as a cross-sectional area, making vivid the idea of a tree trunk dividing into two branches whose combined cross-sectional areas are equal to that of the trunk. I shall henceforth refer to branch weight as branch cross-section. David Baker ([2007], §4.1) has suggested that this picture of a branching universe, or multiverse, involves an unacceptable circularity since it depends on the concept of decoherence for the selection of a measurement basis. Meeting Baker's objection, if it can be met, may require work on the ontological problem I mentioned earlier.

---

[3] Weinberg ([1992], p. 66).

[4] For a current challenge to the ontological coherence of Everett's theory see Maudlin [2010].

[5] See Wallace ([2008], §2) for an account and references. Note that a dendritic evolution may also take place in a different way during cosmogenesis; see Hogan [2007].

The Born rule is employed by stochastic quantum mechanics (QMs) for making predictions and is generally understood to have been corroborated to a high degree by observation. In the case before us the prediction is that any physical system which is considered to be of the same type as the given Stern-Gerlach experimental setup will evolve either into a system showing the result UP with probability pU or a system showing the result DOWN with probability pD. The evidence in favour of or against such a prediction which can be provided by a single observation is very limited and so a series of experiments must be conducted with systems of the appropriate type issuing in an observation of frequencies for the outcomes UP and DOWN. The observed frequencies must fall within a certain range in order to be considered as tending to confirm the prediction and that range is determined by standard statistical procedures. As is well known, any observed frequency is compatible with any predicted probability but an observed frequency has an associated probability of occurring given the predicted probabilities for outcomes. If the probability associated with the observed frequency is low relative to other possible observations then that observed frequency tends to count against the prediction. Quantum mechanics understood stochastically is thus open to the possibility of evidence which would reduce confidence in its truth.

Dendritic quantum mechanics (QMd)[6] predicts that a series of experiments conducted with the Stern-Gerlach apparatus will result in a ramified branching structure where every possible frequency distribution for a given outcome occurs somewhere in the many branches downstream of the experimental run. So there is an immediate problem as to how confirmation and disconfirmation are to be understood. This is not a matter of whether QMd is empirically preferred or not as against QMs but rather of whether quantum mechanics remains a theory open to being tested within the perspective of QMd. If QMd is to be favoured over QMs as an interpretation of quantum mechanics it is, as things stand, a conceptual matter and not one of superior empirical adequacy. It is primordial to the acceptance of the dendritic rather than the stochastic interpretation of quantum mechanics as a physical theory that quantum mechanics should remain subject to empirical test. The recent work of Hilary Greaves [2007b] and together with Wayne Myrvold [2010] has stressed the centrality of this 'evidential problem' for the scientific acceptability of QMd. David Albert also notes its importance, expressing the problem in the following way:

> Everettian pictures of the world are apparently not going to be susceptible of <u>confirmation</u> or <u>disconfirmation</u> by means of experiment - or not (at any rate) by means of anything even remotely like the sorts of experiments that we normally take to be confirmatory of Quantum Mechanics.
> ([2010], §1. Albert's emphases)

What needs to be done in order to resolve this problem in favour of QMd has been clear for some time[7]. Consider the spin-measurement experiment again and suppose that the details of the setup and the application of the Born rule indicate that the results UP and DOWN are equally probable. According to

[6] This term captures Everett's central idea, that of a branching environment, however, the dendritic structure is but a limited aspect of the universal wave-function, albeit one of special interest to observers such as us.

QMs, if the experiment is run a thousand times and the result is UP every time then that should count as evidence against quantum mechanics because the observed frequency is so improbable. But QMd predicts that there will, with certainty, be a branch where UP has been recorded a thousand times. QMd needs an observer on such a branch to be barred from taking the observation as evidence that s/he is on a thousandfold-UP branch for that would be equivalent to an adherent of QMs taking the observation of a thousand UPs as evidence that s/he was seeing an extremely improbable quantum-mechanical event rather than as evidence against quantum mechanics. Such a strategy would make quantum mechanics understood stochastically closed to empirical challenge. Quantum mechanics understood dendritically is likewise rendered immune to contrary evidence if an observer on the thousandfold-UP branch does not have good reason to take the observation as evidence against quantum mechanics and *ipso facto* against QMd.

That is the required conclusion but the problem is how to derive it. If a concept of pre-measurement uncertainty were available for QMd then the needed conclusion would fall out immediately. The informed QMd theorist would assign probabilities to outcomes and the observer on the thousandfold-UP branch would conclude that the observation of a thousand UPs is too improbable to be compatible with the acceptance of quantum mechanics and so of QMd. Simon Saunders and David Wallace [2008] have suggested that a concept of pre-measurement uncertainty is available for QMd. I have criticised that proposal elsewhere [2008] and subsequent work on the project can be found in Saunders [2010]. But whatever may become of that debate it can be shown to be irrelevant to the scientific acceptability of QMd if the needed conclusion about observers in the thousandfold-UP predicament can be derived without any reliance on the pre-measurement uncertainty of a fully informed QMd theorist. Greaves [2007b] and Greaves and Myrvold [2010] have proposed a route to the needed conclusion which proceeds via an analysis of how a subject entertaining QMd could rationally place wagers on the outcomes of quantum measurement processes. This method has been challenged by Albert who writes:

> But remember (and this is the absolutely crucial point) that deciding whether or not to bet on E, in the fission picture, has nothing whatsoever to do with guessing at whether or not E is going to occur. It is, for sure. And so is −E. And the business of deciding how to bet is just a matter of maximizing the payoffs on those particular branches that - for whatever reason - I happen to care most about. And if one is careful to keep all that at the center of one's attention, and if one is careful not to be mislead by the usual rhetoric of 'making a bet', then the epistemic strategy that Greaves and Myrvold recommend suddenly looks silly and sneaky and unmotivated and wrong.
> ([2010], §3(i), Albert's emphases)

These are strong words and there may well be much discussion of the deep clash of views they express. I shall not attempt to resolve the dispute here but what I shall do is argue that Greaves' and Myrvold's work can be enlighteningly amended in a way which sidesteps this line of objection. Albert finds it

[7] See, for instance, Papineau ([1996], p.239).

preposterous that the concept of betting can be applied upstream of branching contexts where all 'possible' relevant outcomes actually occur. And he ridicules the thought that punters' usual assessment of payoffs according to the presumed probabilities of outcomes can be replaced by a concept of pre-measurement future-directed 'caring' introduced by Lev Vaidman ([2002], §6.4) and Greaves ([2004], §2.3). The ideas I shall be presenting revolve around the fact that within the perspective of QMd it is always possible for a post-measurement subject to be held in a state of ignorance of, and so uncertainty about, the result on their branch. In that case, downstream of branching, a punter in the dark can bet in a very ordinary sense on whether E has or has not occurred *on their branch.*

