

Uncertainty and probability for branching selves

Peter J. Lewis

Department of Philosophy, University of Miami, Coral Gables, FL 33124, USA

E-mail address: plewis@miami.edu

Abstract: Everettian accounts of quantum mechanics entail that people branch; every possible result of a measurement actually occurs, and I have one successor for each result. Is there room for probability in such an account? The *prima facie* answer is no; there are no ontic chances here, and no ignorance about what will happen. But since any adequate quantum mechanical theory must make probabilistic predictions, much recent philosophical labor has gone into trying to construct an account of probability for branching selves. One popular strategy involves arguing that branching selves introduce a new kind of subjective uncertainty. I argue here that the variants of this strategy in the literature all fail, either because the uncertainty is spurious, or because it is in the wrong place to yield probabilistic predictions. I conclude that uncertainty cannot be the ground for probability in Everettian quantum mechanics.

Keywords: Quantum mechanics, Everett, many worlds, uncertainty, probability, personal identity.

1. Introduction: The problem

The measurement problem in quantum mechanics arises because the mathematical heart of the theory, Schrödinger's wave mechanics, apparently entails that every possible result of a measurement actually occurs, yet we experience only one of the possible results. Various additions and modifications to Schrödinger's wave mechanics have been proposed to deal with

this problem, but none of them can claim much success.¹ Perhaps the most audacious proposal is Everett's (1957) suggestion that we simply bite the bullet; every possible measurement result *does* actually occur. Then why does it appear to me that only one result has occurred? Everett's answer is that people *branch*; rather than having one successor, I have many, and each of my successors sees exactly one of the possible results.²

A nice feature of Everett's theory is that it can arguably be motivated from within the wave mechanics, without any additions or changes. That is, if people are treated as quantum mechanical systems, then given a particular interpretation of the wave mechanics, it follows that people have branching histories, it follows that there is a single measurement result in each branch, and it follows that the branches evolve independently, so that my successors in different branches will not be aware of each other.³ This feature, among others, makes Everett's proposal an attractive solution to the measurement problem.

The Achilles' heel of Everett's approach is probability. Our experience of the results of quantum mechanical measurements is not simply that exactly one result occurs, but also that the various possible results have different probabilities of occurring. The standard quantum mechanical recipe for calculating probabilities is the Born rule; the probability of obtaining a particular measurement result is equal to the squared amplitude of the wave term corresponding to that result. But if every possible measurement result occurs, as Everett's theory maintains, then it looks like each of them should be ascribed a probability of 1, irrespective of the wave amplitude. There are no ontic chances in Everett's theory—the wave mechanics is entirely deterministic—so probabilities (other than 0 or 1) cannot enter that way. And a person could in

¹ See Albert (1992) for an overview of the options.

² This is not the only way of interpreting Everett's theory. See Barrett (1999) for an overview.

³ These results are not trivial, and depend on a solution to the preferred basis problem. See Barrett (1999, chapter 8) for a presentation of the difficulties facing this approach, and Wallace (2002a) for an attempt to solve them.

principle know the wave state in as much detail as she likes, so probabilities cannot enter via ignorance either. But that apparently exhausts all the options.

One reason this a *problem* is that we want our scientific theories, quantum mechanics included, to act as a guide to action.⁴ Standard quantum mechanics yields probabilities for various future occurrences, and these probabilities can be fed into an appropriate decision theory. But if every physically possible consequence of the current state of affairs is certain to occur, on what basis should I decide what to do? For example, if I point a gun at my head and pull the trigger, it looks like Everett's theory entails that I am certain to survive—and that I am certain to die. This is at least worrying, and perhaps rationally disabling.

The problem, then, is how, if at all, the Born rule can be accommodated within Everett's theory. This problem has two aspects (Wallace 2003, 417). The first is the incoherence problem: If every outcome actually occurs, how can it even make sense to ascribe probabilities (other than 0 and 1) to measurement results? The second is the quantitative problem: Assuming that the ascription of probabilities to measurement results makes sense, why should the probabilities be those given by the Born rule? In this paper, I have nothing to say about the quantitative problem; the most influential recent approach is that of Deutsch (1999). Rather, my concern is entirely with the (logically prior) incoherence problem.

At least at first glance, it looks like the way to solve the incoherence problem is to find some hidden source of *uncertainty* in Everett's theory, so that an agent may be genuinely uncertain what will happen to her, despite the fact that (objectively speaking) every outcome occurs.⁵ Various suggestions have been made along these lines in the literature, and I survey

⁴ A related problem concerns confirmation; if Everett's theory entails *every* possible sequence of measurement results, then no such sequence can confirm it. I have nothing directly to say about this problem here.

⁵ There is an alternative strategy, namely to argue that a probability measure is meaningful even in the absence of subjective uncertainty (Papineau 2004; Greaves 2004). I discuss this strategy briefly in sections 4 and 6 below.

them in the following sections. My contention is that they either fail to locate a genuine source of uncertainty, or locate the uncertainty in the wrong place to ground the required probability measure.

