Language use in a branching Universe

David Wallace [*]

December 2005

Abstract

I investigate the consequences for semantics, and in particular for the semantics of tense, if time is assumed to have a branching structure not out of metaphysical necessity (to solve some philosophical problem) but just as a contingent physical fact, as is suggested by a currently-popular approach to the interpretation of quantum mechanics.

1 Introduction

Aristotle thought that the future was open, in the sense that talk of the future need be neither true nor false. Prior (1957,1967) modelled this openness by giving time a branching structure and constructing a branching-time semantics to match. Thomason (1970), McCall (1984, 1994), Belnap (1992, 2002), and others concur. In each case their motivation is that certain aspects of our everyday world, and our discourse about it, can be modelled only by branching time and branching-time semantics: examples include causality, the objectivity of possibility and probability, free will, and the intuition to say of things not yet determined that they are neither true nor false.

But it is, I think, at least as widely believed that such a radical metaphysical strategy is unnecessary, and that our everyday interactions with time can be understood perfectly well on the conventional picture of time as non-branching. If so (and I am myself inclined to think that it *is* so), philosophy gives us no reason to take seriously the proposal that time is branching, or to adopt a semantics of tense which explicitly accommodates branching.

However, it is a *non sequitur* to suppose that because we are not *compelled* to accept a non-branching structure for time, it is not philosophically *possible* to accept one. This paper is concerned with what semantics we should adopt for tensed discourse in a universe in which branching is *in fact* ubiquitous: its conclusion is that if we live in a branching universe, we are required to adopt a branching semantics.

Why take such a supposition seriously in the first place? One important reason is that 'new developments in physics' do indeed give us (some) grounds

[*]Balliol College, Oxford OX1 3BJ, U.K. Email: `david.wallace@balliol.ox.ac.uk`

to suspect that time may be branching. In particular, the Everett interpretation of quantum mechanics (often called the 'Many-Worlds theory') makes precisely that claim: it supposes that whenever we make a measurement of a quantum system[1] the Universe splits into many copies, one[2] for each possible outcome of the experiment.

It is not my purpose here to argue for the *truth* of the Everett interpretation (for such arguments, see Wallace (2002, 2003a, 2003b)). However, true or not, it is now taken seriously by remarkably many physicists. As such, philosophers have some obligation to ask, if not *is it true*, then at least *could it be true*? That is, is it a possible physical theory, albeit one which could turn out false? Or is it metaphysically incoherent, something to be rejected out of hand without further ado?

One of the more powerful arguments for metaphysical incoherence[3] is: how can it make sense to say, as we habitually do, that we are *uncertain* of the result of a quantum measurement when in fact the Everett interpretation tells us *with certainty* that the result is a division of the universe into many copies? That is, how can we reconcile the apparently indeterministic nature of quantum measurements with the apparently deterministic, in-principle-perfectly-knowable nature of the Everett interpretation?

To expand on the point: suppose an experimenter, asked what he believes the result of the experiment will be, says 'Either outcome A or outcome B will occur, but not both, and I don't know which.' This statement (a perfectly normal response for an experimenter who accepts quantum mechanics) can be construed as the assertion of

- 'A will occur or B will occur', and
- 'Not both A and B will occur'

together with a refusal to assert either

- 'A will occur' or
- 'B will occur'.

However, if quantum mechanics is supposed to mean that the world branches in two, with A occurring in one branch and B in the other, it's unclear how this can be a reasonable set of assertions to make. Do we take A to occur if it occurs *in some branch* (that is, somewhere in the branching Universe)? Then it cannot be true that not both A and B will occur, nor can it make sense to plead ignorance as to which will occur (they both will). Or should we take A to occur if it occurs period, not just occurs-in-some-branch? If so, neither A nor B

[1]And at many other times when measurement is *not* occurring, too — one of the appealing features of the Everett interpretation is that 'measurement' plays no essential role in its formulation.

[2]Rather more than one, actually; but counting branches is a subtle matter. See Wallace (2003b) for more discussion of this point.

[3]Made with some force by Albert and Loewer (1988), amongst others.

occurs; rather, some alien branching event occurs. Either way, the theory seems to fit very badly with how experimental physicists actually behave.

Responses to this objection have generally taken one of two forms: either bite the bullet and accept that we are just wrong when we say that a quantum-mechanical experiment has an unknown outcome, but try to argue that our behaviour should remain unchanged nonetheless (Greaves 2004); or try to make sense of the apparent paradox by appeal to considerations from the philosophy of personal identity (Saunders 1998). I wish to suggest an alternative possibility: that the 'paradox' is a consequence of mistakes in the philosophy of language. In particular, it is due to the mistaken application of a semantic theory designed for a non-branching universe to the discourse of the inhabitants of a branching universe.

In section 2 I review the semantics of linear and branching time; in section 3 I apply each to a branching universe, and show that linear, but not branching, semantics imply that if we live in a branching universe then nearly all our beliefs are trivial, false or meaningless. Section 4 makes the central claim of this paper: that the linguistic principle of charity entails that branching semantics are correct for a branching universe. Section 5 addresses the frequently made objection that talk of 'uncertainty', which applies to a branching universe if branching-time semantics are appropriate, is incoherent in situations where all the objective facts are known. Section 6–8 develop the technical details of the proposal, and section 9 is the conclusion.

2 Branching-time semantics: a review

Suppose that time is linear: that is, suppose that it can be represented as some total ordering $<$ on some set $\mathcal{T}$ of 'times', to be read as "earlier than".[4] Utterances concerned with time can then be divided in to two classes: *tenseless utterances*, such as "p occurs (tenselessly) before q" or "p (tenselessly) occurs at some time t", the truth values of which are unaffected by when they are uttered; and *tensed utterances*, such as "p has happened" or "q is (now) happening", which are true at one time, false at another. Put in terms of sentences rather than utterances: tenseless sentences (ignoring ambiguity) are true or false *simpliciter*, tensed sentences have time-dependent truth values.

There is a well-known relationship between the two sorts of utterance: namely, a tensed utterance made at a particular time $t \in \mathcal{T}$ is true iff a certain related tenseless utterance is true *simpliciter*. For instance, the tensed utterance "p will happen", uttered at time t, is true iff "p happens (tenselessly) at some time $s > t$" is true; "p is (now) happening", uttered at t, is true iff 'p (tenselessly) happens at t" is true. In terms of sentences: each ordered pair of a tensed sentence and a time determines a tenseless sentence; the tensed sentence is true at that time iff the tenseless sentence is true *simpliciter*.