In §2 I shall argue that if a particular assumption is adopted about how such a subject should assign credences to the possible outcomes on their branch then the evidential problem for QMd can be solved and so its status as a properly scientific theory vindicated. In §3 I shall suggest that Greaves' and Myrvold's arguments may be applied directly to justifying the needed assumption via analysis of betting behaviour on the part of a post-measurement, pre-observation subject rather than a pre-measurement subject. In §4 I shall argue that although a justification of betting behaviour on the part of the pre-measurement subject who assumes QMd is not required in order for the interpretation to be a scientifically acceptable theory there is a satisfactory analysis available which depends on a novel appeal to the possibility of post-measurement, pre-observation uncertainty. The argument does not invoke Van Fraassen's Reflection Principle which has been applied to the problem but which is unreliable in this context. Finally, in §5 I shall discuss some consequences of the foregoing arguments which, though bizarre, do not challenge the scientific status of QMd.

2 The Born-Vaidman rule

Consider an observer downstream of an experimental run with the Stern-Gerlach apparatus prepared to have equal amplitudes for the UP and DOWN results as before. This observer, call her Hydra, may or may not have conducted the experiment herself but she is well informed about the setup and has good reason to think that it is reliable. Imagine that Hydra is blindfolded so that she cannot see the frequency readout and suppose first of all that her working hypothesis is QMs. Knowing the details of the experimental setup Hydra can apply the Born rule to assign probabilities to each of the various possible types of frequency distribution she may observe on removing the blindfold. She will set these probabilities as equal to the squared modulus of amplitude of each of the relevant components of the wave function of the system. Vaidman ([1998], p. 253) has suggested that blindfolded Hydra should assign probabilities to seeing particular outcomes on removing the blindfold in exactly the same way if her working hypothesis were dendritic rather than stochastic quantum mechanics.

A few words are appropriate here on a simplification I have been making. According to the analysis of branching in terms of decoherence there will not be one but rather a vast and strictly indefinite number of branches where a thousand UPs occur. These branches will be macroscopically isomorphic but microscopically anisomorphic. The so-called amplitude of the thousandfold-UP branch will in fact be the sum of the amplitudes of all those microscopically anisomorphic branches. Vaidman's idea can be expressed if Hydra is understood to inhabit one of the many thousandfold-UP branches but the picture is simpler if Hydra *qua* observer is imagined to be a single subject somehow inhabiting the totality of the relevant near-isomorphic branches. I shall continue with the simplification for the sake of clarity.

Note that whether it is QMs or QMd which the blindfolded Hydra has in mind it is an assignment of a subjective probability, a credence, which she is making: a judgement about the probabilities of seeing this or that type of frequency distribution on removing the blindfold. But there seems to be an important difference between the Born rule and what I shall call the Born-Vaidman rule in this context. Although they act identically in urging Hydra to set her credences equal to the squared moduli of the relevant amplitudes they differ in their implied explanations for this. According to the Born rule, as used in QMs, the reason Hydra should assign a particular credence to seeing a given result on removing the blindfold is that that is the predicted objective dynamical probability that that result has occurred. According to the Born-Vaidman rule Hydra believes that she has counterparts in other branches who will see different frequency distributions on removing their blindfolds and the reason that Hydra, like each of her blindfolded counterparts, should assign a particular credence to seeing a given result is that that result is predicted to occur on a branch with the corresponding cross-section (squared modulus of amplitude).

So there seems to be a problem here. The explanation for Hydra's credence assignments if she is following the Born rule and assuming QMs looks straightforward: she sets her credence to what she believes is the objective dynamical chance of a particular result having occurred, following what David Lewis has called the Principal Principle ([1980], p. 266). Contrariwise, the explanation for Hydra's credence assignments if she is following the Born-Vaidman rule and assuming QMd looks strange: why should believing the branch cross-sections for particular results to be thus and so motivate her to assign corresponding credences to seeing any one of those results on removing the blindfold?

I shall return to this problem in a moment but first of all we should recognise that *if* Hydra presumes QMd and adopts the Born-Vaidman rule she is just as much open to discovering evidence which should reduce her confidence in quantum mechanics as she would be in presuming QMs and adopting the Born rule. For suppose that Hydra sees that a thousand UPs have occurred on removing her blindfold after a run of measurements with the Stern-Gerlach apparatus. If her working hypothesis is QMs she will take this observation as reducing her confidence in the truth of quantum mechanics, modulo her confidence in the experimental setup, because it would be unreasonable to believe that quantum mechanics is true and a circumstance to which she had previously accorded very low credence pertains, that is, the occurrence of an event with very low objective dynamical probability. Likewise, if Hydra's working hypothesis is QMd she should take her observation as reducing her confidence in quantum mechanics because it would be unreasonable for her to believe that quantum mechanics is true and a circumstance to which she had previously accorded very low credence pertains, that is, her finding herself on a branch of very low cross-section relative to the branch where the experimental run began.

This thought needs some unpacking for it might seem that there is an important difference here between QMs and QMd. According to QMs Hydra on removing the blindfold sees a state of affairs with a low objective dynamical probability of having occurred a few moments earlier and that seems to be an improbable state of affairs in a clear sense. But according to QMd Hydra's being on a branch of very low cross-section is not objectively improbable in any sense, it is simply a circumstance to which blindfolded Hydra accords low credence if she accepts the Born-Vaidman rule. However, credence can do the work in reducing Hydra's confidence in quantum mechanics prior to any concept of objective probability. To see this, consider what has traditionally been regarded as constituting disconfirmatory evidence in classical mechanics where physical processes are taken to be fully determined and so there has, until recently, been no question of objective probability being involved. Probability has there been understood subjectively as an assignment of credences based on limited knowledge of the world, the so-called ignorance interpretation. So in classical mechanics the observation of a state of affairs which has been accorded relatively low credence has been taken as counting as evidence against the theory even thought there is no sense in which that observed state of affairs is objectively improbable.

It is true that there have been challenges to the traditional idea that classical mechanics does not involve objective probability[8] but note that even there it is not suggested that classical mechanics involves a stochastic dynamical evolution, which is the type of objective probability which has been taken to be suggested by the Born rule. Also, Roman Frigg [2008] has fielded counterarguments and the debate remains live. In any case, to dismiss the idea that credence is the fulcrum of confirmation theory in QMd would require rejecting what has traditionally been seen as an unproblematic process of confirmation in that other deterministic physical theory, classical mechanics.

