To appear in the *British Journal for the Philosophy of Science*.

Branching and Uncertainty

Simon Saunders[1] and David Wallace [2]

Abstract: Following Lewis, it is widely held that branching worlds differ in important ways from diverging worlds. There is, however, a simple and natural semantics under which ordinary sentences uttered in branching worlds have much the same truth values as they conventionally have in diverging worlds. Under this semantics, whether branching or diverging, speakers cannot say in advance which branch or world is theirs. They are *uncertain* as to the outcome. This same semantics ensures the truth of utterances typically made about quantum mechanical contingencies, including statements of uncertainty, if the Everett interpretation of quantum mechanics is true. The 'incoherence problem' of the Everett interpretation, that it can give no meaning to the notion of uncertainty, is thereby solved.

1 Introduction

The Everett or many-worlds theory of quantum mechanics (EQM) has been much criticized on the ground that it can make no sense of talk of uncertainty, and hence has no place for the notion of chance (the 'incoherence problem': Greaves [2004], Lewis [2007]. Baker [2007]). This need not prove fatal to EQM, as there may be other attitudes to future contingencies, other than epistemic, which motivate behaviour in much the same way, and indeed by the same rules as maximizing expected utility and Bayesian updating (Greaves [2004], Greaves [2007], Greaves and Myrvold [2007]). But that is to add an additional twist to a theory that is already, to put it mildly, conceptually challenging.

We shall say no more of the prospects for EQM on this strategy, for we claim that the incoherence problem can be solved, in the following sense: there is a ready set of semantic rules according to which our actual extant, ordinary talk of ignorance and uncertainty comes out as true, whether in accordance with conventional quantum mechanics (including the measurement postulates) or in EQM. The rules are simple; they are also conservative, even to the point that they can be stated in a way which is neutral between many- and single-world theories. In the single-world case they are banalities. There is no reason, if EQM is true, to reject or amend them.

[1] Philosophy Centre, Oxford University, e-mail simon.saunders@philosophy.ox.ac.uk.
[2] Balliol College, Oxford University. e-mail: david.wallace@balliol.ox.ac.uk.

This is not to say that EQM has no impact on other everyday beliefs, and certainly not that it has no impact on theoretical beliefs, among them metaphysical beliefs. It will be no surprise if the semantics we are proposing is contested on metaphysical grounds—say on the grounds of a metaphysical theory of reference, or a metaphysics of time, or of personal identity, or of modality. But to this our reply is that we are not proposing a metaphysical theory, but a manual of translation, in something like Quine's sense: the standard of correctness is set by fluency of discourse and the principle of charity, not by metaphysical principles; by its serviceability to our best physical theory. That theory is EQM.

Of course, in a certain sense, EQM is *also* a metaphysical theory, for at the crudest level it just *is* a theory of branching universes—it describes a set of branching worlds, of a sort explicitly considered (although rejected) by Lewis in his *On the Plurality of Worlds*. It will be convenient, in fact, to begin with Lewis' metaphysics, for it sets out the central questions well enough. Specifically, we begin with the topic of personal identity, and branching *within* a world.

2 Metaphysics

2.1 Personal fission

Parfit, in his *Reasons and Persons*, posed a famous challenge to the concept of personal identity. He argued that in certain scenarios there can be no good answer to the question 'which person am I'. These scenarios involved *personal fission*: processes whereby persons, using various speculative procedures, are divided into two (at its crudest, by a surgical procedure). His point was that the concept of personal identity breaks down in such problem cases, so it cannot be judged fundamental. As such, it should not play a foundational role in theories of ethics.

Against this, Lewis [1976] argued that personal fission poses no particular problem to the concept of personal identity. There may be threats from other quarters, but he was able to show how the identity relation can be extended in a conservative way to branching without ambiguity or change in its formal properties. The trick was to suppose in the face of branching, say into two, that *there are two persons present all along*—persons who initially overlap or coincide. This is equivalent to the stipulation that by 'person', roughly speaking, we mean a unique cradle-to-grave continuant, specifically a unique spacetime worm. As for the meaning of 'overlapping', there are plenty of homely analogies: the Chester A. Arthur Parkway, he observed, overlaps with Route 137 for a brief stretch, but still there are two roads.