[4]Recall that a total ordering of $\mathcal{T}$ is an antisymmetric, irreflexive transitive relation on $\mathcal{T}$ such that for any x, $y \in \mathcal{T}$, either $x = y$, $x < y$, or $x > y$.

Those, such as Prior (1957, 1967), who take tense as metaphysically fundamental will, at best, regard this relation between tensed and tenseless sentences as in some sense definitional of the latter; they might wish to go further and reject the framework above altogether (or at least regard it as giving merely the abstract model theory for the logic of tensed utterances). I regard this as untenable, for unoriginal reasons[5], and for the rest of this paper I follow Mellor (1981, 1998) *et al* in regarding the content of tensed utterances as exhausted by an explication of their time-dependent relation to tenseless truthmakers.

What if time is *not* linear, but branching? Formally, this can be modelled by relaxing the requirement that the relation $<$ on the set $\mathcal{T}$ is a *total* ordering. We require instead only that it be a partial ordering[6] which in addition satisfies:

1. *Connectedness*: For any x, y, either $x < y$, or $y < x$, or $x = y$, or for some z, $z < x$ and $z < y$.

2. *No recombination*: If $z < x$ and $y < x$, then either $z < y$, or $y < z$, or $y = z$.

Branching structures for time were proposed by Prior (1957) with the intent of formalising the intuition that (some) future events are not now either true or false. On the branching-time semantics that Prior advocates, a future-tensed sentence (say, 'There will be a sea battle', to borrow Aristotle's example) uttered at t is true iff there is a sea-battle in all branches futurewards of t; false iff there is a sea-battle in no such branch; neither true nor false otherwise.

More formally (and following Thomason (1970)), define a *history* as a maximal totally ordered subset of $\mathcal{T}$. If we restrict our attention just to a particular history, it looks exactly like the 'linear time' we described above. We can easily say what it is for a tensed sentence uttered at t to be true in some history h (where h contains t): the sentence is true at t in h if the associated tenseless sentence is true in h. So in a branching-time picture, the truth values of tensed sentences must be evaluated not just with reference to a time but with additional reference to a history; the truth values of tenseless sentences are not absolute, but require reference to a history. (So 'There (tenselessly) is a sea battle in 480 BC' is true in those histories in which there is a sea battle at that date; false otherwise.)

Actual utterances are made at particular times, but in a branching-time scenario a time does not pick out a unique history. This gives Prior's non-bivalent semantics: a tenseless sentence S uttered at t is true iff it is true in all histories h with $t \in h$ and false iff it is true in no such histories. (Formally, this is a supervaluational semantics; see Thomason (1970) for technical details). We will say, neutrally, that S is *not knowably true or false* at t if S is true in some but not all histories containing t.

Part of the motivation for branching time was to provide an objective framework for future possibilities. In the linear-time framework then possibility is an

[5]See Sider (2001, chapter 1) and Mellor (1998) for recent criticisms of the notion of tense as fundamental.

[6]i. e., a transitive, antisymmetric relation.

epistemic notion: given any two contradictory statements, both may be *possible* but one is true and one is false, and both are possible only because we do not know which is which. If time is branching, however, it is natural to take "it is possible that p" as true at t whenever p is true in some history containing t. (This is actually entailed if we assume that (i) every p is either possible or impossible, and (ii) if something is impossible it is also false). Formal semantics for 'it is possible that' (e. g. Prior (1967)) often take this condition as necessary *and* sufficient, but I shall not do so: it will be convenient to remain neutral as to whether something can be possible but false, as is the case in non-branching time. (I return to this issue in section 8.)

For clarity, I should note two features of this presentation of branching time:

1. Braching-time pictures are often (e. g. , McCall (1994)) combined with essentially tensed views of time, so that as time passes various 'possible' futures succeed or fail to be actualised. This is *not* a feature of the branching time ontology I discuss in this paper: as with my presentation of linear time, the branching structure is taken to have its structure tenselessly.

2. Even if one did not want to accept branching in the actual *structure* of time, it might still be possible to adopt a branching *semantics*. (One might, for instance, say that a future-tense sentence was true at t iff true in all physically possible worlds identical with the actual world up to t, as is advocated by Lewis (1986, pp. 206–9)) By contrast, the picture of branching which I wish to consider is explicitly metaphysical: the real structure of time is branching.

3 Branching time and splitting worlds

Suppose, as does the Everett interpretation, that the universe in which we live is constantly splitting into copies, so that I have not one unique successor but many.

A universe like this has a structure formally identical to that of a branching time. It can be modelled via a set $\mathcal{S}$ of 'instantaneous quasi-classical states', corresponding to the sorts of instantaneous states that 'the classical world' is generally taken to have, together with a relation $<$ on $\mathcal{S}$ satisfying the same axioms as we previously required for branching time and interpreted as causal dependence or some dynamically defined equivalent. Elements of $\mathcal{S}$ describe such things as dead or alive cats, needles on detectors pointing at various determinate values, and so forth; if $x \in \mathcal{S}$ is a state in which a cat is in a box, y, z are similar states in which the cat is respectively alive and dead, and $x < y$, $x < z$, $\sim y < z$ and $\sim z < y$, then we have a model of the "Schrödinger cat" paradox.[7]

[7] In quantum mechanics, $\mathcal{S}$ consists of time-indexed projectors in some decoherence-defined basis; $\langle \widehat{P}_1, t_1\rangle < \langle \widehat{P}_2, t_2\rangle$ iff $t_1 < t_2$ and $\mathsf{Tr}(\widehat{P}_2 \exp(-i/\hbar t\widehat{H})\widehat{P}_1 \exp(-i/\hbar t\widehat{H})) \neq 0$. The splitting process is not part of the fundamental dynamics of the theory, but is emergent from them; see (Wallace 2003a) for details.

Are such universes *actually* examples of branching times? There are two possibilities:

Yes: elements of $\mathcal{S}$ should be identified with times; at each time, the world has a roughly classical description.