To make my case, I will need to refer to three prominent accounts of personal identity in branching situations. According to the first account, a person is a three-dimensional object that persists over time. There is one person prior to the branching event, and there are two persons after the branching event. The pre-branching person is not identical to either post-branching person (since they are not identical to each other), but is psychologically continuous with each of them. Call this the Parfitian account (Parfit 1984). According to the second account, a person is a four-dimensional entity—a space-time worm. There are two persons present both before and after branching; they coincide before the branching event, but diverge thereafter. Hence each pre-branching person is identical to exactly one post-branching person. Call this the Lewisian account (Lewis 1983). According to the third account, a person is a temporal stage of a four-dimensional entity—a time-slice of a space-time worm. There is one person present prior to branching and two persons present afterwards. Strict *identity* does not hold between these persons, since they are distinct temporal stages, but nevertheless the pre-measurement person has two post-measurement temporal counterparts, and the existence of such counterparts makes it the case that the pre-measurement person *will be* each of the post-measurement persons. Call this the Siderian account (Sider 1996).

2. Saunders' argument

Consider an observer watching a spin measurement, with possible outcomes 'up' and 'down'.

Let t_1 be a time just before the measurement, and let she_1 be the observer at t_1 . Let t_2 be a time just after the measurement, and let she_2^\uparrow and she_2^\downarrow be the observer's two successors at t_2 , seeing

outcome ‘up’ and ‘down’ respectively. Who can she_1 expect to become at t_2 ? Saunders (1998) argues that there are three options: (i) she_1 can expect nothing (i.e. oblivion), since she_1 no longer exists at t_2 , (ii) she_1 can expect to become both she_2^\uparrow and she_2^\downarrow , and (iii) she_1 can expect to become one of she_2^\uparrow and she_2^\downarrow (1998, 383). The first option he dismisses as implausible; certainly an expectation of oblivion does not follow from the fact that she_1 no longer exists at t_2 , for such reasoning would apply equally well to ordinary, non-branching persons. The second option he dismisses as “straightforwardly inconsistent”, since “ she_1^\uparrow and she_2^\downarrow do not speak in unison; they do not share a single mind; they witness different events” (1998, 383). So there remains only the third option; she_1 can expect to become one of she_2^\uparrow and she_2^\downarrow , but not both. But this is precisely what it means for she_1 to be *uncertain* about who she will become.

At first glance, Saunders’ three options certainly look exhaustive; ‘neither’, ‘both’ and ‘one or the other’ seem to exhaust the logical space. But consideration of an analogy with a more familiar branching object suggests that things are not so straightforward.⁶ Consider a road that forks at some point, with one fork going to Upton and the other going to Downham. The pre-fork segment of the road is physically continuous with each of the post-fork segments, just as the pre-measurement person in Saunders’ example is physically (and psychologically) continuous with her two post-measurement successors. Let $road_1$ be a road-segment at x -coordinate x_1 , before the fork, and let $road_2^\uparrow$ and $road_2^\downarrow$ be road-segments at x_2 , after the fork, on the branches going to Upton and Downham respectively. The analogous question in the road case is this: Which road (if any) does $road_1$ become at x_2 ? (Note that the question is not “Where will I get to if I drive along the road to x_2 ?”; the analogy is between the road itself and the Everettian person.) And the answers analogous to Saunders’ three options are: (i) neither $road_2^\uparrow$ nor $road_2^\downarrow$, (ii) both $road_2^\uparrow$

⁶ This analogy is suggested by Lewis (1983, 64).

and road_2^\downarrow , and (iii) one of road_2^\uparrow and road_2^\downarrow . Clearly the answer isn't 'neither'; the road doesn't disappear between x_1 and x_2 . And option (ii) is (on the face of it) inconsistent; road_2^\uparrow and road_2^\downarrow have different properties, in particular their respective destinations, so no road can become *both* of them. So, by Saunders' reasoning, we are left with option (iii); road_1 becomes one of road_2^\uparrow and road_2^\downarrow , but not both.

But this conclusion is misleading at best. In effect, the suggested conclusion is that road_1 goes to either Upton or Downham, but not both. If there is a sense in which that is true, it is that I could *take* road_1 and get to either Upton or Downham, but not both. But as noted above, the question of where I could get to along the road is not analogous to the question of what becomes of the Everettian person; rather, the relevant question is where the *road* goes. Furthermore, the conclusion in the road case does not license any inference to *uncertainty*; I am not uncertain about where the road goes, however I may choose to describe it. What we would say in such a case, I think, is the following: road_1 doesn't end up at both Upton and Downham *at once*—that would indeed be contradictory—but rather it ends up at *each* of Upton and Downham, on separate branches. This extra option is available precisely because the road splits, but it doesn't entail any uncertainty. The same, it seems to me, can be said for the case of Everettian branching. she_1 expects to become *each* of she_2^\uparrow and she_2^\downarrow , on separate branches. The fact that there are two successors rather than one means that there is no contradiction between these expectations, even though she_2^\uparrow and she_2^\downarrow have different properties, and even though she_1 expects each *fully* (i.e. with no uncertainty).