But as the spatial illustration shows, there are also times when we want to say that there is only one thing present—for example, when saying how many roads, burdened as you are with the shopping, you have to cross to get to the other side. Even more so in the temporal case: surely we want to say, prior to branching, that there is only one person present. The matter seems context dependent, where the context is set by local or global concerns: fixing on the

local state of affairs, we should say there is one road (one stretch of road) or one person present (one common stretch of persons), whereas considering the global situation, we should say there are two. 'Local', spatially, means we count the roads as identical at x iff they share a spatial part at position x; temporally, suggested Lewis, we should count continuants C_1 and C_2 as identical at t iff they share the same temporal part at time t. The latter he called 'tensed identity'.

Lewis, we believe, was on to something, but it has to be admitted that his answer to Parfit has won few supporters. It seems to have come unstuck, in particular, when it comes to the question of what overlapping persons should say of events on different branches (Parfit [1976]). But here it seems his argument took a wrong turn, rather than that it was in principle wrong-headed. For Lewis went on to attribute thoughts (and we must suppose utterances) at time t to the common temporal parts of speakers at time t. In the absence of branching, utterances are naturally attributed to persons, which, from a four-dimensional point of view, are either continuants or stages. To attribute them to stages, whilst insisting that persons are continuants, seems perverse.

If persons are continuants, we do better to attribute thoughts and utterances at t to continuants C at t. That is, thoughts or utterances are attributed to ordered pairs $\langle C, t\rangle$ or slices of persons $\langle C, S\rangle$, $S \in C$, not to temporal parts S. This is to apply whether or not there is branching. In the absence of branching we obtain the standard worm-theory view; in the presence of branching we conclude that there are *two or more* thoughts or utterances expressed at t, one for each of the continuants that overlap at that time.

Is it to be objected that thoughts or utterances have an irreducibly local significance? We may grant the point that their tokenings are purely local events—and as such, indeed, are identical—but the content of thoughts and utterances is another thing altogether. On even the most timid forms of externalism, or functionalism for that matter, meanings are context-dependent. The sentences produced pre-branching are likely to play different semantic roles for each person subsequently, and likewise their component terms.

This point suggests wide latitude when it comes to the context-dependence of personal pronouns. Here one might make a case for a variety of different semantic rules, but the one we are interested in is this: *the word 'I' refers to the speaker of any sentence in which it occurs.* In the non-branching case, this rule is a banality.

Lewis ruled out this semantics with little or no argument—indeed, he did not even consider it explicitly—offering instead one of his own. Suppose continuants C_1 and C_2 share the temporal part S at t, and suppose C_1 dies shortly after branching, whilst C_2 survives. Then, said Lewis, C_1 and C_2 'cannot share the straight-forward commonsensical desire that he himself survive', because

> The shared stage S does the thinking for both of the continuants to which it belongs. Any thought it has must be shared. It cannot desire one thing on behalf of C_1 and another on behalf of C_2. If it has an urgent, self-interested desire for survival on the part of C_1, that very thought must also be an urgent, self-interested (and not

> merely benevolent) desire for survival on the part of C_2. It is not possible that one thought should be both. So it is not possible for S to have such a desire on behalf of C_1. So it is not possible for C_1 at t to have the straight-forward commonsensical desire that he himself survive. (Lewis 1976 p. 74).

True enough, if there is only one thought. But why not if there are two, as follows if the referent of 'I myself', thought or uttered at time t (at temporal part S) is the continuant who thinks or utters the phrase, as in the non-branching case? Lewis is driving at the conclusion that since the straight-forward desires are not to be had, we should settle for something not so commonsensical, on e.g. the desire 'let at least one of us survive' (what he calls 'weak survival'). But the more conservative option is surely more attractive: why not allow that C_1 and C_2 each desires to survive? And look at what each of them says: 'I will survive', on our proposed semantics, will be true in C_1's mouth, false in C_2's, but since there is nothing to distinguish them prior to branching, neither one can know the truth value of what he says, for neither one can know which of C_1 and C_2 he is.