No: despite the formal similarity, the branching structure of worlds should not be regarded as a branching-time structure. Instead, we should regard it as having a single, linear time, and having instantaneous states which, far from being classical, consist of vast numbers of causally isolated classical sub-worlds.

In a sense, it is a semantic rather than a metaphysical question as to which of these is correct. The actual structure of the world is not in question (it is given by the set $\mathcal{S}$ and its partial order). The practical consequences of declaring that this system of branching worlds is a branching-time or linear-time system is, rather, that we should use branching-time or linear-time semantics respectively to analyse language use within it.

To elaborate on these remarks: it is of course possible simply to talk *about* the branching structure described by the ordering on $\mathcal{S}$, by uttering semantically unproblematic sentences (in which one might say 'some points with property P are futurewards of some points with property Q' for instance[8] Let us call such utterances 'branch-talk utterances' (or utterances of 'branch-talk sentences').

Furthermore, branch-talk sentences can unproblematically be uttered from within the universes that we are considering, and in being made may incorporate some unproblematic context dependence, For instance, someone aware of the branching might say: 'in the future, I will branch into two copies'. In this context, the referent of 'I' will be the person making the utterance, and the statement is true if there are two distinct branches through the point at which the utterance is made, each containing a copy of the person making the utterance.[9]

What is in question, however, is not the semantics of branch-talk utterances, but the semantics of the utterances actually made inside this world: everyday utterances like 'it will rain tomorrow', for instance. It is here that the two possibilities described above lead to very different answers.

If branching-time semantics are correct, the semantic theory for ordinary utterances is very familiar: each history (that is, each maximal totally ordered subset of $\mathcal{S}$) is classical and unbranching, and provides a perfectly conventional (albeit history-dependent) semantics for tenseless sentences, along the lines spelled out in section 2. A tenseless sentence like "Athens [tenselessly] hosts the Olympics in 2500 AD" uttered at $x \in \mathcal{S}$ is true iff it is true in all histories containing x; 'Possibly Athens hosts the Olympics in 2500 AD' is true at x whenever "Athens hosts the Olympics in 2500 AD' is true in some history

[8]In the context of the Everett interpretation, statements like 'the state of the Universe is $|\psi\rangle$' are also of this form.

[9]Certain issues of personal identity, passed over lightly here, will be dealt with in more detail in section 6.)

containing x (and perhaps at some other points as well, since we are leaving it open the extent to which there is a purely epistemic component to possibility even in a branching universe). In turn, tensed sentences like 'it will rain in Oxford tomorrow', uttered at a particular time, have the same truth value as their associated tenseless sentence ('it [tenselessly] rains in Oxford on the 8th of December 2005', were I to utter it now.)

Things are more complicated if we apply linear-time semantics to branching worlds. Recall that in this case, tenseless sentences are true or false *simpliciter*, not true or false relative to histories (equivalently: there is only one *history*, even though there are many chains of classical sub-worlds). Tenseless sentences, in fact, should be 'God's-eye' descriptions of the entire branching structure: the ideal tenseless sentence would be actually be some 'branch-structure sentence' like "Athens hosts the Olympics in some but not all 2500AD-sub-worlds futurewards of 2005-AD subworld k.".

Ordinary people do not utter sentences like that, of course (or even tensed variants of them), so what of 'ordinary' tenseless sentences like 'Athens hosts the Olympics in 2500AD'? As we have noted in the case of branch-structure sentences, the content of such sentences will have some contextual dependence even on linear-time semantics: 'Athens", uttered at some subworld k, clearly refers to the Athens in that subworld or its causally related future subworlds (much as terms like "Me", "George" or "the Queen" have denotations which depend on the context of utterance). Nonetheless, if the basic assumption of linear-time semantics is to hold then these sorts of utterances should have the same truth conditions as some context-independent branch-structure sentence.

The trouble is that there are no very inspiring candidates. Some might include:

- "Athens hosts the Olympics in at least some 2500AD-subworlds futurewards of 2005-AD subworld k" (where k is the subworld at which "Athens hosts the Olympics in 2500AD" was uttered), which is almost certainly true at most plausible k.
- "Athens hosts the Olympics in all 2500AD-subworlds futurewards of 2005-AD subworld k", which will no doubt be false at most plausible k.
- "The mereological sum of all copies of Athens in 2500AD-subworlds futurewards of 2005 AD subworld k hosts the Olympics", which is false since mereological sums scattered across causally isolated sub-worlds cannot be said to host sporting events.

But none of these seem naturally equivalent in content to the original utterance, though. None seem to have a great deal of contingency; on the first and second proposal "Athens hosts the Olympics" and "Athens doesn't host the Olympics" do not come out as contradictory; on the third proposal we seem to be referring to an object which has little to do with our ordinary understanding of 'Athens'.[10] Perhaps the best strategy to adopt is just to declare "Athens

[10] cf Lewis's arguments (Lewis 1986, pp. 209–219) against the relevance of trans-world fusions to issues of trans-world identity.

hosts the Olympics in 2500 AD" as meaningless according to the linear-time semantics, and restrict our usage to more appropriate tenseless sentences like 'Athens hosts the Olympics in some but not all 2500 AD subworlds futurewards of this subworld".

On the linear semantics, it should be apparent, the hypothesis that our world is branching entails that most of our ordinary beliefs about the future are false, vacuously true, or just meaningless. It also entails that our ordinary description of quantum theory's predictions is highly inaccurate: our ordinary reading of the theory is that the correct attitude to (say) a quantum measurement is uncertainty as to which measurement result obtains, but under the Everett interpretation it seems that all we can say is that *all* 'possible' measurement results actually obtain, although only one obtains in each branch.

As was noted in section 1, these radical conclusions have been used by skeptics about the interpretation (such as Albert and Loewer (1988) and Barrett (1999)) to argue that it cannot be taken as explanatory of the empirical data. They have also been accepted by some advocates of the Everett interpretation (such as Tappenden (2000), and possibly Greaves (2004)), who make the point that such a radical change in our ontology ought to have radical consequences for how we think about the world in general and quantum experiments in particular.