Hydra's predicament on seeing the thousandfold-UP result after removing the blindfold would seem to be exactly the same as regards judging the evidence whether she has been using the Born rule and QMs or the Born-Vaidman rule and QMd. Unsurprisingly, since the Born rule and the Born-Vaidman rule both maintain that the credence for seeing that result should be equal to the squared modulus of amplitude of the corresponding component of the wave-function of the system. Hydra has good reason not to believe that she is on a branch of low cross-section if assuming QMd just as she has good reason not to believe that a dynamical evolution with low objective chance has taken place if assuming QMs. So it seems that the evidential problem for QMd dissolves for the temporarily blindfolded post-measurement observer if the Born-Vaidman rule is adopted.

But what if Hydra were not temporarily blindfolded post-measurement? What if an experimental run had been conducted and she were simply observing that a thousand UPs have been registered? To understand Hydra's predicament in this case we should first of all note that on the assumption of QMd without pre-measurement uncertainty she has no record of a prior prediction that seeing a thousand UPs is very improbable[9]. The prediction she has in hand is that a thousand UPs will be registered on a branch of very low cross-section (squared modulus of amplitude) from which it follows that she would assign very low credence to seeing that a thousand UPs have occurred if she were blindfolded and applying the Born-Vaidman rule.

Since Hydra on the thousandfold-UP branch plainly sees a record of a thousand UPs having occurred should she simply accord a credence of unity to seeing that result? Yes and no. Yes in the sense that she does not have reason to doubt what she's seeing, but that is not the sense which is wanted here. For suppose that Hydra were assuming QMs and were seeing that a thousand UPs had occurred. In that case she would assign a credence of unity to her seeing that result, she would believe her eyes, but she

[8] For the most recent see Loewer [2004] and Ismael [forthcoming].

would tend to take the observation as evidence against QM, given that no more plausible explanation were available, because QMs predicts that the occurrence of a thousand UPs is objectively improbable *which entails that she would have assigned low credence to seeing that result if she had been blindfolded beforehand.*

Recall that the only difference between the Born rule as applied by QMs to the post-measurement, pre-observation subject and the Born-Vaidman rule is the explanations they imply for the credence assignments. If the Born-Vaidman rule is acceptable in the context of QMd then a strategy is available to non-blindfolded Hydra for making a judgement about the evidential status of observing that a thousand UPs have occurred. She can base that judgement on the credence she *would have* assigned to seeing a thousand UPs *if* she had been blindfolded prior to making the observation. The very possibility of being in a state of post-measurement, pre-observation ignorance and so applying the Born-Vaidman rule is enough to yield a method for solving the evidential problem of QMd. Hydra is not compelled to adopt that method but if she rejects it she should abandon QMd as unscientific if no other method of judging evidence is available.

In that case it is just the acceptability of the Born-Vaidman rule itself which remains at issue. At this point it is worth saying something about the acceptability of the concept of dynamical chance in QMs. For some reason the idea of an irreducibly stochastic process has come to seem natural and the application of the Principal Principle to derive credences from what are believed to be objective chances seems compelling. But the contemporary familiarity of the concept of stochastic process ought not to make us too complacent. Until the advent of quantum mechanics all physical processes were thought to involve a mechanism of some sort but a stochastic dynamical evolution is a process without a mechanism. Quantum mechanics understood dendritically respects the tradition in science of explaining processes in terms of mechanisms. However, there remains a problem: what motivates adoption of the Born-Vaidman rule?

I shall come to this in the next section but first some further reflection on the picture developed so far. We left Hydra on the thousandfold-UP branch. We saw that if her working hypothesis is QMd she has good reason to dismiss the idea that she's on a thousandfold-UP branch and take what she observes as evidence against quantum mechanics so long as she has full confidence in the experimental setup. That's because she knows that if she had been blindfolded and were applying the Born-Vaidman rule she would have assigned a relatively very low credence to seeing the thousandfold-UP result. However, the hypothesis is that Hydra is on a thousandfold-UP branch even though she has reason not to believe that. What this shows is that the branch cross-section of Hydra's environment relative to the branch where the experimental run began is radically invisible to her. There is no known way she can measure it directly. If she could she would be able to ascertain that she is indeed on a thousandfold-UP branch, incredible as that is according to the Born-Vaidman rule. The lack of any means of direct measurement of relative branch cross-section in QMd corresponds to the lack of direct access to dynamical chances in QMs. According to QMs the observation of a relatively very unlikely event must be taken as evidence against

[9] An interesting discussion of the predicament of an observer post-branching is to be found in Saunders ([1997], pp. 64-6).

quantum mechanics unless that observation is counterbalanced by a mass of other evidence suggesting otherwise. There is no way of detecting the supposed fact that an objectively improbable quantum-mechanical event has occurred.

The concept of branch cross-section, the squared modulus of amplitude of a branch, is as some extant physical quantity hitherto unknown to science. That is the implication of QMd, of course, since Everett's key thought was that the wave-function represents actualia rather than possibilia[10]. It is suggestive of some sort of novel physical extension. The idea is not new; Michael Lockwood ([1989], p. 232) has called this quantity a superpositional 'dimension', not meaning that to be taken too literally, and attributes the origin of the idea to David Deutsch. I believe that developing an understanding of the concrete physical nature of branch cross-section is central to addressing the fundamental ontological problem for QMd which I mentioned at the outset, but I cannot pursue that thought further here.

3 Greaves on uncertainty

We are left with the problem of how to motivate the Born-Vaidman rule if it is not simply to be taken as a fundamental postulate of QMd. I believe that this problem might be addressed via work undertaken by Greaves and Myrvold but that a change of perspective may be required which arises from a new supplementary argument. So far, Greaves' and Myrvold's project has been applied to the analysis of a probabilistic-like attitude to be adopted by the pre-measurement subject. Greaves has argued for future-directed assignments of 'caring measure' or 'quasi-credence' which arises out of trying to resolve both the evidential problem which I have been discussing and the problem of how a subject might rationally place bets when facing the prospect of a payoff regime associated with the outcomes on various branches. If Greaves' and Myrvold's work were applied to the assignment of credences by the post-measurement, pre-observation subject its relevance to the evidential problem need not involve an irreducible concept of caring measure, quasi-credence or 'α-function'[11]. Whatever this concept is called, it amounts to the idea of a non-credence attitude which mimics credence in the process of pre-measurement decision-making. But such a concept looks puzzling. Why should I 'care' about the amplitude of a future branch when I know that for the inhabitants of that branch the relative amplitude makes no difference whatsoever to their quality of experience? An experience of joy or pain is in no way diluted by occurring on a relatively low-amplitude branch according to Greaves' and Myrvold's view and yet prior to branching an agent must be guided by value assignments based on those future branch amplitudes.