The knowledge each of them lacks is evidently *self-locating knowledge*, what Lewis called knowledge *de se*; each may know everything there is to know, when it comes to knowledge *de dicto*—propositional content—and yet, like Perry's amnesiac lost in the library at Stanford, each does not know whom he is. Unlike in Perry's example, however, it is impossible for either to know—short of the moment of death—which person (unique continuant) he is.

We conclude: on this semantics, in the example Lewis considers, neither C_1 nor C_2 at t can know if he will survive. Each at time t should say he is uncertain if he will survive.

2.2 Branching worlds

Notice that on our semantics talk of branching persons is in many cases the same—a number of sentences have the same truth values—as were they diverging, rather than branching. 'Divergence' is Lewis' term for continuants (whether persons, things, or worlds) which are spatiotemporally disjoint, but which for some initial segment are intrinsically, qualitatively, the same. Clearly if C_1 and C_2 are diverging persons, qualitatively exactly the same up to time t but not after, neither at t can know which of them he is. Each is uncertain as to what he will see after t. The semantics we are proposing applies here as with branching. Yet Lewis was at pains, in *On the Plurality of Worlds*, to distinguish branching from divergence.

Of course diverging persons in a non-branching world are not seriously comparable to branching persons in a non-branching world. The two are observationally entirely distinct—diverging persons do not have to undergo surgery or cloning, for example. These are respects in which branching and divergence clearly differ. But shift our topic to *worlds as wholes* and the difference between branching and diverging worlds, on the semantics we propose, is considerably

more subtle. There is no longer any question of any *observational* distinction between the two.

Our claim (on the semantics we propose) is that rather large swathes of ordinary discourse now come out as true, whether worlds branch or merely diverge. Indeed, as Lewis recognized, when it comes to the initial segments of diverging worlds the relation of 'similarity' *ex hypothesi* holds exactly; and as such it is transitive as well as reflexive and symmetric. For such segments, the formal discrepancy between similarity and identity disappears (Lewis [1986 p.209]). That observation effectively neutralizes the intuition that talk of branching must show up linguistically as different from talk of divergence. The salient distinction, between similar and identical initial segments, is invisible at the level of syntax and logical form.

Lewis has two objections to branching worlds. The first is that it makes nonsense of ordinary beliefs about the future: instead one has to say (to use Aristotle's example), supposing that after branching in one world there is a sea battle, but not in the other, that there is a sea battle tomorrow *and* there is no sea battle tomorrow; or, perhaps, that there is nothing at all tomorrow because one doesn't have a unique tomorrow—and so on, for various other semantics. In ever case he considered, he found the result deficient—either strange or apparently self-contradictory. Denizens of such worlds, he concluded—for Lewis allows that the notion of 'branching' is physically meaningful, at some level, so there are such worlds out there in logical space—will be epistemically unfortunate, and, if they speak like us, deeply confused. We, thankfully, are not in their position (although presumably our world is *diverging* from countless others, in logical space).

But this problem, we hope, we are slowly dispelling. Lewis here as in the context of personal fission overlooked the semantics that we have given. On our semantics branching into the future brings with it nothing stranger than the notion of uncertainty. It is a conservative semantics in which ordinary beliefs about the future come out as true.

Lewis's second objection to branching worlds is that they are of no use to modal discourse, or not if taken as providing *de re* transworld identities of the sort sometimes wanted by metaphysicians. Thus, if what is possible for C are the properties that C has in all possible worlds that contain C, then C has all its intrinsic properties in all worlds possible for C—so none of its intrinsic properties (which do not depend on relations with other things) can be accidental. He called this 'the problem of accidental intrinsics'.

Evidently the problem only arises if *de re* identity is to play a central role in accounting for possibility. As such, again it is a problem that can be stated for diverging worlds as well (replacing identity by exact similarity): if what is possible for C are the properties that C has in all possible worlds in which it has an exactly similar counterpart, the same unwelcome result follows. There is, however, a difference. Under divergence, the antecedent of the conditional has a natural weakening to *approximate* similarity, whereupon the consequent no longer follows. Allow that the properties possible for C are the properties of approximately similar counterparts to C, and the difficulty disappears. But

this trick will hardly work with identity, for what could 'approximate identity' possibly mean?