This is exactly the conclusion reached by Greaves (2004, 440). Greaves appeals to an analysis of personal identity in branching situations which is essentially the Siderian account sketched above. According to this account, she_1 will be she_2^\uparrow , and she_1 will be she_2^\downarrow ; hence if

she₁ is informed concerning her situation, she knows that she will see ‘up’ and she knows that she will see ‘down’, and there is nothing for her to be uncertain about. However, this conclusion is challenged by Wallace (2005, 11). Wallace contends that Greaves’ argument misses its mark, since Greaves assumes a different account of personal identity from that implicit in Saunders’ argument. Wallace takes Greaves to assume a Parfitian account, i.e. that the concept of personal identity breaks down in cases of branching. On the other hand, Saunders (according to Wallace) assumes that the relevant relation between the person-stages at t_1 and t_2 is precisely that of identity; she₁ can expect to become she₂[↑] only if she₁ is identical to she₂[↑]. Hence his argument against option (ii) above: If she₁ were identical to both she₂[↑] and she₂[↓], then she₂[↑] and she₂[↓] would have to be identical to each other (by the transitivity of identity), and hence would have to have all their properties in common. The same is clearly true if she₁ is identical to *each* of she₂[↑] and she₂[↓], and hence Saunders cannot accept the position endorsed by Greaves (and myself), namely that she₁ expects to become each of she₂[↑] and she₂[↓].

I think that Wallace slightly misstates his case, since Greaves’ account of personal identity is closer to Sider’s than to Parfit’s, and hence Greaves’ argument doesn’t obviously rely on the failure of the concept of personal identity.⁷ On the Siderian account, a personal-identity-over-time relation (a symmetric analog of the will-be relation) holds between she₁ and she₂[↑], and between she₁ and she₂[↓]. Nevertheless, the Siderian personal-identity-over-time relation is not transitive; she₂[↑] does not bear the relation to she₂[↓], as Greaves notes (2004, 439). So if it the assumption of transitivity that is the key to Saunders’ argument, then Wallace is right that Greaves rejects it. Furthermore, Wallace is right that there is an account of personal identity in

⁷ It is worth noting that Greaves’ argument would go through under a Parfitian account of personal identity. Under such an account, she₁ is not identical to either she₂[↑] or she₂[↓], but she₁ survives through each of them, and again there is nothing about which she₁ is uncertain.

branching situations which preserves transitivity of identity, namely the Lewisian account. Since Greaves does not discuss this possibility, her argument against Saunders remains inconclusive. So before drawing any firm conclusions, we must first examine Wallace's explicitly Lewisian reconstruction of Saunders' argument.

3. Wallace's Lewisian argument

Wallace (2005, 13) claims that the Lewisian account of personal identity can be used to defend Saunders' account of uncertainty in Everettian branching. According to the Lewisian account, there are two persons present at t_1 , one of whom will become she_2^\uparrow and one of whom will become she_2^\downarrow , but since the two persons are physically identical up to the measurement, neither of them can discover which one she will become, except by waiting to find out. Hence Saunders' option (iii) is, after all, the right way to describe each of the two pre-measurement persons, and there can be genuine uncertainty about the future in an Everettian universe.

I agree that the Lewisian account of personal identity is tenable, and I agree that it yields a kind of indeterminacy about who an Everettian observer will become, but I don't think this indeterminacy underwrites any kind of uncertainty. Consider again the road analogy. It is certainly possible to identify a road with its whole *course*, rather than with a road-segment. On this view, there are two roads even at x_1 , one of which ultimately goes to Upton and the other of which ultimately goes to Downham. These two roads have all physical properties in common up to the fork, but thereafter they have different physical properties. Indeed, this way of individuating roads is sometimes reflected in road numbering systems; for example, it could be that the 15 goes to Upton and the 501 goes to Downham, but they coincide for a few miles

around x_1 .⁸ In that case, the question “Where does this road go?”, when uttered at x_1 , may have no determinate answer. An appropriate reply might be “Do you mean the 15 or the 501? The 15 goes to Upton and the 501 goes to Downham.” But notice that this indeterminacy doesn’t yield any uncertainty. Once I am apprised of all the relevant facts, it would be very strange for me to *still* insist that I don’t know where the road goes. In particular, I cannot reasonably harbor uncertainty as to whether my use of ‘this road’ at x_1 refers to the 15 or to the 501, since my demonstrative indicates certain physical features of the world at x_1 , and those are common to the 15 and the 501.

Furthermore, I don’t think that the identification of a road with its whole course is privileged over other ways we identify roads. We have other conventions for naming roads that suggest that sometimes we identify a road with a road segment. That is, the segment of road around x_1 might be called *High Street*. Strictly speaking, High Street doesn’t go to either Upton or Downham, since High Street ends well before x_2 (and let us assume, before the fork). Still, I might reasonably ask where High Street leads to, and be told that it leads to a fork in the road, beyond which the right fork leads to Upton and the left fork leads to Downham. This answer would, if I accept it, eliminate my uncertainty. Notice that this system does not identify either the road to Upton or the Road to Downham with High Street; rather, High Street bears a certain relation to them. That is, the system of local names gives up on strict *identity* between road-segments, and makes do with *leading to* relations between them instead. The local name system for individuating roads in a sense contradicts the road number system of the prior paragraph; according to the number system there are two roads at x_1 , and according to the local name system there is just one. But the contradiction is superficial; it is not that there is some *true* number of

⁸ At least, the conventions for road naming in the US work this way. In other countries, each road-segment has exactly one number, even if this means that some road-numbers refer to disjoint entities.

roads at x_1 that we must capture, but rather that the number of roads at x_1 depends on how we choose to count them.