However, it is important to note that such 'radical consequences' *only* follow if the appropriate semantics for a branching world is linear-time semantics. If instead we assume that branching-time semantics are appropriate, virtually all of our existing beliefs remain intact (the main exceptions are technical statements *about* the topology of time or the interpretation of quantum mechanics). Furthermore, the semantics straightforwardly recovers the ordinary description of measurement: if different measurement results occur on different branches, then for each possible result r, "r will obtain" is not knowably true or false, and "r will possibly obtain" is true.

So: the question of whether the Everett interpretation entails massive error on the part of experimentalists and ordinary people comes down to the question of what semantics are appropriate for a branching world.

It is not at all obvious that the burden of proof here lies with the advocate of branching-time semantics. While there has been an extensive debate on the topology of time within mainstream philosophy, it has been almost exclusively concerned with whether branching time is *necessary* to recover some aspect of the everyday world (freedom of will, objectivity of possibility, causation, etc — see Prior (1967) and Belnap (2002) for samples of such arguments). I am not aware of any attempt to argue for the *a priori* necessity of linear-time semantics.

It is, however, possible to give a more constructive argument in favour of branching-time semantics for a branching world. The next section will present it.[11]

[11] Alternative arguments, from rather different perspectives, have been proposed by Saunders (1998) and Vaidman (1998); see Wallace (2005) for a critical evaluation.

4 The charity argument for branching-time semantics

Imagine a certain society, living in a Universe apparently not so different from ours, in which philosophers and physicists have reached a consensus about the correct metaphysics and semantics of time. Specifically: the physicists are confident that their universe is a linear, non-branching reality (whose instantaneous states are rather like 'classical states' of the world in our own universe); the philosophers are confident that the universe is tenseless at a fundamental level and that the correct semantics for tensed talk is linear-time semantics as described in section 2. In their analysis, 'possibility' is an entirely epistemic notion: there is a unique actual future but other futures are 'possible' in a sense that is ultimately "merely" a matter of the ignorance of individuals about the actual future.

This notion of possibility can be expanded and quantified by the use of probabilities: rational natives of the society, the philosophers argue, may be idealised as assigning a probability between 0 and 1 to the truth of each tenseless sentence.[12] (This in turn may be derived from more primitive, qualitative assumptions about which sentences they believe are more likely to be true, using the methods of decision theory; see, e. g. , Jeffrey (1983) and Savage (1972) for two influential proposals as to how this is to be cashed out; see (Kaplan 1996; Joyce 1999) for further details.).

The catch is that the natives' physicists are completely wrong. Their universe is not, as they believed, described by the evolution of matter in an approximately classical way: in reality it is constantly — trillions of times per second — branching into quasi-classical sub-worlds, so that their planet and everything on it is constantly being copied into duplicates almost identical to one another.

Since none of the natives know this, it should be unsurprising that certain aspects of their world-view are mistaken. The question is: which aspects? There are two possibilities.

According to the *Elite View*, despite the errors of the physicists in describing the physical structure of the world, the philosophers' analysis of semantics and rationality remains correct. That is: tenseless sentences have history-independent truth conditions, modulo some referential context-dependence of terms like 'Athens'. As we saw in the previous section, there are no candidate 'branch-talk' sentences to express such truth-conditions without making most of our 'ordinary' future-directed sentences (like 'it will rain tomorrow') trivial or meaningless. Since rationality requires agents to act in accordance with probabilities over tenseless sentences (again ignoring referential indexicality), it follows that 'probability' remains an essentially epistemic notion, measuring an agent's degree of belief in facts which from a God's-eye view are fixed.

To adopt the Elite View is to continue to apply linear-time semantics to

[12] I ignore complications raised by the 'problem of the essential indexical', and although I follow Davidson (1990) in ascribing probabilities to the truth of sentences rather than (e. g.) to propositions, this is purely for convenience and does not affect the argument.

this universe despite the physicists' radical discoveries about its structure. This move reduces to falsehood, triviality or nonsense almost every ordinary opinion held by natives of this universe.

The alternative *Charitable View* accepts that the philosophers are wrong in some important respects in their description of rationality and their linear-time semantics. Tenseless sentences are true or false not *simpliciter* but only relative to a history; actual utterances are therefore to be classified as true, false or not-knowably-true-or-false as dictated by branching-time semantics.

As such, rational agents should define probabilities to the truth of 'tenseless sentences' understood in this way. The price of this move is of course that it abandons the idea that (future-directed) beliefs and desires should be interpreted in terms of ignorance about, and preferences between, possible tenseless states of the entire Universe, and that it abandons any attempt to provide determinate truth-values to future-tense sentences. On the other hand, by doing so, it saves the future-directed beliefs of agents from widespread falsehood. For instance, suppose someone says 'in 2500 AD, either Athens will be hosting the Olympics, or some other city, but not both.' This can be analysed as having the tenseless (but history-dependent) truthmaker

> (Athens is hosting the Olympics in 2500 AD $\vee$ A city other than Athens is hosting the Olympics in 2500 AD) $\wedge\neg$ (Athens is hosting the Olympics in 2500 AD $\wedge$ A city other than Athens is hosting the Olympics in 2500 AD)

which is true at x (according to the Charitable View) iff for every history containing x, in 2500 AD either the Olympics is being hosted by Athens or by some other city but simultaneously by Athens and another city — which let us suppose is in fact the case. So (unlike under the Elite View) this banality does indeed come out true. The sentence 'In 2500 AD Athens will host the Olympics', on the other hand, comes out as not knowably true or false, since in some histories through x it is true and in others false. But this is perfectly in accord with everyday impressions of the openness of the future and with our native's perfectly sensible decision to give (as it happens) probability 0.4 to this outcome, rather than 1 (determinately true) or 0 (determinately false).

So, what is the correct move: to abandon the philosophers' original (linear-time) analysis of semantics, or to abandon the assumption that ordinary natives by and large have reasonable beliefs about the future? In one sense there is no definitive answer to give. For the natives' conceptual scheme, as a whole, fails to fit perfectly with reality, and so needs to be modified; yet if we accept (with Quine, Davidson, Lewis, et al) that the criterion of validity for a conceptual scheme are ultimately somewhat holistic, then there may be no absolute fact of the matter as to which part of a holistically wrong scheme is to be abandoned.

