Probability, Self-Location, and Quantum Branching

Peter J. Lewis†‡

The main problem with the many-worlds theory is that it is not clear how the notion of probability should be understood in a theory in which every possible outcome of a measurement actually occurs. In this paper, I argue for the following theses concerning the many-worlds theory: (1) If probability can be applied at all to measurement outcomes, it must function as a measure of an agent's self-location uncertainty. (2) Such probabilities typically violate reflection. (3) Many-worlds branching does not have sufficient structure to admit self-location probabilities. (4) Decision-theoretic arguments do not solve this problem.

1. Introduction. The many-worlds theory is widely regarded as the best solution to the foundational problems of quantum mechanics, and with some reason. It requires no additions or modifications to the standard unitary dynamics, and as a consequence, it is not subject to the conflicts with relativity that plague other solutions to the measurement problem. But it faces a conceptual problem of its own; it is not clear how the notion of probability should be understood in a theory in which every possible outcome of a measurement actually occurs. In this paper, I argue that the only plausible role for probability in the many-worlds theory is as a measure of an agent's self-location uncertainty. I show that self-location uncertainty can generate probabilities for simple branching scenarios and that the probabilities involved typically violate reflection. However, I also show that self-location uncertainty cannot generate the standard quantum probabilities in the context of the many-worlds theory. Finally, I argue

†To contact the author, please write to: Department of Philosophy, University of Miami, P.O. Box 248054, Coral Gables, FL 33124-4670; e-mail: plewis@miami.edu.

‡Research for this paper was supported by a National Science Foundation Scholar's Award (SES-0724760). I also acknowledge the hospitality and support of the Centre for Time at the University of Sydney and the Department of Philosophy at Durham University. I would like to thank audiences in Sydney, Canberra, and St. Andrews and at the British Society for the Philosophy of Science 2008 and PSA 2008 meetings for helpful feedback.

Philosophy of Science, 76 (December 2009) pp. 1009–1019. 0031-8248/2009/7605-0018$10.00
Copyright 2009 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/605805 Published online by Cambridge University Press

that the difficulties facing the many-worlds theory are not adequately addressed by the Deutsch-Wallace decision-theoretic approach, as the central proof presupposes reflection. I conclude that the probability problem facing the many-worlds theory remains unsolved.

2. Probability and Uncertainty. By "the many-worlds theory" I mean the recent version in which branching is an emergent phenomenon induced by decoherence (Wallace 2003). That is, at the fundamental physical level there is a single universe whose state evolves according to the standard unitary dynamics. This state can be written in various bases and will generally take the form of a complicated superposition of terms. For certain choices of basis, decoherence ensures that the interference between the terms is negligible. For such a choice, via a suitable coarse-graining of micro-properties into macro-properties, one can find the structure of a classical world in each term—that is, cats, pointers, observers, and such. It is these terms that constitute the *worlds* of the many-worlds theory.

A quantum measurement occurs when a micro-property of a system is correlated with a macro-property such as the position of a pointer. When a quantum measurement occurs in a world, the underlying unitary dynamics causes the relevant term to split into several terms, one for each possible outcome of the measurement. Hence the world *branches*, and an observer, as an inhabitant of a world, also branches; the premeasurement observer has several postmeasurement successors, one for each possible measurement result. So while at the fundamental level there is always just one physical world, at the macroscopic level the history of an object (such as a person) has a structure like the branches of a tree.

It is important that any account of quantum mechanics reproduce the probabilistic predictions of the standard theory, since it is the agreement between these probabilities and the observed frequencies of measurement outcomes that constitutes our evidence for the theory. According to the standard theory, the probability of a measurement outcome is given by the squared amplitude of the corresponding term in the quantum state. But the many-worlds theory takes these terms to be worlds, and it is far from clear why the amplitude of a world should be taken to be its probability. The underlying dynamics of the many-worlds theory is entirely deterministic; amplitude is not related to the objective chance of something, as would be the case for a collapse theory. So any probabilities in the theory will have to be epistemic, reflecting the observer's uncertainty about something or other. But what can the observer be uncertain about? Since an observer can look forward to having a successor who sees each possible measurement result, it is hard to see how the probability of her seeing a particular result can be anything other than 1.