I maintain that parallel considerations apply to Everettian persons. We can, if we wish, adopt the Lewis criterion for personal identity in Everettian contexts, in which case there are (in that sense) two persons prior to measurement. We could even give these two persons names, say she^\uparrow and she^\downarrow , so that she^\uparrow refers to the person who sees ‘up’ and she^\downarrow refers to the person who sees ‘down’. If I walk into the lab and ask “What result is she about to see?”, I might be told “ she^\uparrow will see ‘up’ and she^\downarrow will see ‘down’, and at the moment she^\uparrow and she^\downarrow coincide”. That ought to eliminate my uncertainty, provided that I understand and accept Everett’s theory. In particular, I can’t wonder further whether my use of the pronoun ‘she’ when pointing at the observer picks out she^\uparrow or she^\downarrow ; since she^\uparrow and she^\downarrow coincide at the moment, I am pointing at both of them.

One might object here that I *can* wonder whether my use of the pronoun ‘she’ when pointing at the observer picks out she^\uparrow or she^\downarrow , since under Wallace’s Lewisian proposal there are *two* persons doing the pointing, me^\uparrow and me^\downarrow . That is, the branching of the observer produces two copies not only of the observer herself, but of everything that comes in contact with her. If that is the case, then I can legitimately wonder which person I am— me^\uparrow pointing at she^\uparrow , or me^\downarrow pointing at she^\downarrow . But even if there are two persons doing the pointing and two persons being pointed at, I don’t think it follows that I am uncertain who I am pointing at. As Tappenden (2006) points out, there is a single *act* of pointing, since temporally local events are not individuated as persons are. This act is performed by a person-stage that is common to me^\uparrow and me^\downarrow , and it points at a person-stage that is common to she^\uparrow and she^\downarrow , but there is no further fact concerning which of me^\uparrow and me^\downarrow is doing the pointing and which of she^\uparrow and she^\downarrow is being

pointed at. The single act of pointing is performed by *both* me^\uparrow and me^\downarrow , and is directed at *both* she^\uparrow and she^\downarrow .

Alternatively, one might object that the uncertainty is only available from a first-person perspective; she^\uparrow and she^\downarrow can each ask “What will I see?”, and the uncertainty arises because the ‘I’ in one person’s utterance refers to she^\uparrow and in the other person’s utterance refers to she^\downarrow , and neither of them knows which person her use of ‘I’ refers to (Wallace 2005, 12–13). But again, there is a single utterance here, made by a person-stage that is common to she^\uparrow and she^\downarrow , and directed at a person-stage that is common to she^\uparrow and she^\downarrow . There is no further fact concerning which of she^\uparrow and she^\downarrow makes the utterance, and which of them it is directed at; it is made by them both, and directed at them both, and hence there is no room for uncertainty.⁹

To make the same point in a slightly different (although perhaps more controversial) way, I am not convinced that the pronoun ‘I’ picks out a *person* in any deep metaphysical sense, any more than the pronoun ‘she’ does, or any more than the demonstrative phrase ‘this road’ picks out a particular numbered road. Recall that in the case of roads, the number system was one way of individuating roads, but not the only way. We can also individuate roads by local name, in which case there is only one road at x_1 . One system may be more convenient than the other for various purposes, but neither has a better claim to embody the way roads really are. So it is, I contend, with persons. We can individuate persons in the Lewis way, in which case there are two persons at t_1 , or we can individuate persons in the Parfitian or Siderian way, in which case there is one person at t_1 . For a reductivist about persons, I see no reason to regard one of these as embodying the unique correct account of persons. But recall from the previous section that

⁹ See Tappenden (2006) for the details of this argument. Lewis apparently concurs; in the case of two persons C_1 and C_2 , who currently coincide but later diverge, he writes that “the ‘me’ in their shared thought ... cannot refer to C_1 in C_1 ’s thought and C_2 in C_2 ’s thought, for these thoughts are one and the same” (1983, 75).

neither the Parfitian nor the Siderian account yields any uncertainty about what she₁ will see. So it should strike us as suspicious, I think, that Wallace's Lewisian account apparently does yield uncertainty. After all, there is no deep truth about personal identity for the pronoun 'I' to latch on to. In particular, it is not a deep truth that there are two persons at t_1 , but merely a convenient way of tracking branching entities across time. Hence I think there can be no fact about which of the two persons 'I' really refers to. If one chooses to regard the situation at t_1 as one in which there are two persons, then 'I' refers to them both, and there can be no uncertainty about which of them is really me.

Let us take stock. Saunders argues that there is genuine uncertainty arising from Everettian branching. Greaves maintains that Saunders is wrong, basing her case on a Siderian account of persons. Wallace suggests that an alternative (Lewisian) account of persons is available, according to which there is genuine personal identity over time, and within which Saunders position on uncertainty can be defended. My argument here is twofold. First, neither Sider nor Lewis can claim the whole truth about persons; both are individuation schemes with features to recommend them, but neither has a better handle on underlying reality. Second, neither account yields genuine uncertainty in cases of Everettian branching. Wallace's Lewisian account yields an appearance of uncertainty, but this appearance is dissolved by a careful treatment of reference.