Nonetheless, as a practical matter there are enormously strong reasons why the 'correct' move is to abandon the Elite View. For one thing, in doing so we have to accept as wrong only those parts of the natives' understanding of belief and desire talk which come from a fairly theoretical analysis — furthermore, one motivated (*ex hypothesi*) by a wildly wrong picture of what the universe as a whole is like. By contrast, in clinging to the Elite View we preserve the truth-value of the writings of native philosophers, but at the cost of abandoning the truth of most ordinary conversation about the future.

(As an analogy, consider a variant on Putnam's Twin-Earth example (Putnam 1973). There is a general consensus that, if transported to Twin-Earth and given a glass of XYZ, I'm incorrect to say 'that's water', and correct to say 'water is H_2O'. We interpret my language so that 'water' refers to H_2O — even at the cost of making most of my post-arrival statements about water come out false — because this is the most explanatory interpretation of my linguistic community's overall language use.

But now suppose that I never left Earth at all, but that in fact the stuff in the rivers and seas has *always* been XYZ. Only the tireless efforts of an International Conspiracy of Chemists has kept the public — even quite educated members of the public, such as those with chemistry degrees — duped into thinking that this stuff is H_2O. When investigative journalists expose the Conspiracy, how should we react? Surely not by clinging to our technical, theoretical belief that water is H_2O — thus preserving a few technical truths at the cost of making almost all the population's beliefs about water turn out false — but rather by acknowledging that water was XYZ all along and accepting that the technical connection between water and H_2O is simply false.)

Furthermore (though it is a related point), if we ask what justifies our semantics of the natives' language, it is hard to see how we could motivate the Elite View given the actual (i.e. branching) nature of the natives' universe. For it has been powerfully argued (by Quine (1960) and Lewis (1975), for instance) that some sort of principle of charity must apply to translation and interpretation of a genuinely foreign language. That is, if we have a proposed interpretation of such a language according to which nearly all uses of a large part of that language turn out not to be true (or possibly not even to be meaningful) then that seems simply to count against the accuracy of that proposed interpretation. Judged by this standard, the Elite View should be abandoned as being overwhelmingly less charitable than the Charitable View.

For these reasons, I claim, the correct semantics for the natives' language, and the correct analysis of their beliefs and desires, is that given by the Charitable View: that is, by applying branching-time semantics. The native philosophers, whose proposed semantics is motivated by an entirely reasonable but entirely false view of their world, are simply wrong in their understanding of their own language's semantics. And of course, since *their world* could well be our own world except for some trivial simplifications, if our own world has the branching structure that quantum mechanics suggest then the correct semantics of our own languages are really branching.

5 Probability and uncertainty

An objection remains to be responded to: I have argued that charity considerations require us to adopt branching-time semantics and thus to take an attitude of uncertainty towards the results of experiments which have different results in different branches. But how can it even *make sense* to 'take an attitude of uncertainty' towards the future when we know all the relevant facts?

This objection is frequently brought up in informal discussions of the Everett interpretation; it has been made explicitly by Greaves (2004), who writes:

> I can feel uncertain over P only if I think that there is a fact of the matter regarding P of which I am ignorant.

This objection can be responded to on its own terms, by noting that it is ambiguous between a truism of ordinary language and a technical claim. As the former, it is indeed plausible to say that if a supposed concept of 'uncertainty' says that we can be uncertain about something without there being something to be uncertain about, then it does not deserve the name 'uncertainty'. However, there *is* something about which I am uncertain if I am uncertain about whether Athens will host the Olympics: trivially, that 'something' is *whether Athens will host the Olympics*. Or if preferred, it is about whether "Athens will host the Olympics" will be true.[13]

The less trivial reading of Greaves' comment is that it asserts that we cannot be uncertain about something unless there is a third-party fact — a fact expressible via branch-talk sentences — about which to be uncertain. But now it is far from clear what motivates Greaves to accept this technical claim. This doubt is strengthened when we note that most philosophers do not accept the claim even in the absence of branching: the 'problem of the essential indexical' (Perry (1979), Lewis (1979)) seems amply to demonstrate that there are other sorts of ignorance which are not ignorance of third-party-accessible facts but are unavoidably indexical.

It would be to miss the point to respond to this objection by retreating to the claim that we cannot be uncertain about something unless there is *either* a third-party-accessible fact about which to be uncertain, or an indexical aspect to our ignorance. The point is that we didn't realise that some ignorance was indexical until it occurred to authors such as Perry to consider some such situation, to describe it in everyday language, and to note that we intuitively classify such situations as ignorance even though no relevant third-party-accessible facts are unknown. Essentially the same has occurred in the case of Everettian branching: describe it in everyday language (using the charity argument to tell us that 'everyday language' has a branching-time semantics) and it is intuitively clear that it is another situation of ignorance.

(I should note in passing that on the 'Lewisian semantics' to be described in the next section, indexical ignorance is actually all that we need: the 'essential indexical' is *my history*. This strategy will not work for the other semantic proposals made in that section, though.)

There is also a more positive defence of the use of probability and uncertainty talk in the face of branching. After all, even in the absence of branching no

[13]'Will be true', not 'is true', of course; to find out whether it *will be true*, wait and see. See Prior (1967, pp. 122-3) and Ryle (1964, pp. 15–34) for more on the mild linguistic tangles involved in this wait-and-see notion of truth-conditions.

reductive analysis is given of what it means to assign probability 0.4 to the truth of some such proposition. Rather, the meaning of the statement is given implicitly, as part of a decision-theoretic analysis which treats 'more likely than' and 'preferable to' as theoretical terms whose engagement with empirical fact is via their role in the best explanation of my actual behavioural dispositions. The same applies to the case of branching-time semantics — it's just that the set of propositions has a slightly different conceptual sense. (For more on this issue see section 5 of Wallace (2005)).

6 'Not-knowably-true-or-false': semantic details

So far, I have presented the formal structure of branching semantics and provided an argument from the principle of charity that they are applicable in a branching universe. However, I have so far been rather unspecific as to how this semantics is to be understood — and in particular as to what meaning is to be ascribed to the 'not-knowably-true-or-false' label which I applied to sentences which are true in some but not all histories. In this section I will attempt to remedy this, by providing three different strategies by which 'not-knowably-true-or-false' may be understood.