There may of course be other tricks that work with identity (hence with exact similarity). In fact one is rather promising: 'C might have had P' is true if and only if C has a temporal part that is a temporal part of a continuant C' that has P (in terms of exact similarity, '....if C has a temporal part that is exactly similar to a temporal part...'). Thus Al Gore might have won the 2000 US presidential election, if he has a temporal part which is a part of a person who won.

Simpler still is to suppose *de re* identities play no particular role in one's treatment of the modalities, just as Lewis makes no particular use of exact similarities in his. But that undercuts what is supposed to be the principal attraction of the picture of branching worlds, which is precisely that it promises to provide the needed *de re* identities. Combine this with the worry about making nonsense of ordinary beliefs about the future, and it is understandable that branching worlds have never been much in favour, not even among those who, like Lewis, believe all possible worlds to be real. (Of course antirealists about possible worlds never so much as consider branching worlds, or exactly similar worlds—for the real world cannot be exactly similar to an ersatz world.)

One can put the matter in terms of tensed identity. Branching and diverging worlds are made to look the same using tensed identity. Elsewhere Lewis wrote of tensed identity as a somewhat unnatural and *ad hoc* device (Lewis [1986 pp.218-9]) as applied to branching within a world, but in that case, post-branching, the number of branches is directly observable. We have already remarked on this—no such observations are available if it is worlds as wholes that branch, no more so than if they merely diverge. All the more reason, then, to count by tensed identity—not just persons, but things and worlds themselves. There is no doubt that if Al Gore had won the election he would have tried to save the Amazonian rainforest. Which rainforest would that be, exactly?—why, his rainforest, the one in the world in which he was elected. So there is no doubt about how many rainforests Al Gore was trying to save: just the one, his own—the same number as by counting with tensed identity.

That would seem to suggest that things are not denizens of more than one world after all. Didn't we say that branching as opposed to divergence brings with it *de re* identities? There is no formal contradiction; the *de re* identities we made use of were of temporal parts, not things; but we seem to end up with the difficulty, on our semantics, of putting the very doctrine we are interested in into ordinary words (it appears to require *metaphysically technical* terms, like 'temporal parts'; it appears *purely* metaphysical). And there are other puzzles that naturally arise: what of the uncertainty involved in branching (equivalently with divergence)—does it comes in degrees? If so, presumably it brings with it the notion of probability—what are these probabilities, exactly, and what have they to do with any statistical data? And what about the converse of branching and divergence, namely recombination and convergence of worlds?

These are deep waters for metaphysicians. No wonder branching has been looked on with suspicion by all parties to the debate (although as we have seen

many of the same questions arise with divergence). All parties, that is, save one.

3 Physics

What if one takes seriously modern physics, and specifically, the most successful fundamental theory of modern physics, quantum mechanics? For quantum mechanics, under the only interpretation to date that can lay claim to being a *realist* interpretation (the only interpretation under which we have a *serviceable universal theory at all*), appears to be saying that the world is constantly branching—if not branching into all possible worlds, then branching into all *physically* possible worlds. And it appears to be saying that it is ubiquitously a *branching* process, not a recombining or converging of worlds, at least at the macroscopic level.

That theory is Everettian quantum mechanics (EQM). We earlier said the conventional theory includes the measurement postulates: it is the latter that pose problems for realism. They are clearly unsatisfactory as dynamical principles, however, in certain circumstances, that is just what they must clearly be (the infamous 'projection postulate'). Stated as they are in terms of 'measurement' or 'observation' (or 'the observer')—or, as in the more recent literature, in terms of a choice of a 'decoherent history space'—they appear variously as out-and-out instrumentalism, idealism, or as requiring top-down causation. Realists would like to do better—yet without some such device as the measurement postulates, it seems impossible to account for or even so much as describe the events actually observed.