This puzzle can be dissolved if the underlying reason for the assignment of quasi-credences is seen to be the possibility of assignments of copper-bottomed credences by the post-measurement, pre-observation subject. I do not mean that the pre-measurement subject does not need to 'care' about future branches in the way that Greaves and Myrvold suggest, rather I am arguing that the air of mystery surrounding the idea of future-directed caring disappears if it is seen as arising out of a conceptual

[10] Strictly speaking, the wave-function instantiates actualia in QMd and represents possibilia in QMs.
[11] The term currently used by Greaves and Myrvold [2010].

dependence on the very possibility of the sort of post-measurement ignorance which Vaidman was the first to notice.

Certainly there remains the problem of why the Born-Vaidman rule should apply but if Greaves' and Myrvold's arguments are correct then purchase on that problem should be gained by applying their method to the gambling predicaments of post-measurement, pre-observation subjects. That strategy may also be complimented by other work which I shall mention shortly. In the next section I shall explain how the problem of guiding the informed pre-measurement punter's betting can be resolved by an independent argument linking possible post-measurement, pre-observation credences with a rationale for pre-measurement decision-making.

So far I have argued that the evidential problem for QMd can be solved if the Born-Vaidman rule is accepted. This is recognised by Greaves when she writes:

> Under relatively uncontentious auxiliary assumptions, a necessary and sufficient condition for solving the Epistemic Problem is as follows: establish that the agent who believes that QME is true and that branching has occurred, but who has gained no post-split information as to which branch she is on, should set her self-locating credences equal to the assumed branch weights. ([2007b], §6).

The term 'QME' stands for 'Everettian quantum mechanics', equivalent to my term 'QMd' for 'dendritic quantum mechanics', an idea which has been developed a great deal since Everett's original proposal and could well do with a more descriptive label. Greaves here calls the evidential problem the 'Epistemic Problem' but she and Myrvold have more recently used the term 'Everettian evidential problem' ([2010], §1). Greaves then goes on to discuss a possible objection to her use of the concept of an agent's pre-measurement 'caring measure' in her defence of QMd and she suggests a 'Response 1' to this objection. That response makes two points. One is in effect that since the application of the Born rule in QMs assumes the Principal Principle without argument to demand more for the correlate of the Principal Principle which is assumed by the application of the Born-Vaidman rule in QMd is to apply double standards. The other point is that nevertheless this further support for the 'Everettian Principal Principle' might be available via a programme initiated by Deutsch[12]. Greaves goes on to write:

> Response 1, however, leaves it somewhat mysterious whether the arguments adduced in this paper have any role to play in the defence of the Everett interpretation: if we can argue from the pre-measurement perspective that the Everettian branch weights play the same role as chances in practical action, and we can argue from the post-measurement perspective that Everettian branch weights play the same role as chances in theory confirmation, then we have solved both the Practical Problem and the Epistemic Problem; who cares, then, about the link between the two?

We must care about the link for the following reason. If, having independently advocated a particular strategy for rational action and a particular epistemic strategy, we were to find that the combination of the two led to failure of intertemporal consistency ……… then our argument as a whole would have landed us in paradox. ([2007], *ibid*.)

What Greaves calls the practical problem is the problem of how an agent who believes QMd can act rationally prior to measurement. For example, if a subject believing QMs is presented with a gamble on the outcome of a quantum process s/he can make a judgement about odds and possible payoffs which can guide the laying of a stake; a manifestation of the practical problem is how a subject believing QMd can make a corresponding judgement about what Barry Loewer has dubbed the 'bramble' which is the dendritic equivalent of the stochastic setup with branch cross-sections corresponding to chances. The argument I shall be presenting in the next section should make it clear that the Born-Vaidman rule can also, as it happens, be used to underwrite a rationale for pre-measurement betting, though that is not a requirement for the scientific acceptability of QMd. Recall that the Born-Vaidman rule is equivalent to the application of the Born rule to the post-measurement, pre-observation subject so far as the assignment of credences is concerned, the difference is the Born-Vaidman rule presumes QMd. Again, I should emphasise that the implication is not that Greaves' and Myrvold's arguments have no bearing on the evidential problem for QMd but rather that their work is not necessary to providing the link between the evidential and practical problems and so can be applied directly to the evidential problem by a decision-theoretic analysis of the predicament of a post-measurement, pre-observation subject, something which Greaves herself acknowledges as possible.

Before going on I want to dwell for a moment on another aspect of Greaves' motivation for avoiding any appeal to the post-measurement, pre-observation uncertainty which Vaidman invokes. Greaves ([2007b], §5.2.1) gives two reasons for not relying on this form of uncertainty. The first and lesser reason is that it does not generally occur. The second and more important reason for her is:

…exclusive reliance on the adapted Minimum Information argument would promote the idea that genuine *uncertainty* about which outcome (if not pre-measurement, then post-measurement) occurs is an *essential* part of the Everettian's epistemological story, even in cases in which all details of the branching structure are known. One of the central aims of this paper is to dispel this notion. To fulfil this aim, we must supply an argument that does not depend on this sort of 'branch uncertainty' at all. ([2007b], *ibid*.. Greaves' emphasis)

Why is Greaves so intent on banishing any dependence on post-measurement, pre-observation uncertainty? A possible explanation is that she wants to tell the Everettian epistemological story in the logically purest way, purged of any unnecessary assumptions. I am arguing that it is the mere possibility of post-measurement, pre-observation uncertainty which may be essential. That is a *necessary* feature of

[12] See Greaves ([2007a], §2 & §3) for an account of this work and references. See Wallace [2010] for the most recent manifestation of the project. This is the work which I mentioned may compliment Greaves' and

QMd, just as it is of QMs. But there is a clue in her earlier writing to another aspect of what may drive Greaves to banish uncertainty. She considers a thought experiment:

> Suppose (somewhat artificially) that I am confronted with the choice of whether or not to sign a particular form. After I make my choice, a Stern-Gerlach experiment will be carried out. If I signed, my spin-up successor will be force-fed olives, and my spin-down successor will be given chocolate. If I didn't sign, no-one is fed. But I hate olives. How do I decide whether having one successor worse off and another better off is preferable to having all my successors stuck with the status quo? Do I sign or not? ([2004], §4.3)

And she goes on to write:

> I can now say why, in my view, the Reflection approach obscures the logic of Everettian decisions. Suppose, for the sake of argument, that, in the above scenario, the relevant outcomes (chocolate, hunger, olives) are equally spaced in terms of utility. Then if I strictly prefer signing the form to not signing it (perhaps in a fission setup in which the spin-down successor has higher quantum-mechanical weight), that is so *because* I care more about my spin-down successor than I do about my spin-up successor. It is not *because* my post-fission selves will suffer weighted disorientation, although (as emphasized by the Reflection argument) that is also true. (*ibid.* Greaves' emphasis)

I shall say more about what Greaves refers to as the Reflection argument shortly. What is important to notice here is that she accords explanatory primacy to the notion of pre-measurement caring but gives no supporting argument; she merely expresses an opinion. What she means by 'weighted disorientation' is the state of uncertainty experienced by a post-fission, pre-observation subject. So Greaves seems to accept the possibility of an argument with the force of the one I shall be presenting in the next section but summarily dismisses it whilst expressing no concern for the strangeness of the unprecedented concept of pre-measurement 'caring' which is a non-credence attitude operating as if it were a credence.