The most straightforward such strategy might be called *minimalist semantics*, and it takes 'not-knowably-true-or-false' at face value, as a third truth-value: 'indefinite' (with the semantics of that truth value understood according to supervaluationist logic: that is, an utterance of a tenseless sentence s is true (false) at x iff s is true (false) in all branches containing x). This has the advantage of simplicity; it also accords well with the intuitions of an 'open future' which led Prior to tense logic in the first place. On the other hand, it requires us to accept the failure of bivalence, with the consequent logical difficulties (in particular, a rejection of Tarski's truth biconditionals, which lead to contradiction in the absence of bivalence[14]).

An alternative strategy, which we might call *ambiguity semantics*, was suggested to me by Simon Saunders (in conversation). Suppose that just as a tensed sentence is associated with a certain tenseless sentence, a tenseless sentence is associated with a certain 'branch-talk' sentence (e.g. 'in branch x, Y happens'). However, whereas in the tensed case it is usually unambiguous which tenseless sentence is associated with an utterance, in the branching case it is often highly ambiguous. This ambiguity often means that even a tenseless sentence cannot be given a truth value; however, if it comes out true under all the disambiguations then it is true, period. (Saunders has used the example of a proper name: if I use 'David' as a name for the 4-dimensional cradle-to-grave David, then

[14] The Tarski biconditional for the sentence 'it will rain tomorrow' is " 'It will rain tomorrow' is true iff it will rain tomorrow", which has the form " 'It will rain tomorrow' is true iff P"; similarly, the biconditional for 'It will not rain tomorrow' is " 'It will not rain tomorrow' is true iff it will not rain tomorrow", which has the form " 'It will rain tomorrow' is false iff $\sim P$". If it is neither true that it will rain tomorrow nor true that it won't, we can derive $(P\& \sim P)$. See, e. g. , Williamson (1994, pp. 187–197) for a more detailed presentation of this argument.

it is unavoidably ambiguous which of the vastly many Davids I am actually referring to. However, the semantics of a branching universe cannot *reduce* to the ambiguity of proper names, because many future-directed statements we might wish to make (e. g., 'there will be a war next year') contain no names at all.) Ambiguity semantics 'demystifies' the 'indefinite' truth value, making it simply an unavoidable consequence of our inability to refer uniquely to future branches; it has the technical advantage of maintaining bivalence, at least at the theoretical level.

Motivation for yet a third strategy[15] comes from the philosophy of personal identity: in particular, from considerations of 'personal fission', the dividing of an individual into two or more copies (by means of, for instance, dividing the brain in two or using a Star Trek-style teletransporter), which has been used by Parfit (1984) to criticise the whole concept of personal identity. David Lewis, seeking (in Lewis 1976) to reconcile common-sense views of identity with the case of fission, proposed that when fission occurs in an person's life there are two persons present *all the time* — it's just that they have shared person-stages until the fission. Extending this to a branching universe, the natural move is to regard a classical world not as an instantaneous thing (an element of $\mathcal{S}$, in our previous notation), but as an entire history consisting of instantaneous 'world-stages'. It would follow that there is no branching of worlds, simply various periods when worlds cease to share world-stages.

The strategy, which we could call *Lewisian semantics*, might seem to suffer from extravagant multiplying of entities. If there are already multiple identical-at-a-time worlds prior to branching then there are multiple utterances, multiple agents, multiple labs, etc. On the other hand, if we subscribe to an ontological framework where higher-level ontology is to be understood as structures present in lower-level ontology (as defended in general by Dennett (1991) and Ross, Ladyman, Spurrett, and Collier (2006), and in the case of quantum mechanics by Wallace (2003a)), it does not seem unreasonable to regard such a structure as a property of a history rather than an instant — especially as it is likely to be quite temporally extended relative to the branching timescales (at least, this is true in the case of quantum mechanics).

Furthermore, the resultant semantics is extremely straightforward, and has no need of truth-value gaps: "The result of the experiment will be Up", uttered in world w before branching, is true iff in w the result is indeed Up. The simplification is possible because the physical event which serves as a world-stage for the utterance in w is also a world-stage for *another* utterance in another world w'. In one of w and w' the result is Up, in the other it is Down. No rational agent prior to branching should commit to which result he will see, but there is still a fact of the matter.

The difference between Lewisian and ambiguity semantics maps reasonably well onto Sider's (Sider 2001, pp. 60–1) distinction between the 'Worm View' and the 'Stage View' of how our ordinary talk about persisting objects relates to the sequence of temporal parts making up those objects' histories. The Worm View,

[15] Suggested in conversation by various people, notably Simon Saunders.

identifying an object with its entire history, leads (*ceteris paribus*) to Lewisian semantics; the Stage View, identifying an object with an instantaneous temporal part, leads to ambiguity semantics, or possibly to minimalism.[16]

Detailed considerations from philosophy of language are likely to bear on which (if any!) of these three alternatives is preferable. For instance, Lewisian semantics gives a very clear *formula* for reference: in that semantics, an utterance can be regarded as an ordered pair of a microphysical event x and a complete history h, and a referential expression in the utterance refers to whatever thing in h it would have referred to on our conventional view of semantics if we regarded h as the only history. On the other hand, it is unclear what *causal* story could be told as to how this reference occurs (unless we wish to regard causation as itself emergent and branch-relative), since nothing intrinsic about the microphysical x ties it to a particular history.[17]

By contrast, no unique formula for reference can be given at all in ambiguity semantics: indeed, it is an explicit requirement of the system that most referential terms fail to refer uniquely when used of the future. This seems to put ambiguity semantics in direct conflict with what Gareth Evans has called Russell's Principle (Evans 1982): the principle that in order to refer successfully, I must know to what unique thing I am referring.

However, I do not wish to attempt here to adjudicate between the various semantic proposals that I have suggested: further discussion would involve too many controversies and would go beyond the scope of this paper. I rather suspect that there may not really even be any fact of the matter as to which is correct: here perhaps, even if nowhere else, radical interpretation can lead to indeterminacy of translation, as Quine (1960) urged long ago.

7 Propositions

In presenting branching semantics, and the arguments for it, I have chosen to speak of tenseless sentences and their truth-values in the hope of acheiving a presentation as neutral as possible between various different semantic schemes. The conclusion of the charity argument — that one who knows his future involves branching should adopt an attitude of uncertainty towards his post-branching state — does not, so far as I can see, depend on any specific theory of meaning.