This problem ('the problem of measurement') has of course an enormous literature, but a consensus of sorts has emerged. It is that if EQM hangs in the balance—and may ultimately fail, whether on the basis of the incoherence problem, that we are currently pursuing, or some other—two other approaches are agreed by almost all to be credible, namely adding hidden variables obeying supplementary equations, or changing the dynamical equations for the wave-function. What results, in either case, is a new theory—not an interpretation of the existing formalism—whether a hidden-variable theory, or a dynamical collapse theory, or (perhaps) a combination of the two. There are successful theories of the first and second type, although restricted to the non-relativistic regime, namely the dBB theory (de Broglie-Bohm theory, also called pilot-wave theory and Bohmian mechanics) and the GRW theory (after Ghirardi, Rimini, and Weber).

In all three theories, EQM, dBB, and GRW, the fundamental ontology is the universal state (the 'wave-function of the universe'), a function on a space of enormously high (and in quantum field theory infinite) dimensionality. Extracting interesting structures from this wave-function—structures which in their macroscopic features look like the observable universe—is the business of decoherence theory, in EQM and (in the decoherent histories formalism) in standard quantum mechanics. What results in the latter cases is a quasi-classical de-

cohering history space, representing the evolving wave function as a system of branching worlds. In the dBB theory one has precisely this same structure—the fundamental equations for the state are unchanged—but in addition to this there is a particular sequence of particle configurations deterministically threading through (and thereby picking out) a single branch in this history space. In GRW theory the equations are changed, but precisely so as to rapidly suppress ('collapse') all save one of the branches by a stochastic dynamical process.

In dBB the probabilities concern the values of the hidden-variables (they are like probabilities in classical statistical mechanics, whatever *they* are), whereas in GRW they enter in the new stochastic dynamics controlling the collapses. In both cases one is ignorant—of the precise nature of the one and only branch that is actually realized. A notion of 'probability', then, is available as degree of belief, or likelihood, of one branch arising rather than some other: as to how it will go with the way things are objectively arranged. Curiously, in both theories, these degrees of belief are to be carefully matched to the values of certain perfectly definite quantities in the branching structure—to the way things are objectively arranged—namely to the *branch weights*, the modulus squares of the amplitudes of the components of the wave-function, referred to the quasiclassical history space. No reason is given for this, other than that it works on empirical grounds.

What probabilities mean in EQM (if they mean anything at all) is of course the question we are examining—the incoherence problem. If there is ignorance and uncertainty, then it is not like ignorance and uncertainty in GRW and dBB theories, which is about the way the universe is arranged at the fundamental level. In EQM (as in dBB and GRW) we suppose the wave-function is completely known at some time, but (like GRW but not dBB) since there is no other fundamental ontology, there is no further ignorance—and unlike GRW, since the equations are deterministic, there is no ignorance of the future either. So what in EQM is the source of ignorance and uncertainty?

Step back, at this point, from the details of these theories. EQM, in broad brush, says that worlds are subject to branching, and that branches come with *weights*—quantities whose sum is preserved under branching. As a metaphysical picture, it is precisely the one considered in Section 2.2, save that we have these additional quantities, 'weights'. The incoherence objection is in effect that this picture makes no epistemic or metaphysical sense, at least not in terms of probability. One might have expected to find a verdict on this objection in the philosophical literature, but branching as we have seen has been largely ignored by metaphysicians, whilst weighted branching is unheard of. No postulate of this kind could ever have been seriously presented by a metaphysician as an *a priori* hypothesis about modal space; it has too much of the flavour of a physical theory. So, alas, we learn little of direct relevance to this objection from the metaphysicians.

In any case, now that we are coming at this metaphysical picture as a theory of physics, the nature of the argument, on both sides, is transformed. As goes the incoherence problem of EQM, it is now rather clear, from Section 2, of what we are ignorant: we don't know which world—which branch, big-bang to

end-of-time—is ours. It is lack of knowledge *de se*, uncertainty of where we are located, not as a stage S but as a world-stage $\langle W, S\rangle$ or world-time $\langle W, t\rangle$, among the branching worlds. Ignorance on this score makes rather obvious sense in the case of diverging worlds, and now we are in a position to see that it makes just as much sense, on our semantics, in the case of branching worlds. That is, on our semantics, much the same sentences in much the same linguistic contexts have much the same truth values, whether worlds are branching or merely diverging (high-level theoretical sentences, like 'worlds diverge and do not branch' or 'worlds are branches of a decohering history space', do not count). This ignorance is more like GRW ignorance than dBB, where there is something (the values of the hidden variables) which are in principle unobservable. In EQM, as in GRW, even knowing everything there is to know up to t, one's future is still unknown.