There are several proposals in the literature arguing that the premea-

https://doi.org/10.1086/605805 Published online by Cambridge University Press

surement observer can nevertheless be uncertain about which result she will see, but none of these, I think, are ultimately successful (Lewis 2007). However, as Vaidman (1998) has argued, there is an uncontroversial source of *post*measurement uncertainty. Suppose the observer closes her eyes during the measurement, so that she cannot see the measurement result. Then after the branching event, each postmeasurement observer can be unsure which result she will see when she opens her eyes. This is a kind of self-location uncertainty; each successor becomes uncertain about her own location in the branching structure. In fact, it seems that postmeasurement self-location uncertainty is the only potential source of uncertainty in the many-worlds theory, since a sufficiently well-informed observer can know all the facts about the branching structure of the universe, in which case it is only her own location within this structure that is opaque to her.

Perhaps I have overlooked a source of premeasurement uncertainty. But note that any such uncertainty must be compatible with Vaidman uncertainty; it can only serve to make this postmeasurement uncertainty accessible to the premeasurement observer. Hence the results below follow even if there is a hidden source of premeasurement uncertainty.

3. Branching. So probability enters the many-worlds theory, if it enters at all, as a measure over an observer's postmeasurement self-location uncertainty. By considering simple branching scenarios, it is easy to see how such uncertainty can arise. Putting aside quantum mechanics for the time being, consider a person who splits into two copies, for example via a malfunctioning *Star Trek* transporter. Suppose that one copy is produced wearing a T-shirt bearing the number 1 and the other wearing a T-shirt bearing the number 2. Then each postsplit person can be uncertain what the number on her T-shirt is. Furthermore, by symmetry considerations, it is plausible, at least, that she should partition her uncertainty equally between the two possibilities—that she should ascribe a subjective probability of 1/2 to each. More generally, if she splits into n copies, wearing numbers between 1 and n, it seems that she should ascribe a subjective probability of $1/n$ to each possibility.

On further thought, though, it is not so clear that self-location uncertainty in simple branching scenarios admits a coherent probabilistic interpretation. Consider the following case of iterated branching (which I learned from David Wallace). The person initially branches into two copies, labeled 1 and 2, and a little while later the copy labeled 2 branches into two further copies, labeled 3 and 4. Symmetry considerations suggest the following reasoning: after the first branching event, the person should ascribe a probability of 1/2 to each outcome. After the second branching event, her credence in outcome 2 should again be divided equally between

https://doi.org/10.1086/605805 Published online by Cambridge University Press

the further outcomes 3 and 4. Hence at the end of the experiment her credence in 1 should be 1/2, and her credence in 3 and 4 should be 1/4 each. Compare this case with one in which the person branches directly into three copies labeled 1, 3, and 4; here, by the argument of the prior paragraph, each outcome should be ascribed a credence of 1/3.

But note that these two cases are continuous with one another; if the time between the successive branchings is made smaller and smaller, the first case approaches the second. The probability assignments, on the other hand, are discontinuous between the two cases; however tiny the interval between the two branching events, outcomes 3 and 4 should be ascribed a probability of 1/4, not 1/3. This is problematic; if the time interval is small enough, there may be no fact of the matter about whether this is a case of the first or the second kind, given a plausible vagueness about the individuation of macroscopic events. More to the point, the agent concerned surely doesn't care whether this is a case of the first or second kind; she has branched into three copies, and it is hard to see how it can make a radical difference to her state of uncertainty whether the branching occurred all at once or *almost* all at once. Hence it looks like the self-location uncertainty involved in simple branching scenarios cannot be quantified using a probability measure.

4. Reflection. However, there is an assumption of intertemporal coherence at work in the above reasoning that may well be unwarranted in cases of self-location uncertainty. This assumption is a corollary of reflection—the principle that your current credences should equal your future credences when those future credences are known to you. The corollary is that if all that happens to you is a branching event, where a possible self-location splits into several possible self-locations, then the subjective probabilities of the new self-locations should add to the subjective probability of the old self-location. This assumption is clearly at work in the reasoning that leads to the asymmetric probabilities in the first case above.