One particular semantic proposal, however, seems worth discussing in more detail, both because it is interestingly affected by the assumption of a branching world and because it is currently rather popular in many circles. This is the

[16] A third possible view, not considered by Sider but suggested to me in conversation by Simon Saunders, is that an object at a time should be identified with its history *up to that time*; this view avoids the double-counting of the Worm View in the presence of branching, and leads to the same semantics as for the Stage View given that we are allowing only futurewards branching.

[17] For these sort of reasons, Lewis himself claimed (Lewis 1983, pp. 73–76) that neither one of a pair of individuals who will be undergoing fission can refer to himself and not his twin until the fission occurs; as such he would in effect reject 'Lewisian' semantics! (I am grateful to Oliver Pooley for this observation.)

'possible worlds' semantics developed by Kripke (1981), Stalnaker (1999), Lewis (1975), *et al*, according to which, in crude outline:

- Utterances express propositions;
- The proposition expressed by an utterance is a function of the context in which that utterance is made;
- Propositions are sets of possible worlds, possible 'ways the world might be';
- A proposition P is true at a world w if $w \in P$.

In this framework the difference between tensed and tenseless utterances is that the proposition expressed by the former is dependent on the time of utterance, the proposition expressed by the latter is not.[18]

In branching-time semantics, it is clear, a possible world is 'a way a particular history could be', not a way that the entire branching structure could be. For the difference between the Elite and the Charitable Views, then, is that on the Elite View a possible world is a way that the entire branching structure could be; on the Charitable View it is a way that a single history in the branching structure could be. This has the consequence that in the Elite View most of our everyday utterances do not seem to express propositions at all, whereas in the Charitable View it is straightforward which proposition is expressed, but less straightforward how the truth-value of that utterance is determined from the history-dependent truth-value of the proposition expressed by it.

It is instructive to see how the various semantic proposals of section 6 look in the language of propositions. Lewisian semantics is straightforward: to each utterance is assigned a unique history, and the utterance is true iff the proposition expressed by the utterance is true in the history in which the utterance was made. Minimalist semantics, characteristically, offers simply a brute rule: an utterance made at x is true iff the proposition expressed by the utterance is true in all histories containing x.

Ambiguity semantics is more interesting. The natural strategy would be to regard it as ambiguous which proposition was expressed by a given utterance, and this is particularly natural in the case of proper names which Saunders considered. For suppose that my colleague Fred is to make a quantum measurement with two possible results: Up and Down. Call the two post-measurement Freds Fred_1 (who sees Up), and Fred_2 (who sees Down). If, pre-measurement, I utter 'Fred [tenselessly] sees Up', perhaps it is ambiguous whether my utterance of 'Fred' refers to Fred_1 or to Fred_2, and hence whether the proposition which I express is the proposition that Fred_1 sees Up (which is true) or the proposition that Fred_2 sees Up (which is false). It is less clear where the ambiguity in utterances like 'there will be a war next year' lies, but presumably it lies in the implicit 'on this planet' which comes with such utterances.

[18]Ignoring examples where the time of utterance has relevant 'non-tense' implications: for instance, 'the Queen' in 'the Queen is [tenselessly] called Elizabeth' denotes a different object when uttered in 1850 than in 2005.

I should note that even if possible worlds are usually to be construed as ways a history might be, there will be some contexts in which they must be understood as the Elite View requires, as ways that the entire branching reality might be: in particular, we may want to discuss possible states of that branching reality; we may also wish to use possible worlds in this second sense to ground the semantics of the branch-talk sentences we use when we actually *describe* this branching reality. However, this does not seem particularly problematic either for the Charitable View or for possible-worlds semantics in general.

8 Possibilities

When I say that some future event is possible, sometimes I speak simply of epistemic possibility, so that 'possibly P' means just 'I do not know that P is false'. As we have seen, the charity argument tells us:

- That if I know that it is true that P on some history containing my current location in the branching structure, then I should assent to 'possibly P'.
- That if I know that it is true that P on no such history then I should reject 'possibly P'.

Sometimes, however, we seem to use possibility in a non-epistemic sense (as might be the case for physical, historical, logical possibility, etc.) It is fairly clear that if P is true of some history through x, then 'possibly P' must also be true at x: that which cannot in principle be ruled out should not be impossible. However, nothing in the charity argument entails that the converse holds: that for 'possibly P' to be true at x, P must be true at some history through x.

This is not to say that some philosophically relevant notion of possibility *could* be defined via

'possibly P' is true at x iff P is true in some history containing x;

indeed, as far as I can see anyone tempted by the purely philosophical arguments for branching time should embrace this notion of possibility, for it seems to provide all that the advocates of such arguments desire of 'objective possibility'.

But this is not to say that it is *required* that this particular notion of probability has conceptual significance. As I mentioned in the Introduction, I am myself sceptical that free will, causation *et al* require a branching-time picture. This raises the question of how those analyses of possibility which traditionally have no notion of branching time fare under the branching-world hypothesis.

In this context, recall that Lewis (1986) defined possibility relative to some pre-defined class X of 'possible worlds' (such as the physically permitted worlds, or the physically permitted worlds compatible with some data x, etc.): P is possible relative to this class iff it is true for some world w in the class. (See also Taylor and Dennett (2001) for a development of this position.) For future-directed possibility statements (that is, for statements of the form 'it is possible that Athens will host the Olympics in 2500 AD', rather than of the form 'it is possible that Hitler won the Second World War') it is a requirement that the actual world is in X.

Branching-time semantics makes only a slight alteration to this framework. In such a semantics (ignoring Lewisian semantics) there is no unique 'actual history' but rather a set of such histories: all those through the point x at which the possibility statement is considered. Any future-directed notion of possibility, defined via a set X of possible worlds (i. e. , histories) and assessed at some point x, is coherent only if the entire set of histories through x is contained within X. As far as I can see, this alteration does not have profound effects for the possible-worlds analyses of possibility. (Also note that on Lewisian semantics, everything is almost exactly as in the absence of branching, with exactly one world 'actual' relative to a given utterance.)

9 Conclusion

If the only consistent ways to explain the semantics of tensed talk all assume a certain structural property of time, then that would indeed be a reason to assume that time does have that structure. This is not the case for the question of the linearity or otherwise of time: though I see no reason to doubt that nothing in our experience *rules out* linear time and linear semantics, I have argued here that if time (or 'the world' if you prefer) is actually branching, then the correct semantics must also be branching.