The argument is transformed on the side of metaphysics as well. As we concluded in Section 2, some serious questions remained—most notably the worry as to how to characterize branching itself, of how to quantify these self-locating uncertainties, and of what to say about recombination, respectively convergence.

Stating the theory is now not a problem: whatever expressions of it in everyday words may or may not be possible, in accordance with our semantics, we now have available the language of quantum mechanics too. Compare four-dimensionalism, and the difficulty of stating that doctrine in everyday words, under a semantics which preserves the truth of ordinary talk of change. (For the parallels between disputes over branching, re uncertainty, and four-dimensionalism, *re* becoming, see Saunders ([1995], [1998]) and, in terms of radical translation, Wallace [2005].)

As for how to *quantify probabilities*, branch-weights which are 'equivariant' (whose sum is preserved by the dynamics), quantities never considered in the philosophy literature, are just the ticket. More than that, a potential rival probability measure, which actually leads to severe problems with diachronic consistency—to take the worlds produced on branching to be *equiprobable*—is revealed as a will o' the wisp, relying on numbers that aren't even approximately defined by dynamical considerations (they are rather defined by the *number of kinds* of outcome, oblivious to the number of outcomes of each kind). This point has been made a number of times in the literature (see e.g. Saunders [1998], Wallace [2003]), although it is often ignored or forgotten. Thus Lewis ([2004]), in his one foray into quantum mechanics (published posthumously), and Putnam ([2005 pp.630-1]), in a reprise of his earlier famous paper 'A Philosopher Looks at Quantum Mechanics', made much of this supposed alternative to branch weights in quantifying probability. (See Saunders [2005], Wallace [2007] for recent and detailed criticisms on this putative probability measure.)

Finally, worries about how to handle *recombination* (branching worlds) and *convergence* (diverging worlds) now have a rather different character. Recombination, where it can be realized, is quantum erasure, and is (at least to date) common ground to EQM as to the standard formalism. Insofar as it may be predicted in novel circumstances by EQM, it is likely to lead to novel experi-

ments.

Such are the virtues of naturalized over a prioristic metaphysics.

4 Objections

It will not do, in this changed situation, to object that the semantics we have introduced is counter-intuitive: intuition, if we are talking of physical discoveries, doesn't come into it. Nor does it matter if there is some other, alternative semantics to ours in the branching case, according to which the incoherence problem remains. *Of course* there is; our goal was to find a semantics in which it is solved. We are not looking for deep metaphysical truths about identity, the referent of 'I', the nature of persons, and so on; we are looking for serviceability.

We do not, in particular, accept that there is a *correct* semantics, whether or not consistent with ours, as determined by metaphysical principles. For we do not believe there are any metaphysical truths when it comes to the meanings of everyday words like 'person' and 'I', over and above those that are fixed by observable linguistic useage. Our picture of the meanings of words like these is like Quine's: giving truth conditions for their use, in terms of our best physical theory, is similar to compiling a manual of translation, subject to the principle of charity. That and ease of use, clarity, and consistency, are for us the only relevant criteria. Our proposal, on none of these criteria, is in contention. Our rules are too simple to harbour concealed contradictions.

Is it a worry that our proposed semantics may lead to revisions, if EQM is true, of truths customarily assumed in testing and confirming quantum mechanics (or any physical theory) in the first place? That would be a serious difficulty if it were true. But to guard against this is a principal rational for the principle of charity, insofar as it is can be applied under theory change. Call it conservativism, the principle of charity, or the principle of minimal mutilation, the upshot is the same. From the beginning our goal was to show that EQM can be defended against the incoherence objection.