But reflection and its corollaries have been called into question for reasoning about self-location, in particular in the literature on the Sleeping Beauty problem (Elga 2000). The case is familiar: Sleeping Beauty believes that while she is sleeping on Sunday night a coin will be tossed, and if it comes up heads she will be woken briefly on Monday, but if it comes up tails she will be woken briefly on Monday and Tuesday, and on Monday night her memory of Monday will be erased. What should her credence in tails be on waking? The consensus view is that it should be 2/3. But this is a prima facie violation of the corollary of reflection cited above. On Sunday, Sleeping Beauty thinks that the probability of tails is 1/2. After Sunday this possible location branches into two further locations,

https://doi.org/10.1086/605805 Published online by Cambridge University Press

namely "Monday and tails" and "Tuesday and tails," to which she ascribes a probability of 1/3 each. Hence the probabilities ascribed to the further locations do not add up to the probability of the original location from which they have branched. Similarly, on Sunday Sleeping Beauty thinks that the probability of heads is 1/2, but after Sunday she thinks the probability of heads—that is, the credence she attaches to the "Monday and heads" location—is 1/3. That is, even though she knows that her future credence in heads will be 1/3, this gives her no reason to shift her current credence from 1/2, contradicting reflection.

Does this lesson carry over to the branching scenarios considered above? There are undoubted similarities. In the tails case, Sleeping Beauty's memory erasure means that her epistemic access is limited; there are links of memory between Monday and Sunday and between Tuesday and Sunday, but none between Tuesday and Monday. In this sense, it is akin to branching into two persons. In any case, in virtue of the memory erasure Sleeping Beauty becomes unsure of her location in time just as the agent in the above branching scenario becomes unsure of her location among the branches.

However, there are also dissimilarities between the cases. In particular, the initial coin toss in the Sleeping Beauty case does not induce branching or self-location uncertainty; her uncertainty concerns a categorical (non-self-locating) fact about the world, not her location in it. But it is not hard to produce a case of genuine iterated branching by modifying the Sleeping Beauty case. For example, suppose there is no coin toss, but instead she is told that she will be woken briefly on every day between Monday and Friday. On Tuesday night she is administered a drug that causes her to forget the Tuesday waking, and on Wednesday night she is administered a drug that causes her to forget everything after Sunday. This protocol divides her wakings into two sets of subjectively indistinguishable days. On Monday and Thursday, she remembers Sunday as if it were yesterday; call these her *first* wakings. On Tuesday, Wednesday, and Friday she remembers Sunday as if it were the day before yesterday; call these her *second* wakings.

When she wakes up remembering Sunday, she is unsure whether today is Monday or Thursday and ascribes them each a credence of 1/2. When she wakes up remembering Sunday plus one further day, she is unsure whether today is Tuesday, Wednesday, or Friday and ascribes them each a credence of 1/3. The Tuesday and Wednesday wakings have "branched" from the Monday waking, and yet their credences do not add to her credence in Monday on her first waking. The Friday waking follows from the Thursday waking with no further branching, and yet her credence in Friday on her second waking is not equal to her credence in Thursday on her first waking.

This modified Sleeping Beauty story is strongly analogous to the iterated

https://doi.org/10.1086/605805 Published online by Cambridge University Press

branching case above. It seems, then, that we should not appeal to reflection or its corollaries in the iterated branching case but, instead, should simply appeal to symmetry. That is, after the first branching event, outcomes 1 and 2 should be ascribed a probability of 1/2 each, and after the second branching event, outcomes 1, 3, and 4 should be ascribed a probability of 1/3 each. The credences ascribed to outcomes 3 and 4 do not add to the credence of outcome 2, from which they have branched. Similarly, the credence ascribed to outcome 1 changes from 1/2 to 1/3 even thought there is no further branching. But these are just the kinds of violation of reflection that consideration of the Sleeping Beauty case suggests are inevitable.

This solves the discontinuity problem described above; as the time between the two branching events gets smaller and smaller, the interval in which outcome 1 gets a probability of 1/2 gets smaller and smaller until it disappears entirely. The final probabilities ascribed to outcomes 1, 3, and 4 are 1/3 regardless of whether there are two branching events or a single three-way branching event. Hence the agent's credences do not depend on vaguely defined features of the world she cannot know about and does not care about.