If the Reflection argument worked then all that would be required to resolve the puzzle about caring would be to suggest that the direction of explanation may go in exactly the opposite direction to the one Greaves favours. In that case the non-credence concept of caring would be seen not to be primitive but rather derived from that of post-measurement, pre-observation credences. However, it is not clear that the Reflection argument does work. Here is how Greaves presents the idea:

> But now consider the following *decision-theoretic reflection principle*:
>
> > If, at time t, I decide rationally to pursue a certain strategy at a later time t′, and if I gain no new information relevant to that strategy between times t and

Myrvold's if applied to justifying the Born-Vaidman rule.

t′, then it is rational [i.e. rationally compelling] not to change my choice of strategy at t′. (Wallace (2002b), p.58)

This principle, I think, is a reasonable one, and it holds the key to the second line of response to the applicability problem: if we want to apply a decision theory whose application is justified under uncertainty, we can apply it directly to the post-fission observer stages on the basis of their genuine uncertainty, and appeal to reflection in order to argue that the pre-fission observer-stage should adopt the same strategy for action. ([2004], §4.2)

Greaves' quote from Wallace [2002] here is the passage in which Wallace puts in his own words Van Fraassen's Reflection Principle ([1995], *passim*) and Wallace goes on to note that this principle has been challenged by Adam Elga [2000]. That the principle is not universally applicable in the context of QMd can quickly be seen by noting that in an actual measurement situation it can be arranged for some of an observer's downstream successors to be immediately and painlessly killed post-measurement so the informed downstream post-measurement, pre-observation survivors do indeed gain relevant information between t and t′: they know that they are not dead. Furthermore, in any measurement situation we would expect there to be non-zero amplitude for deaths of that sort.[13]

So it is not clear that the principle which Greaves endorses as reasonable is, after all, acceptable. Wallace did not accord any importance to the argument as he himself saw it but remained nonetheless open to the possibility that some effective argument might be given for the relevance of post-measurement uncertainty. Having dubbed as the 'fission program' the approach to QMd which aims to do without any concept of pre-measurement uncertainty for the informed subject, Wallace writes, with reference to the strategy of appealing to post-measurement uncertainty[14]:

> This strategy provides, so far as I can see, the only promising means to salvage the fission program; however, I do not find it wholly satisfactory, for two (admittedly very inconclusive) reasons. ([2006], §4.2)

In the following section I aim to fulfil the promise in a way which is satisfactory and which does not invoke the Reflection Principle.

4 Retroactive counterfactual uncertainty

A convenient starting point for tackling the practical problem for QMd is the analysis of betting behaviour. Believing conventional QMs, a subject is generally seen as being able to unproblematically

[13] This was pointed out to me by David Baker.

[14] The term 'fission program' suggests itself because any concept of pre-measurement uncertainty for QMd requires that branching does not involve the fission of worlds but rather the divergence of previously isomorphic worlds. For an early application of the idea of divergence to QMd see Deutsch ([1985], p.20) and for a current application see Saunders [2010].

make judgements about bets on the outcomes of quantum measurements given the probabilities of the outcomes and the payoffs associated with them. If would be perplexing if a subject following QMd were not to have good reason to behave in at least a fairly similar fashion. However this is a pragmatic matter which is not central to the admissibility of QMd as a physical theory which, as we have seen, follows from the resolution of the evidential problem. Also, the idea of betting on quantum measurements is rather contrived. Even if we live in Everett's multiverse, when you pull the handle of a normal fruit machine the final resting places of the fruit wheels can no doubt be computed in advance given reasonably accurate knowledge of their initial speed and the friction acting on them. There will not be a branching of the machine and its environment into all the classically possible outcomes with branch cross-sections corresponding to the classical probabilities. The branching which takes place will be hugely dominated by the quasi-classically determined outcome. In that case it would be appropriate for prospective uncertainty to guide betting behaviour exactly as conventionally understood. In most everyday situations it would be very challenging to compute to what extent macroscopic processes about which decisions are to be made are subject to relevant branching. So for all practical purposes a subject tending to believe QMd will remain uncertain about whether a particular context involves relevant branching or not simply because of ignorance of the highly complex quantum-mechanical details. Still, it is clearly possible to bet on quantum measurements rather than quasi-classical processes and it is natural to want to know how a punter should behave if believing QMd rather than QMs.

Betting normally involves behaving in a way which is believed to maximise personal gain and so involves the concepts of personal expectation and identity over time. I shall discuss later how it is possible to go beyond these constraints but in the meantime it will be useful to work with them. The reason is that the adoption of QMd as a view of our world poses a challenge as to how we can coherently speak of persons and objects persisting through time and it is useful to be aware of ways in which that challenge can be met. So let us take Hydra back in time. She is about to make a measurement with the Stern-Gerlach apparatus and suppose now that it is set up so as to operate as a 'quantum die'. On pressing a button the apparatus measures the spin of a particle of a certain type relative to a particular direction so that according to QMs there is a chance of one sixth that it will show the result UP and of five sixths that it will show DOWN. According to (simplified) QMd the quantum die, Hydra and her environment bifurcate into branches of unequal amplitude, the squared moduli of the amplitudes being one sixth for the UP branch and five sixths for the DOWN branch. On the UP branch there will be a 'successor' of Hydra who sees the result UP. Call her HydraUP. Likewise there will be the successor HydraDOWN on the DOWN branch.

What should Hydra expect to experience if she presses the button on the quantum die? I shall proceed on the assumption that she should expect to see the result UP and the result DOWN simultaneously as distinct experiences. This is a strange idea. If QMd is indeed coherent it is not for nothing that it is taking so long to establish the point. A first thought is that Hydra can only expect to have experiences which will be hers and so she must in some sense become HydraUP and become HydraDOWN in order to expect having each of their experiences simultaneously. It can perhaps be argued that personal expectancy need not depend on continuing personal identity but there are various theories in the literature on persistence which might be used to maintain Hydra's identity. Since it would be difficult to live with QMd if it were not possible to track the identities of persons and objects in the multiverse I shall adopt a way of speaking which allows for such tracking. Examples of theories which ostensibly are able to cope with persistence through fission and might be applicable to what follows are John Perry's 'lifetime' theory [1972], Ted Sider's 'stage' theory [1996, 2001] and André Gallois' 'occasional identity' [1998]. I shall use Sider's stage theory because it is perspicuous and fairly well established though of course not without its critics.