4. Wallace's semantic argument

Despite his defense of Saunders' account of uncertainty, Wallace doesn't regard this as an adequate solution to the problem of uncertainty. The reason is that he thinks the question of how we would *extend* our identity concepts were we to discover cases of branching is somewhat beside the point. Rather, the question is how we should understand our current concepts, given

that we in fact live in an Everettian universe (2005, 14). That is, given that we are living in an Everettian universe, unbeknownst to us, what does ‘uncertainty’ mean *now*?

To that end, Wallace asks us to consider a race of beings who inhabit an Everettian universe, but don’t know it (2005, 14). When confronted with (what is in fact) a branching event, these beings are disposed to say “I am *uncertain whether A will happen” when A happens in some but not all branches, where ‘*uncertain’ is a term in their language. Their philosophers, ignorant of the true branching nature of the universe, claim that one should only be *uncertain of something if there is some fact about which to be *uncertain. Wallace asks: What is the real meaning of ‘*uncertain’? He distinguishes two views on the matter, the Elite View and the Charitable View. The Elite View is that the philosophers are right, but this has the consequence that the beings say many false or truth-valueless things. They say “I am *uncertain whether A will happen” when there is no relevant fact about which the beings could be uncertain; in a branching world, there is no fact about whether A occurs, or whether the relevant being will see A, and so on. So either their *uncertainty claims are typically false, or they rely on a false presupposition—namely that there is a single unique future—and hence lack a truth value. The Charitable View, on the other hand, is that the beings are generally using ‘*uncertainty’ correctly, and that it is the philosophers who are wrong. That is, ‘*uncertainty’ doesn’t mean what the philosophers think it means, and in particular, *uncertainty doesn’t require a fact about which one is *uncertain. Rather, “I am *uncertain whether A will happen” is true just in case A is true in some but not all of the future branches relative to the utterer.

Wallace thinks the Charitable View is clearly correct, for the Davidsonian reason that the best interpretation is the one that makes most of the community’s utterances true (2005, 15). So it follows that if, unbeknownst to us, we are living in a branching universe, then the Charitable

View is the best interpretation of our utterances, too. That is, if we are living in a branching universe, then uncertainty has never required some fact about which one is uncertain, but rather, “I am uncertain whether A will happen” has always meant that A will happen in some but not all future branches. Hence our claims to be uncertain about whether various future events will occur are true even if this world is a branching one, and no fancy metaphysics of personal identity is required to find a place for uncertainty in an Everettian universe.

Despite these arguments, though, I think there is something to be said for the Elite View. In Wallace’s imagined world, the philosophers’ insistence that *uncertainty requires a fact about which one is *uncertain is presumably not arbitrary; certainly the insistence on this point by many actual philosophers in our world (e.g. Greaves 2004, 441) is not arbitrary. Rather, the Elite View expresses part of a web of conceptual connections, linking our concept of uncertainty with other concepts. That is, we think of uncertainty claims as expressing a lack of knowledge; if I am uncertain whether A will occur, then there is some fact about the future that I am ignorant of. Calling this view ‘elite’ is something of a misnomer; the Elite View is not an ivory-tower invention, but is built into the way we (the masses) deploy the term ‘uncertain’. Similarly, calling the opposing view ‘charitable’ is somewhat misleading; it is charitable regarding the truth of our claims, but not regarding the conceptual connections between them. The Elite View may have the consequence that most of our uncertainty claims are false or truth-valueless, but the Charitable View has the consequence that our uncertainty claims do not carry the implications that we ordinarily take them to have. In other words, when we reinterpret our current linguistic practices under the assumption that we are living in a branching world, we have to give up *something*—either the truth of some of our claims, or the conceptual connections between them.

Hence it is not so clear whether the Elite View or the Charitable View is to be generally preferred in situations like this. Historical precedents go either way. On the one hand, the fact that tables and such are mostly empty space did not lead us to interpret our prior claims that tables are solid as false or truth-valueless. Rather, we came to regard as mistaken certain connections between solidity and other concepts, such as impenetrability; a material particle could, after all, traverse a solid object. On the other hand, the null result of the Michelson-Morley experiment did lead us to interpret our prior claims that light is a vibration in the ether as false or truth-valueless. We chose instead to retain the conceptual connections according to which the ether was regarded as a material medium. Furthermore, neither of these precedents is particularly clear-cut; Eddington (1928, xiv) famously argues that our belief in the solidity of tables is really false, and Kitcher (1993, 147) suggests that historical claims that light is a vibration in the ether could be interpreted as true. My suspicion is that there is no pressing need to decide these matters; the Elite View and the Charitable View constitute pragmatic choices about how to use language, and we can make different choices in different contexts, provided we make our choice clear.