The violation of reflection on which this solution depends seems to be endemic to self-location probability. Indeed, it is not hard to see why this should be so. Reflection requires that one's epistemic situation does not spontaneously deteriorate, and for ordinary categorical facts about the world this is arguably the case (Evnine 2007). But for facts about one's self-location this is clearly false (Arntzenius 2003). To take a trivial example, when one loses track of time, one's uncertainty about one's location in time spontaneously increases; time passes without one's stir, and moments of time are intrinsically all alike. Normally, one has some relational information about one's temporal location via one's memory, but in the absence of external input (e.g., from a clock), memory is a rough and fallible guide, so uncertainty about one's temporal location gradually increases. In Sleeping Beauty's case even this source of information is taken away from her, so her increase in self-location uncertainty is sudden and dramatic. Similarly in the case of branching: the branches are, by hypothesis, subjectively identical (at least before the agent looks at the measurement outcome), so during a branching event the agent's uncertainty about her location among the branches suddenly and spontaneously increases. Reflection is arguably inapplicable in such contexts.

5. Weight. In the simple branching scenarios considered above, a coherent ascription of probabilities to branches by way of self-location uncertainty is possible, albeit one that violates reflection. But what we are really interested in here is many-worlds quantum mechanics. If the many-worlds

https://doi.org/10.1086/605805 Published online by Cambridge University Press

theory is to recover the probability ascriptions of standard quantum mechanics, then the probability of a measurement outcome must be given by the *weight* of its branch—that is, by the squared amplitude of the corresponding term in the quantum state. But the self-location uncertainty considered above ascribes each branch the same probability; there is no role for the weights. Is there any way that a branch-weight probability measure can be given an analysis in terms of self-location uncertainty?

Temporal analogs are suggestive here. Consider a Sleeping Beauty variant in which she believes the following: there is no coin toss, and she will be woken on Monday for 1 hour, and on Tuesday for 2 hours, with memory erasure in between; furthermore, during both wakings, she will be administered a drug that impairs her sense of temporal passage, so that she can't tell whether she has been awake for 5 minutes or 90 minutes. During her waking period, what credence should she assign to its being Monday? A reasonable case can be made that the answer is 1/3. When she wakes, she knows that she is situated somewhere within her 3 hours of waking. Furthermore, by symmetry, her credence distribution over the 3-hour period should presumably be flat, assigning equal probabilities to equal temporal durations. So since she is awake for twice as long on Tuesday as on Monday, she should assign a credence of 1/3 to the latter.

Clearly the probability Sleeping Beauty assigns to Monday can be made to be anywhere between 0 and 1 by a suitable choice of waking intervals on Monday and Tuesday. Hence we have a temporal analog of weighted branching, where the proportion of waking time occurring on Monday corresponds to the weight of the Monday outcome. The weighted branching is realized by creating a continuum of indistinguishable temporal locations, partitioned into two sets, where temporal duration provides the appropriate measure of these sets.

Unfortunately, no analogous analysis is available for the branching involved in the many-worlds theory. What it would take is for each branch to be associated with a continuum of subjective locations, where branch weight provides the appropriate measure of these locations. But structure of this kind is not present in the many-worlds theory. Many-worlds branching, recall, is not part of the fundamental physical description of the universe but is an emergent phenomenon. A macroscopic state of affairs, such as a measurement outcome, is taken by the advocates of the many-worlds theory to be a particular pattern in the underlying physical stuff of the world, embodied in a branch of the universal quantum state (Wallace 2003). The same goes, presumably, for subjective states like seeing a particular measurement result; a subjective state is a physical state, which in turn is a pattern embodied in a branch of the universal quantum state. This branch has a certain weight, and in this sense the pattern can be said to have a certain weight. But it makes no more sense to say that a

https://doi.org/10.1086/605805 Published online by Cambridge University Press

higher-weight pattern contains more subjective locations than to say that a pattern on a larger television screen contains more pictures. Distinct subjective locations require distinct patterns in distinct branches.

Self-location uncertainty, then, can at best ascribe probabilities to many-worlds branches according to a "branch-counting" measure that ignores branch weight and simply divides the observer's credence evenly over her possible branch locations. This is the wrong measure; the probabilities so obtained are not those of standard quantum mechanics. Furthermore, a consequence of the emergent nature of the many-worlds branches is that there is no objective number of branches for a given quantum state; there are reasonable choices that differ over the number of branches and agree only over the total branch weight assigned to a given outcome. So the branch-counting measure is not just wrong; it is ill defined (Wallace 2007).

One might be tempted to recover the standard quantum probabilities by simply postulating the required extra structure in the world. Perhaps the weight of a branch is to simultaneous subjective locations as the duration of a period of time is to successive subjective locations; a higher-weight branch has more subjective locations at a time just as a longer duration has more subjective locations over time. This involves denying the above identification of a subjective state with a pattern, but so what? After all, subjective states are mysterious things.