According to stage theory persons are momentary stages so Hydra, HydraUP and HydraDOWN are distinct persons. However, HydraUP and HydraDOWN are future 'temporal counterparts' of Hydra which means that Hydra bears the relation *will be* to HydraUP and HydraDOWN severally and HydraUP and HydraDOWN bear the relation *was* to Hydra. This idea is an adaptation of David Lewis's [1968] introduction of 'counterpart theory' to modal logic according to which Hubert Humphrey bore the relation *might have been* to being elected president if and only if he had a modal counterpart who became president. On Sider's view ([2001], p. 201), which is applied to thought experiments about personal fission rather than QMd, Hydra can expect to become HydraUP and she can expect to become HydraDOWN. This analysis of fission applies generally to transtemporal identities not just that of persons so ApparatusUP, which shows the result UP, and ApparatusDOWN, which shows the result DOWN, are distinct apparatuses but each of them was the very apparatus which Hydra prepared to make the quantum measurement.

Returning to the quantum die, let QMd be Hydra's guiding hypothesis. After having pressed the button, if blindfolded and applying the Born-Vaidman rule, HydraUP and HydraDOWN would each assign a credence of one sixth to seeing the result UP on removing the blindfold. And being betting women they would both be prepared to lay a stake if told that they would be paid seven-to-one if they were to see the result UP. Now suppose that Hydra is blindfolded before playing and is told that the blindfold won't be removed until a minute after she presses the button. She is also told that there will be no payoff unless the stake is laid before playing, as is usual when betting on dice rolls. In this situation the well-informed Hydra can confidently expect that immediately after pressing the button she will assign a credence of one sixth to seeing the result UP and would want to lay a stake for a seven-to-one payoff on that result. Hydra can be confident that after pressing the button she will regret not having laid a stake beforehand if she has not done so. That seems to be a compelling reason for Hydra to lay a stake before pressing the button. So blindfolded Hydra has good reason to lay a stake in advance, just like a person who takes the quantum die to have a conventional stochastic dynamical evolution rather than a dendritic evolution.

Some qualification is called for here. Hydra will know that HydraUP will be aware that whilst she has won there is a closely related person who has lost, which may temper HydraUP's pleasure. Likewise, Hydra will know that HydraDOWN will be aware that whilst she has lost there is a closely related person who has won, which may temper HydraDOWN's displeasure. This thought could have an effect on what odds blindfolded HydraUP and HydraDOWN judge to be sufficient to make it worth laying a stake. However, on Sider's analysis HydraUP and HydraDOWN are distinct people who are not in causal contact. Conventionally, if a punter were convincingly told by a demon that she had a doppelganger on a planet in a distant galaxy and that things would be arranged so that if she won a bet her doppelganger would lose and *vice versa* would that make a difference to how the punter bet? Perhaps. It would be a matter of individual judgement but many people may not be too bothered about the fate of a distant doppelganger. It is true that HydraUP and HydraDOWN have a common history but that has been and gone by the time the crucial quantum measurement has been made and decoherence has in a sense brought more distance between them than being in different galaxies.

I shall return to the theme of possible anomalous consequences for betting behaviour in the final section but first we need to consider what attitude Hydra can reasonably adopt if she is not blindfolded when offered bets on the quantum die. What if she does not expect to be in a state of post-measurement, pre-observation ignorance? The challenge to Hydra in this situation is how to weight the utilities of her successors and it is plausible that there is no reason for this to differ from the way she would make judgements if blindfolded. To see why, imagine that Hydra were told that so long as she laid a trivial stake before a measurement HydraUP would get an enormous payoff and HydraDOWN would simply lose the stake. In that case it would seem compelling that Hydra should lay the stake for there would be so much to gain and so little to lose. But if the ratio between the stake and the payoff is progressively reduced where should Hydra draw the line between laying and not laying a stake in advance? The possibility of post-measurement, pre-observation ignorance provides a guide here, so long as the Born-Vaidman rule is assumed. Hydra knows that if she were blindfolded after the measurement she would want a stake to have been laid beforehand if the payoff on seeing UP is better than six-to-one, setting aside the proviso I mentioned. Hydra knows that HydraDOWN will be disappointed to have lost but she also knows that HydraDOWN cannot reasonably regret having laid a stake on a bet which she would have thought good had she been blindfolded.

So it looks as though Hydra without the blindfold still has good reason to lay the stake in advance, just like a subject believing QMs, always bearing in mind the qualification about possible empathy between winner and loser. It's not that Hydra is compelled to take into account what her attitude would be if she were blindfolded post-measurement but it is an option open to her which is not obviously unreasonable and provides a coherent strategy. Post-measurement, pre-observation uncertainty can constrain pre-measurement betting behaviour for a subject adopting QMd even if that uncertainty does not actually occur because simply considering the predicament of the post-measurement, pre-observation subject is enough to guide a punter's reasoning. So counterfactual post-measurement uncertainty can be effective in determining pre-measurement betting behaviour in a way which bears a good deal of similarity to the pre-measurement uncertainty of conventional stochastic quantum mechanics.

The argument has been based on an example of betting for personal gain but it can be detached from strictly personal interests[15]. Imagine that the demon convinces Hydra that some time after the quantum die measurement takes place a distant country will be blighted by plague in all worlds issuing from the DOWN branch, of relative cross-section five sixths, unless she makes a significant sacrifice before pressing the button, in which case it will be in the worlds issuing from the UP branch that that country will be blighted by plague. If Hydra reflects on the predicament of inhabitants of that country who are ignorant of the origin of their world but aware of the general setup she will realise that any such

[15] As Greaves and Myrvold point out ([2010], §3).

subject accepting the Born-Vaidman rule would assign a credence of one sixth to having issued from the UP branch and five sixths to having issued from the DOWN branch. All those future subjects would, in general, hope that Hydra had made the appropriate sacrifice and Hydra would know that in advance, which might be enough to motivate her to make the sacrifice even though she has no self-centred interest in the outcome.

We have seen that a subject adopting QMd as a guiding hypothesis need not be rendered incapable of making rational future-directed decisions in the special quantum-measurement circumstances where the various outcomes of an action are presumed to exist simultaneously in different branches of a multiverse. The very possibility of applying the Born-Vaidman rule post-branching to assign credences to outcomes can serve as a substantive guide to pre-branching action. However, there are bizarre consequences of the overall picture which I must discuss before closing.