But I do not need to defend this view here, since the main point I wish to stress is that whether we adopt the Charitable View or the Elite View regarding uncertainty, the use of uncertainty to ground the concept of probability in Everett worlds is blocked. This is obviously the case for the Elite View, as Wallace notes; if our claims to uncertainty as to the results of measurements are not true, then that uncertainty cannot provide the basis for a probability measure over those results. But if we adopt the Charitable View, the conceptual connection between uncertainty and ignorance of fact is lost, and it is arguably via this conceptual connection that uncertainty is related to (subjective) probability. A subjective probability

measure is, after all, a measure of degree of belief, so unless uncertainty entails ignorance about the measurement result, no probability other than 1 can be assigned to any result. This may yield a probability measure in a trivial sense, but it is certainly not the Born rule.

To clarify, consider the road analogy again. Suppose we discover that people branch in much the same way that roads branch. How are we likely to interpret our prior uncertainty claims concerning what will happen to us? When we're in full possession of the facts, we don't say that we're uncertain where the road goes, so we might be inclined to dismiss our prior uncertainty claims as false. But as noted in the previous section, we might notice that there is a sense in which we can regard the destination of a forking road as indeterminate, since the road has no unique destination. In that case, we might try to recover the truth of our uncertainty claims by assimilating them with indeterminacy claims; when I said "I am uncertain what will happen to me", what I said was true, and what I meant was something close to "It is indeterminate what will happen to me". But whichever option we take (and I suspect we might take both, on different occasions), there is no role for probability here. First, there is a perfectly good interpretation of our prior uncertainty claims under which they are false, so there is a perfectly straightforward sense in which probability is inapplicable. Second, even if we choose to interpret uncertainty claims as true, we do so by assimilating uncertainty to indeterminacy, and this severs the connection between uncertainty and ignorance of fact that could justify the application of a probability measure.

The conceptual connection between uncertainty and ignorance could be recovered by adopting a Charitable View of the meaning of *ignorance* claims too; "I am ignorant concerning whether A will happen" has *also* always meant that A will happen in some but not all future branches. But this just pushes the problem further back; now the severed connection is between

ignorance and any question concerning the actuality of A. As a last-ditch strategy, one might insist that an observer can attach a probability (other than 0 or 1) to A even in the absence of any question concerning the actuality of A. But now I wonder why we were looking for *uncertainty* in the first place, since what is doing the work here is a reconceptualization of probability. This latter strategy is that of Papineau (2004) and Greaves (2004); Greaves, for example, suggests that one think of a probability measure in an Everett world as a measure of how much an observer *cares* about her successors. I have nothing to say about this strategy here, except that it makes the search for uncertainty in Everettian quantum mechanics entirely beside the point.

Hence I don't think that Wallace's argument from interpretive charity succeeds. While it may be the case that we would (or could) interpret our current uncertainty claims as true even if it turns out that we live in an Everett world, the concept of uncertainty could not emerge unscathed from this transformation in our understanding. The reason that so many authors are trying to find a source of uncertainty in Everett's theory is that there is a *prima facie* connection between uncertainty and probability. But once we engage in the radical reinterpretation of our uncertainty claims that Wallace suggests, this *prima facie* connection is lost, and it is no longer clear how uncertainty is a route to probability.

5. Uncertainty after the fact

There is, however, a place for uncertainty about measurement results in Everett's theory that is quite uncontroversial—that doesn't rely on any special views about the metaphysics of persons or the semantics of 'uncertain'. This place—or rather, this *time*—is after the measurement is complete. Suppose an observer closes her eyes during a spin measurement. She ends up with two successors, each of whom can ask herself "I wonder whether I will see 'up' or 'down' when I open my eyes?". Here, then, is genuine, unproblematic uncertainty in an Everett world,

uncertainty that *does* involve ignorance about a matter of fact. However, this uncertainty may never be actualized in a particular case, for example if the observer keeps her eyes open and the measurement proceeds quickly enough (Wallace 2005, 21); this is something that we will have to bear in mind as we proceed.

The argument that this after-the-fact uncertainty provides the foundation for probability in Everett worlds is made by Vaidman (1998; 2002). Clearly Vaidman succeeds at finding *something* in Everett's theory that we can coherently quantify using a subjective probability measure. What is not so clear is whether this is the probability we were looking for. Recall that the reason it is important to find a place for probability in Everett's theory is that the use of quantum mechanics to make decisions seem to require a sense in which Everett's theory can make probabilistic predictions. But notice that Vaidman's after-the-fact probabilities are not predictive in any obvious sense; the observer may be uncertain about what *has* happened, but not about what *will* happen. That is, the uncertainty seems to be in the wrong place to play any role in decision.

Vaidman (2002) is aware of this difficulty, and suggests the following solution. First, he suggests that we *define* the observer's pre-measurement probability for a given outcome to be the post-measurement ignorance probability of her successors. This gives us the right numbers for the pre-measurement probabilities (assuming the quantitative problem can be solved), but it isn't clear what these numbers represent, since there is no pre-measurement uncertainty. As Vaidman (2002) admits, "since all outcomes of a quantum experiment are actualized, there is no probability in the usual sense" prior to the experiment. Nevertheless, he argues that the observer "should behave as if there were certain probabilities for different outcomes". That is, even though the numbers assigned to the various outcomes prior to measurement aren't *genuine*

subjective probabilities, they are *effective* probabilities, in the sense that the observer should behave as if she already had the degrees of belief that she can expect her post-measurement successors to have. Hence post-measurement uncertainty can play the required role in pre-measurement decision.