Is it that the semantics, *qua* an account of meaning and reference, violates customary norms and principles in the philosophy of language and philosophical logic? That too would be damaging if true, if somewhat less so than the case just considered, as not directly self-undermining. But is it true? Apparently not from the point of view of philosophical language: the semantics is bivalent; utterances are attributed to persons; and the referents of personal pronouns are context-dependent, depending (at least in part) on the utterer of the sentence in which they appear. Nor does it appear true from the point of view of philosophy of language, for its salient features are these: meanings are determined by global linguistic usage, rather than by metaphysical principles; not all knowledge is propositional; not all content of thought is narrow content. These are all standard principles.

To conclude: if—as Lewis proposes—in cases of personal branching we say there are two persons present even before the branch, it is at least somewhat natural to attribute two sets of thoughts to those persons; in the case of worlds

branching, it becomes entirely natural. As a result, talk of uncertainty in the face of branching comes out as true. Whether or not branching thereby finds new applications in modal metaphysics is an open question, but since our best theories of *physics* seem to describe branching worlds, it finds a natural application to the physical universe.

References

Baker, D. [2007]: 'Measurement outcomes and probability in Everettian quantum mechanics', *Studies in the History and Philosophy of Modern Physics*, **38**, pp. 153-69.

Greaves, H. [2004]: 'Understanding Deutsch's probability in a deterministic multiverse', *Studies in the History and Philosophy of Modern Physics*, **35**, pp. 423-56, <philsci-archive.pitt.edu/archive/00001742/>

Greaves, H. [2007]: 'On the Everettian epistemic problem', forthcoming in *Studies in the History and Philosophy of Modern Physics*, <philsci-archive.pitt.edu/archive/00002953/>

Greaves, H., and Myrvold, W. [2007]: 'Everett and evidence', forthcoming.

Lewis, D. [1976]: 'Survival and identity', in A. Rorty, (ed.), *The Identity of Persons*, University of California Press. Reprinted in D. Lewis, *Philosophical Papers*, Volume 1, Oxford: Oxford University Press, 1983.

Lewis, D. [1986]: *On the Plurality of Worlds*, Oxford: Blackwell.

Lewis, D. [2004]: 'How many lives has Schrödinger's cat?', in *Lewisian Themes: The Philosophy of David K. Lewis*, F. Jackson and G. Priest, eds., Oxford: Oxford University Press.

Lewis, P. J. [2007]: 'Uncertainty and probability for branching selves', *Studies in the History and Philosophy of Modern Physics*, **38**, pp. 1-14, <philsci-archive.pitt.edu/archive/00002636>

Parfit, D. [1976]: 'Lewis, Perry and what matters', in A. Rorty, ed., *The Identity of Persons*, University of California Press.

Putnam, H. [2005]: 'A philosophy looks at quantum mechanics (again)', *British Journal for the Philosophy of Science*, **56**, pp. 615-634.

Saunders, S. [1995]: 'Time, quantum mechanics, and decoherence', *Synthese*, **102**, pp. 235-66, <philsci-archive.pitt.edu/archive/0211138>

Saunders, S. [1998]: 'Time, quantum mechanics, and probability', *Synthese*,**114**, pp. 373-404, <philsci-archive.pitt.edu/archive/00000465/>

Saunders, S. [2005]: 'What is probability?', in *Quo Vadis Quantum Mechanics*, A. Elitzur, S. Dolev, and N. Kolenda, (eds), Springer, <philsci-archive.pitt.edu/archive/00002011/>

Wallace, D. [2003]: 'Everett and structure', *Studies in the History and Philosophy of Modern Physics*, **34**, pp. 87-105, <philsci-archive.pitt.edu/archive/00000681/>

Wallace, D. [2005]: 'Language use in a branching universe', forthcoming <philsci-archive.pitt.edu/archive/00002554/>

Wallace, D. [2007]: 'Quantum probability from subjective likelihood: improving on Deutsch's proof of the probability rule', *Studies in the History and Philosophy of Modern Physics*, **38**, pp. 311-32, <philsci-archive.pitt.edu/archive/00002302/>