5 Odds and ends

When there is more at stake than a fairly trivial bet the reservation on Hydra making judgements about decisions on the basis of Born-Vaidman rule credences looks more pressing. An adaptation of Huw Price's 'Legless at Bondi' thought experiment can illustrate the point ([2010], §6)[16]. Imagine that Hydra had lost a leg to a shark and were offered a new form of treatment which involved a quantum measurement where, to simplify, there were two relevant downstream branches. On one branch of very high amplitude she would awake from the treatment with two good legs and on the other very low-amplitude branch, as an unavoidable consequence of the treatment, she would awake with no legs at all. Supposing that a post-operative patient could be held in a state of wakefulness whilst unaware of whether she were legless or not, the above argument suggests that Hydra's pre-operative decision should be based on the post-operative Born-Vaidman rule credence assignments. Given a credence of only one in a billion of finding oneself legless most people would opt for the treatment. But if assuming QMd Hydra can be sure in advance that the treatment will result in there being a legless person. Price suggests, with good reason, that a case like this bears some similarity to the famous Trolley Problem, mentioning the version where a fat man might be pushed off a bridge to his death in order to stop a runaway tram and thereby save several passengers. Price notes that most people reject killing the fat man. Likewise it might be though Hydra would similarly have to reject the leg treatment given that she were confident that there would be a legless person paying a terrible price for the relief of the bipedal one. There is, of course, an important difference between the QMd case and the Trolley Problem analogue. The patient offered the quantum branching treatment, on Sider's analysis of transtemporal identity, can believe that it is she who will be each of the downstream people. If Hydra goes ahead with the treatment she will know that the person finding herself legless will have only herself to blame for having accepted the treatment in the sense that the legless person was the person who signed the consent form.

[16] This problem for the fission program is discussed by Greaves and Myrvold ([2010], Objection 8).

But I do not wish to play down too much the differences between decision-making according to QMs and QMd. It is plausible that for ordinary rambles through brambles Hydra's betting judgements if assuming QMd need not differ very much, if at all, from those she would make if assuming QMs but there may well be special circumstances where that is not so. That does not constitute a threat to the scientific credentials of QMd as it only makes practical action problematic, not the evaluation of evidence. My point is that nonetheless a subject assuming QMd is far from being entirely at sea in respect of practical action. The very possibility of Born-Vaidman rule credence assignments can serve as a substantive guide to pre-measurement decision making, a guide which in non-extreme circumstances may be pretty much as good as conventional assignments based on presumed chances.

Finally, we need to consider a class of betting scenarios even more extreme than Price's example. Imagine that the quantum die is modified so as to kill HydraDOWN instantly and painlessly before she observes the result in her branch whilst HydraUP gets a large payoff. This is quantum Russian roulette (qRr) with five bullets in the cylinder. If Hydra knows that HydraDOWN is to die instantly it can seem that Hydra should expect with certainty to get the payoff. I have argued in support of that conclusion elsewhere [2004] and that argument would also apply to Greaves' and Myrvold's view[17]. However, we should hesitate in taking this as a way for Hydra to make easy money since she would know that her young children would be left motherless in the fatal branch, something she would care about very much. But even if not necessarily a route to easy money this view is a radically different perspective on qRr scenarios from that yielded by QMs. According to QMs, if you play qRr you have a good chance of bringing an end to your personal experience. That is not so on the analysis I have presented.

Furthermore, the post-measurement, pre-observation credence assignments must be renormalized in scenarios of this sort and cease to be numerically equal to branch cross-sections (squared moduli of amplitudes). For instance, if HydraUP is well informed and blindfolded and knows that the quantum measurement has taken place she will be confident that she will see UP when the blindfold is removed since she knows that she cannot be HydraDOWN, who dies immediately. However, the relative cross-section of HydraUP's branch is one sixth. She must renormalize this value to unity in making her blindfolded judgement of the credence for seeing UP.

If Hydra had full confidence in QMd and were wholly egocentric and if she were offered the chance of playing qRr with a guaranteed reward on the UP branch then it would seem that she ought not hesitate to play. Disturbing as it is, that is an unavoidable consequence of the view I have developed. There is a further bizarre consequence. Suppose that Hydra were offered a round of qRr with guaranteed reward for the survivor but she were told that she must wear a blindfold. Also, when the readout has come up on the quantum die HydraUP and HydraDOWN will hear a bell and a little later HydraDOWN will be killed and HydraUP will get the reward. In that case Hydra can be certain that if she plays she will find herself in a state of great anxiety, being uncertain whether or not the quantum die in her downstream branch has shown DOWN or not. So even if Hydra cares for no-one but herself it will make a huge

[17] Greaves' and Myrvold's work would benefit from consideration of qRr. I think it is a mistake that qRr is often left out of discussion of the Everett interpretation as it can only help to clarify what is proving to be a very tortuous subject.

difference to her whether she is offered the round of qRr with or without the blindfold and delay mechanism or some equivalent variant of that setup.

Again, this discussion of Hydra's expectations about her future has nothing to do with the effectiveness of QMd in making physical predictions and providing a means for their verification. As we saw in §2, that is quite impersonal. In qRr contexts QMd assigns branch cross-sections to outcomes which are numerically equal to the Born rule values. The fatal branches exist. It is only when the foregoing analysis is applied to the personal experience of a subject that that subject's expectations part company with what would be demanded by conventional wave-collapse theory in qRr scenarios.

The anomalies simply indicate how physical theory may have surprising consequences for our understanding of the human predicament. David Lewis warned that acceptance of QMd could have terrifying implications for our conception of death [2004]. I believe that he made a plausible case, but one cannot reject a physical theory just because it seems unpalatable, as Lewis himself of course recognised. However, these anomalies might be avoided if there were available a concept of pre-measurement uncertainty for QMd. The argument here has only been that pre-measurement uncertainty is not necessary for the acceptability of QMd as a scientific theory so long as a satisfactory explanation of the Born-Vaidman rule is available. Bizarre consequences do not make a theory unscientific.

Recall Max Born's words in the opening quote: 'somewhere in our doctrine is hidden a concept, unjustified by experience, which we must eliminate to open up the road'. Ironically, the Born rule ushered in the idea that our observed environment evolves stochastically in quantum measurement contexts. That is not justified by experience.

Acknowledgements

My thanks to two anonymous referees, David Albert, David Baker, David Deutsch, Hilary Greaves, Peter J. Lewis, Mark Moyer, Wayne Myrvold, Avinash Puri, Simon Saunders and Lev Vaidman for valuable comments. I am particularly grateful to Víctor Durà-Vilà for a good deal of ongoing advice, much of which I have followed.