Greaves (2004, 442) also endorses this line of argument, noting that it rests on a particular kind of reflection principle.¹⁰ It is not an epistemic reflection principle; the post-measurement uncertainty does not give the pre-measurement observer reason to be uncertain. Rather, it is a decision-theoretic reflection principle; the pre-measurement observer should adopt the same decision strategy as her post-measurement successors. As noted above, there may be no time at which a particular observer is uncertain about a given measurement result, since she may be aware of the measurement result as soon as the measurement is performed, but I don't think this is an insurmountable problem. The reflection principle can simply be modified to say that the pre-measurement observer should adopt the decision strategy that the post-measurement observer *would* have had, had she not learned the result.

But I think the road analogy shows that any such decision-theoretic reflection principle is inappropriate here. Suppose I am traveling down the road as a passenger in a car, and I close my eyes, so that I don't know which way the car goes at the fork. After the fork, I am genuinely uncertain whether this road is the road to Upton or the road to Downham. This uncertainty may ground a subjective probability measure; I may judge that it is twice as likely that this is the road to Upton as that this is the road to Downham (for whatever reason). This subjective probability assignment may in turn lead to certain behaviors or strategies; I may be prepared to bet at 2:1

¹⁰ Greaves attributes the reflection principle to Wallace (2002b, 58). She endorses the reflection argument in the sense that she thinks it constitutes a successful solution to the incoherence problem. However, it is not her preferred solution, since she feels that it "obscures the real logic of the argument" (2004, 443). Her preferred solution is to adopt the caring account of probability mentioned at the end of the previous section.

odds that this is the road to Downham. But now consider my behavior before the fork. Should I behave as if I judge that the chance that this is the road to Downham is $1/3$ at that point as well? Surely not; I *know* that this is the road to Downham, just as I know that it is the road to Upton. If I am offered a bet that returns \$1 for a \$10 stake if this is the road to Downham, then before the fork I should take it, and after the fork I should not.

Note, again, that the analogy is between the *road* and the Everettian person. In the road case, it may well be rationally incumbent on me to adopt the same probability that I am *going* to Downham before the fork as after it, assuming I learn nothing relevant along the way. But my trajectory along the road is not analogous to the history of an Everettian observer; I go one way or another, the Everettian observer branches. So the analogy is *not* between “Will I see ‘up’ or ‘down’?” in the Everettian case and “Will I go to Upton or Downham?” in the road case, but between “Will I see ‘up’ or ‘down’?” and “Does this road lead to Upton or Downham?”. And given that this is the right analogy, then the application of a decision-theoretic reflection principle is equally inappropriate in each case. While there is genuine after-the-fact uncertainty, it has no obvious relevance to before-the-fact behavior.

Ismael (2003) also attempt to ground probability in post-measurement uncertainty, but she stresses the indexical nature of the uncertainty. Consider an analogy with time. If I lose track of time during a talk, I become uncertain as to what time it is *now*. Analogously, if I close my eyes during a measurement, I lose track of what branch I occupy, and hence I become uncertain as to what the result of the measurement is *here*, where ‘here’ is not a spatial or a temporal indexical, but a branch indexical. Indeed, Ismael insists that the kind of uncertainty appropriate to Everett’s theory can only be expressed indexically.

The trick, as before, is to parlay the acknowledged post-measurement uncertainty into a pre-measurement uncertainty, and this is what Ismael attempts. She argues that there is a sense in which I can legitimately wonder what the result of the measurement *will* be (2003, 778). The difficulty, she explains, is that before the measurement I have no way to express what I am wondering about. After the measurement, there is no such difficulty; I can say “I wonder what the result is *here*”. But before the measurement, I cannot say “I wonder what the result will be *there*” without specifying *where* in the branching structure I mean to refer to. The only means at my disposal for describing a location in the branching structure is via the measurement result in that branch; I can use a locution such as “the branch in which the result is ‘up’”. But this is clearly no help at all; I can’t be uncertain what the result will be in the branch in which the result is ‘up’.

Ismael contends that this is just a linguistic peculiarity, “and that the way to see that it is a linguistic peculiarity ... is that if worlds had names, or came in different flavors (so that ‘the chocolate world’ or ‘the grape world’ were identifying descriptions) there would be no problem about interpreting the Born probabilities” (2003, 781). In other words, there are *facts* that I am uncertain of before the measurement, facts concerning which branch (or ‘world’ in Ismael’s terminology) contains which measurement result. If branches had flavors, then such a fact could be expressed as “The grape-flavored branch contains the ‘up’ result”. But because the only way I can refer to a future branch is via the measurement results it contains, these facts are inexpressible in advance of the measurement. Inexpressible facts, though, are still facts, and I can still be uncertain about what they are.