But surely they are not so mysterious; we may not be able to say clearly *why* a particular physical pattern constitutes a subjective state, but that doesn't mean that we are free to deny *that* it does. A perfectly sensible physical story can be told, at least in sketch form, about why a period of time contains a continuum of subjective states and why the duration of that period provides the appropriate measure of those states. The story appeals to the evolution of the underlying physical stuff. But no such story would be available for the proposed association of a continuum of subjective states with a quantum branch; as argued above, there is no underlying physical structure to appeal to. And even if such a continuum of subjective states were to be postulated, no story could be told about why the weight of the branch constitutes the appropriate measure of these states; this would have to be an additional brute postulate.

On the other hand, the history of attempts to make sense of quantum mechanics suggests that no brute postulate is too outlandish. One could bite the bullet here and find a place for self-location probabilities in many-worlds quantum mechanics via the route just suggested. But note that this would constitute a big step toward dualism; subjective states can no longer be accounted for solely in terms of the evolution of the underlying physical state. In effect, one would be moving away from the many-worlds theory and toward the many-minds theory (Albert and Loewer 1988), which postulates a continuum of minds associated with each observer in

https://doi.org/10.1086/605805 Published online by Cambridge University Press

addition to the quantum state. Hence one would be jettisoning the main advantage of the many-worlds theory—that it involves no additions or modifications to the physics. One certainly could not claim to be finding a place for probability *within* the many-worlds theory via this route.

Let me summarize the argument so far. I have argued that the only potential source of uncertainty in the many-worlds theory is self-location uncertainty. I have argued that self-location uncertainty admits of a coherent probability measure in simple branching scenarios, albeit one that violates reflection. Finally, I have argued that the identification of probability with branch weight required by the many-worlds theory cannot be given an analysis in terms of self-location uncertainty. I conclude that the probabilistic predictions of standard quantum mechanics cannot be recovered in the context of the many-worlds theory.

6. Rationality. This conclusion might quite rightly be regarded as hasty, since I have so far ignored the most influential proposal for recovering standard probabilities within the many-worlds theory, namely, the decision-theoretic program of Deutsch (1999) and Wallace (2007). But I think the above considerations undermine this program. Deutsch and Wallace prove that if one assumes some standard axioms of rationality, the only assignment of credences to measurement outcomes satisfying those axioms is the one that identifies credences with branch weights. But note that the Deutsch-Wallace proof rests on some form of reflection; in Deutsch's original (1999) proof, the substitutability axiom is equivalent to the corollary of reflection considered above, and in Wallace's (2007) version, the equivalence principle ensures that reflection is satisfied. Given the foregoing, this reliance on reflection is prima facie suspicious.

But perhaps this suspicion can be assuaged; Wallace (2007) attempts to justify his equivalence principle (and hence reflection), arguing that it follows from the deeper principle that an agent shouldn't care about branching per se. That is, it should make no difference to her credences whether an outcome occurs in a single branch or in several subjectively identical branches. Wallace argues for this deeper principle on the basis that "it is a (trivially) provable result of decision theory that introducing 'irrelevant' uncertainty of this kind is indeed irrelevant" (2007, 327). But note that this result is only trivially provable in contexts in which reflection holds; in cases like Sleeping Beauty, introducing "irrelevant" self-location uncertainty *does* affect the agent's credence that the outcome of the coin toss is heads at the location in question. Hence the standard decision-theoretic result cannot be used to justify reflection or its variants in the current context.

Wallace's provides a second argument that an agent shouldn't care about branching per se based on the point mentioned above that the

https://doi.org/10.1086/605805 Published online by Cambridge University Press

number of branches is ill defined in the many-worlds context. One clearly shouldn't care about the value of an ill-defined quantity. But this objection cuts both ways. The world according to the many-worlds theory doesn't contain enough physical structure to allow one to count the branches associated with a particular outcome; and, by the same token, it doesn't contain enough physical structure to allow branch weight to measure self-location uncertainty. That is, even if Wallace's arguments here are successful, their conclusion—that credence depends (only) on branch weight—is something that the physical structure of the theory cannot deliver.