References

Albert, D. Z. [2010]: 'Probability in the Everett picture'. In *Many worlds: the Everett interpretation of quantum mechanics.* Saunders, S., Barrett, J., Kent, D., Wallace, D. (*eds.*), Oxford: Oxford University Press.

Baker, D. [2007]: 'Measurement Outcomes and Probability in Everettian Quantum Mechanics', *Studies in the History and Philosophy of Modern Physics* **38**, pp. 153-69.

Born, M. [1954]: Nobel lecture. <nobelprize.org/nobel_prizes/physics/laureates/1954/born-lecture.pdf>

Deutsch, D. [1985]: 'Quantum theory as a universal physical theory', *International Journal of Theoretical Physics* **24**, pp. 1-41.

Elga, A. [2000]: 'Self-locating belief and the sleeping beauty problem', *Analysis* **60**, pp. 143-7.

Everett, H. [1957]: '"Relative state" formulation of quantum mechanics', *Reviews of Modern Physics* **29**, pp. 454–62.

Frigg, R. [forthcoming]: 'Probability in Boltzmannian Statistical Mechanics'. In *Time, Chance and Reduction: Philosophical Aspects of Statistical Mechanics*. Ernst, G. and Hüttemann (*eds.*), Cambridge: Cambridge University Press.

[2008]: 'Chance in Boltzmannian Statistical Mechanics', *Philosophy of Science* **75**, pp. 670-81.

Gallois, A. [1998]: *Occasions of Identity.* Oxford: Clarendon Press.

Greaves, H. [2004]: 'Understanding Deutsch's probability in a deterministic multiverse', *Studies in History and Philosophy of Modern Physics* **35**, pp. 423-56.

[2007a]: 'Probability in the Everett interpretation', *Philosophy Compass* **2**, pp. 109-28.

[2007b]: 'On the Everettian epistemic problem', *Studies in History and Philosophy of Modern Physics* **38**, pp. 120-52.

Greaves, H. and Myrvold, W. [2010]: 'Everett and Evidence'. In *Many worlds: the Everett interpretation of quantum mechanics.* Saunders, S., Barrett, J., Kent, D., Wallace, D. (*eds.*), Oxford: Oxford University Press.

Hogan, C. J. [2007]: 'Quarks, electrons and atoms in closely related universes'. In *Universe or Multiverse?* Carr, B. (*ed.*), Cambridge: Cambridge University Press.

Ismael, J. [forthcoming]: 'Probability in Deterministic Physics', *Journal of Philosophy*.

Lewis, D. K. [1968]: Counterpart Theory and Quantified Modal Logic. *Journal of Philosophy, 67*, pp. 427-46.

[1980]: 'A Subjectivist's Guide to Objective Chance'. In R. C. Jeffrey (ed.), *Studies in Inductive Logic and Probability,* Vol. 2, pp. 263-93. Berkeley: University of California Press.

[2004]: 'How many Lives has Schrödinger's Cat?', *Australasian Journal of Philosophy* **82**, pp. 3-22. (Originally presented as the 2001 Jack Smart Lecture). Also in *Lewisian Themes*, Jackson, Frank and Priest, Graham (*eds.*), Oxford: Oxford University Press, 2004.

Lockwood, M. [1989]: *Mind, Brain &the Quantum*, Oxford, Blackwell.

Loewer, B. [2004]: 'David Lewis's Humean Theory of Objective Chance', *Philosophy of Science* **71**, pp. 1115-25.

Maudlin, T. [2010]: 'Can the world be only wave-function?'. In *Many worlds: the Everett interpretation of quantum mechanics.* Saunders, S., Barrett, J., Kent, D., Wallace, D. (*eds.*), Oxford: Oxford University Press.

Papineau, D. [1996]: 'Many Minds are No Worse than One', *British Journal for the Philosophy of Science* **47**, pp. 233-41.

Perry, J. [1972]: 'Can the self divide?', *Journal of Philosophy* **69**: pp. 463-88.

Price, H. [2010]: 'Decisions, Decisions, Decisions: Can Savage Salvage Everettian Probability?'. In *Many worlds: the Everett interpretation of quantum mechanics.* Saunders, S., Barrett, J., Kent, D., Wallace, D. (eds.), Oxford: Oxford University Press.

Saunders, S. and Wallace, D. [2008]: 'Branching and Uncertainty'. *British Journal for the Philosophy of Science,* **59**, pp. 293-305.

Saunders, S. [1997]: 'Naturalizing Metaphysics', *The Monist* **80**, pp. 44-69.

[2010]: ‘Chance in the Everett interpretation’. In *Many worlds: the Everett interpretation of quantum mechanics.* Saunders, S., Barrett, J., Kent, D., Wallace, D. (eds.), Oxford: Oxford University Press.

Sider, T. [1996]: ‘All the world's a stage’, *Australasian Journal of Philosophy* **74**, pp. 433-53.

[2001]: *Four Dimensionalism*. Oxford: Oxford University Press.

Tappenden, P. [2004]: ‘The ins and outs of Schrödinger’s cat box; a response to Papineau’, *Analysis* **64**, pp. 157-64.

[2008]: ‘Saunders and Wallace on Everett and Lewis’, *British Journal for the Philosophy of Science* **59**, pp. 307-14.

Vaidman, L. [1998]: ‘On Schizophrenic Experiences of the Neutron or Why We Should Believe in the Many-Worlds Interpretation of Quantum Theory’, *International Studies in the Philosophy of Science* **12**, pp. 245-66.

[2002]: ‘The Many-Worlds Interpretation of Quantum Mechanics’, *The Stanford Encyclopedia of Philosophy,* Zalta, E. N. (*ed.*).

van Fraassen, B. [1995]: ‘Belief and the problem of Ulysses and the sirens’, *Philosophical Studies* **77**, pp. 7 - 37.

Wallace, D. [2002]: ‘Quantum Probability and Decision Theory, Revisited’.

[2006]: ‘Epistemology quantized: circumstances in which we should come to believe in the Everett interpretation’, *British Journal for the Philosophy of Science* **57**, 655-89.

[2008]: ‘The Quantum Measurement Problem: State of Play’. In *The Ashgate Companion to Contemporary Philosophy of Physics*, Rickles, D. (*ed.*) London: Ashgate.

[2010]: ‘How to prove the Born rule’. In *Many worlds: the Everett interpretation of quantum mechanics.* Saunders, S., Barrett, J., Kent, D., Wallace, D. (*eds*.), Oxford: Oxford University Press.

Weinberg, S. [1992]: *Dreams of a Final Theory,* New York: Pantheon.