I agree that there are post-measurement facts of which an observer may be genuinely uncertain, but I don’t think Ismael’s characterization of these facts—as facts about which branch

contains which result—is appropriate. Recall again that there is nothing in Everett’s theory over and above the quantum state. To say that the state has branches is just to say that it can be written as a sum of more-or-less independent terms, where each term is taken as a description of a state of affairs—say, the state of affairs in which the measurement result is ‘up’. So the state of affairs *is* the branch—it is not *contained in* the branch—and it makes no sense to conceive of the same state of affairs in a different branch. So there are no facts about which measurement result appears in which branch; whatever it is I am uncertain about, this is not it. Ismael is of course free to *postulate* the existence of ‘worlds’ as entities over and above the quantum mechanical state, but this would be to give up on the ‘nothing but the physics’ approach that she clearly endorses (2003, 779).

Then what is the fact that I am uncertain about, after the measurement? It is the fact that can only be expressed indexically, the fact that measurement result *here* is ‘up’. But the point about such facts, as argued above, is that I can’t wonder about them in advance. Just as I can’t wonder about what time it is until I lose track of time, so I can’t wonder about what branch this is until I lose track of my branch. Prior to the measurement, I know with certainty where I am and what will happen to me in the future; but I also know that if I close my eyes during the measurement, I will lose track of where I am in the branching structure, and hence become uncertain about the measurement result *here*. So while Ismael, like Vaidman, locates a genuine source of uncertainty in Everettian quantum mechanics, this uncertainty cannot ground the pre-measurement probability measure that we are after.

6. Conclusion

I have examined several arguments that inhabitants of Everett worlds may be uncertain about what will happen to them. None of these arguments, I think, succeed in locating the kind of

uncertainty that could ground non-trivial probability assignments to measurement outcomes. Arguments from personal identity fail to identify any genuine source of uncertainty. Arguments from interpretive charity can perhaps recover the truth of uncertainty claims in Everett worlds, but only at the expense of severing the connection between uncertainty and probability. Vaidman and Ismael locate a genuine source of after-the-fact uncertainty—inhabitants of Everett worlds may be uncertain about what *has* happened—but neither succeeds in connecting this uncertainty to a predictive probability assignment.

It is possible that there is some other source of uncertainty in Everett worlds that has so far been overlooked, but if the analogy between branching persons and branching roads developed above is appropriate, this seems unlikely. There is, however, an alternative strategy for grounding probability in Everett worlds that does not require any subjective uncertainty. That is, one postulates a relation between an observer and her successors, which might be described as how much she *cares* about each of them, and uses this relation to ground a probability assignment (Papineau 2004; Greaves 2004). Examining the tenability of this approach, however, will have to wait for another day.

Acknowledgements

I would like to thank Paul Tappenden, Jenann Ismael and two anonymous referees for helpful comments on earlier drafts of this paper.

References

Albert, David Z. (1992), *Quantum mechanics and experience*. Cambridge, MA: Harvard University Press.

- Barrett, Jeffrey A. (1999), *The quantum mechanics of minds and worlds*. Oxford: Oxford University Press.
- Deutsch, David (1999), “Quantum theory of probability and decisions”, *Proceedings of the Royal Society of London A* 455: 3129–3137.
- Eddington, Arthur (1928), *The nature of the physical world*. Cambridge: Cambridge University Press.
- Everett, Hugh III (1957), “‘Relative state’ formulation of quantum mechanics”, *Reviews of Modern Physics* 29: 454–462.
- Greaves, Hilary (2004), “Understanding Deutsch’s probability in a deterministic multiverse”, *Studies in History and Philosophy of Modern Physics* 35: 423–456.
- Ismael, Jenann (2003), “How to combine chance and determinism: thinking about the future in an Everett universe”, *Philosophy of Science* 70: 776–790.
- Kitcher, Philip (1993), *The advancement of science*. Oxford: Oxford University Press.
- Lewis, David (1983), “Survival and identity”, in *Philosophical Papers, Volume I*. Oxford: Oxford University Press, 55–77.
- Papineau, David (2004), “David Lewis and Schrödinger’s cat”, *Australasian Journal of Philosophy* 82: 153–169.
- Parfit, Derek (1984), *Reasons and persons*. Oxford: Oxford University Press.
- Saunders, Simon (1998), “Time, quantum mechanics and probability”, *Synthese* 114: 373–404.
- Sider, Theodore (1996), “All the world’s a stage”, *Australasian Journal of Philosophy* 74: 433–453.
- Tappenden, Paul (2006), “Everett branching and David Lewis’s analysis of personal fission”, <http://philsci-archive.pitt.edu/archive/0000????/>.

- Vaidman, Lev (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: 245–261.
- (2002), “Many-worlds interpretation of quantum mechanics”, in Edward Zalta (ed.), *The Stanford Encyclopedia of Philosophy*.
<http://plato.stanford.edu/archives/sum2002/entries/qm-manyworlds/>.
- Wallace, David (2002a), “Worlds in the Everett interpretation”, *Studies in History and Philosophy of Modern Physics* 33: 637–661.
- (2002b), “Quantum probability and decision theory, revisited”, <http://philsci-archive.pitt.edu/archive/00000885/>.
- (2003), “Everettian rationality: Defending Deutsch’s approach to probability in the Everett interpretation”, *Studies in History and Philosophy of Modern Physics* 34: 415–439.
- (2005), “Epistemology quantized: Circumstances in which we should come to believe in the Everett interpretation”, <http://philsci-archive.pitt.edu/archive/00002368/>. Forthcoming in *British Journal for the Philosophy of Science*.