A further possibility worth mentioning is that there is no need to ground probability in any source of uncertainty. This possibility is developed by Greaves (2004) and endorsed as a fallback position by Wallace (2007). The idea is that, even in the absence of uncertainty about what will happen to one's successors in future branches, one can still adopt a differential attitude toward them, which one might describe as how much one cares about each successor. The amount one cares can be quantified using a measure that sums to 1 and has all the other mathematical properties of a probability measure; on this proposal, probability measures degree of care, not degree of belief. Then one can appeal to the Deutsch-Wallace proof to establish that the probabilities so grounded must be identified with the branch weights, on pain of irrationality.

Greaves' proposal bypasses the need for uncertainty in the many-worlds theory, but it does not eliminate the worry about the applicability of reflection. This is because any subjective probability measure can be transformed into a caring measure, even those that violate reflection. For example, one can regard Sleeping Beauty's attitudes toward her various successors as expressing how much she cares about them. On Sunday, she cares equally about her heads successor and her tails successors, but when she wakes on Monday she cares more about the tails successors than the heads successor, and she cannot be accused of irrationality for so doing. Indeed, when what one cares about is, in part, one's own location, the arguments above suggest that one's caring measure should not obey reflection. In that case, Greaves cannot appeal to the Deutsch-Wallace argument to establish that an agent's degrees of care should equal the many-worlds branch weights, since that argument assumes reflection.

Finally, I note that the analogy between many-worlds branching and the modified Sleeping Beauty cases developed here depends on tricky issues concerning personal identity. Only in the many-worlds case does the observer genuinely branch into distinct persons; in the Sleeping Beauty case, there is a single person whose epistemic access to her past has a branching structure. One might argue that this disanalogy affects the argument that an agent's credences may rationally violate reflection in the many-worlds case. This worry can be mitigated, I think, by noting

https://doi.org/10.1086/605805 Published online by Cambridge University Press

that rationality considerations do not constrain *whom* one cares about, and hence issues of personal identity cannot be relevant to decision-theoretic arguments (Price 2008). In particular, it cannot be a matter of rationality that I care only about my successors relative to this branch and not about my counterparts in other branches. This promises to eliminate the force of the disanalogy between the two cases; however, these issues are complex and must be put off for another occasion.

7. Conclusion. In conclusion, then, I have argued that any probabilities in the many-worlds theory must be probabilities concerning an agent's self-location and that such probabilities typically violate reflection. A straight attempt to ground the standard quantum probability assignments in self-location uncertainty fails for the many-worlds theory, because the theory does not have the relevant physical structure to allow the identification of branch weight with self-location uncertainty. The more indirect Deutsch-Wallace decision-theoretic approach fails because the decision-theoretic axioms they rely on include reflection, in a context in which reflection is in doubt. Hence the claim of the many-worlds theory to recover the predictions of standard quantum mechanics remains unsubstantiated.

REFERENCES

Albert, David, and Barry Loewer (1988), "Interpreting the Many-Worlds Interpretation", *Synthese* 77: 195–213.

Arntzenius, Frank (2003), "Some Problems for Conditionalization and Reflection", *Journal of Philosophy* 100: 356–370.

Deutsch, David (1999), "Quantum Theory of Probability and Decisions", *Proceedings of the Royal Society* A 455: 3129–3137.

Elga, Adam (2000), "Self-Locating Belief and the Sleeping Beauty Problem", *Analysis* 60: 143–147.

Evnine, Simon (2007), "Personhood and Future Belief: Two Arguments for Something Like Reflection", *Erkenntnis* 67: 91–110.

Greaves, Hilary (2004), "Understanding Deutsch's Probability in a Deterministic Multiverse", *Studies in History and Philosophy of Modern Physics* 35: 423–456.

Lewis, Peter J. (2007), "Uncertainty and Probability for Branching Selves", *Studies in History and Philosophy of Modern Physics* 38: 1–14.

Price, Huw (2008), "Decisions, Decisions, Decisions: Can Savage Salvage Everettian Probability?", http://philsci-archive.pitt.edu/archive/00003886/.

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.

Wallace, David (2003), "Everett and Structure", *Studies in History and Philosophy of Modern Physics* 34: 87–105.

——— (2007), "Quantum Probability from Subjective Likelihood: Improving on Deutsch's Proof of the Probability Rule", *Studies in History and Philosophy of Modern Physics* 38: 311–332.

https://doi.org/10.1086/605805 Published online by Cambridge University Press