As such, the semantics of tense are more naturalistic than meets the eye. If the right semantics for tense depends on the structure of time, we can be no more certain of our semantic theory than of our theory of temporal structure. (Perhaps this shouldn't be too surprising: if linguistic meaning supervenes on linguistic behaviour and if the correct third-party description of linguistic behaviour depends on facts about the world about which we are ignorant, then our theory of linguistic meaning should inherit that ignorance.)

In particular, if the Everett interpretation of quantum mechanics is true, *our universe* is branching. In such a situation, we are required to assume that the semantics for our language is itself branching. The apparent conceptual paradoxes of the branching-universe are caused by insisting instead, without justification and against the principle of charity, on adopting linear-time semantics.

Once we recognise the semantic implications of the branching-universe hypothesis, we realise that the statement 'either A or B will occur, but not both, and I don't know which', uttered at some point x in the branching structure of the Universe, becomes a perfectly rational thing for someone to say if A occurs in some of the branches futurewards of x and B occurs in the others. A rational agent at x (and who knows the above facts) would assert the truth of 'either A or B will occur' (for in every branch one or the other occurs) and the falsity of 'both A and B will occur' (for in no branch do both occur) but will assert neither the truth nor the falsity of 'A will occur' and 'B will occur', for each occurs in some but not all branches futurewards of x. As such, an interpretation of quantum mechanics in which experiments cause the Universe to branch has no difficulty in explaining why experimenters are entirely justified when, as they routinely do, they express ignorance as to the outcome of their experiments. If so, one of the major conceptual objections to the Everett interpretation of quantum mechanics is unfounded.

References

Albert, D. and B. Loewer (1988). Interpreting the Many Worlds Interpretation. *Synthese 77*, 195–213.

Barrett, J. (1999). *The quantum mechanics of minds and worlds*. Oxford: Oxford University Press.

Belnap, N. (1992). Branching space time. *Synthese 92*, 385–434.

Belnap, N. (2002). Branching histories approach to indeterminism and free will. Available online from http://philsci-archive.pitt.edu.

Davidson, D. (1990). The structure and content of truth. *The Journal of Philosophy 87*, 279–328.

Dennett, D. C. (1991). Real Patterns. *Journal of Philosophy 87*, 27–51. Reprinted in *Brainchildren*, D. Dennett, (London: Penguin 1998) pp. 95–120.

Evans, G. (1982). *The Varieties of Reference*. Oxford: Clarendon Press.

Greaves, H. (2004). Understanding Deutsch's probability in a deterministic multiverse. *Studies in the History and Philosophy of Modern Physics 35*, 423–456. Available online at http://arXiv.org/abs/quant-ph/0312136 or from http://philsci-archive.pitt.edu.

Jeffrey, R. C. (1983). *The Logic of Decision (2nd edition)*. Chicago: University of Chicago Press.

Joyce, J. N. (1999). *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.

Kaplan, M. (1996). *Decision Theory as Philosophy*. Cambridge: Cambridge University Press.

Kripke, S. A. (1981). *Naming and Necessity (revised and enlarged edition)*. Oxford: Blackwell.

Lewis, D. (1975). Languages and language. In K. Gunderson (Ed.), *Minnesota Studies in the Philosophy of Science*, Volume VII. Minnesota: Minnesota University Press. Reprinted in David Lewis, *Philosophical Papers*, Volume I (Oxford University Press, Oxford, 1983).

Lewis, D. (1976). Survival and identity. In *The Identities of Persons*. Berkeley: University of California Press. Reprinted in David Lewis, *Philosophical Papers*, Volume I (Oxford University Press, Oxford, 1983).

Lewis, D. (1979). Attitudes *De Dicto* and *De Se*. *Philosophical Review 88*, 513–543. Reprinted in David Lewis, *Philosophical Papers*, Volume II (Oxford University Press, Oxford, 1986).

Lewis, D. (1983). Postscripts to "Survival and Identity". In *Philosophical Papers, Volume 1*, pp. 73–77. Oxford: Oxford University Press.

Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Blackwell.

McCall, S. (1984). Counterfactuals based on real possible worlds. *Nous 18*, 463–477.

McCall, S. (1994). *A model of the universe: space-time, probability, and decision*. Oxford: Clarendon Press.

Mellor, D. H. (1981). *Real Time*. Cambridge: Cambridge University Press.

Mellor, D. H. (1998). *Real Time II*. New York: Routledge.

Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.

Perry, J. (1979). The problem of the essential indexical. *Nous 13*, 3–21.

Prior, A. N. (1957). *Time and Modality*. Oxford: Clarendon Press.

Prior, A. N. (1967). *Past, Present and Future*. Oxford: Clarendon Press.

Putnam, H. (1973). Meaning and reference. *Journal of Philosophy 70*, 699–711.

Quine, W. v. O. (1960). *Word and Object*. Cambridge, Mass.: MIT Press.

Ross, D., J. Ladyman, D. Spurrett, and J. Collier (2006). *Everything Must Go: Information-Theoretic Structural Realism*. Oxford: Oxford University Press.

Ryle, G. (1964). *Dilemmas*. Cambridge: Cambridge University Press.

Saunders, S. (1998). Time, Quantum Mechanics, and Probability. *Synthese 114*, 373–404.

Savage, L. J. (1972). *The foundations of statistics* (2nd ed.). New York: Dover.

Sider, T. (2001). *Four-Dimensionalism: An Ontology of Persistence and Time*. Oxford: Clarendon Press.

Stalnaker, R. C. (1999). *Context and Content: Essays on intentionality in speech and thought*. Oxford: Oxford University Press.

Tappenden, P. (2000). Identity and probability in Everett's multiverse. *British Journal for the Philosophy of Science 51*, 99–114.

Taylor, C. and D. Dennett (2001). Who's afraid of determinism: rethinking causes and possibilities. In R. Kane (Ed.), *Oxford Handbook of Free Will*. New York: Oxford University Press.

Thomason, R. H. (1970). Indeterminist time and truth-value gaps. *Theoria 3*, 264–281.

Vaidman, L. (1998). On Schizophrenic Experiences of the Neutron or Why We Should Believe in the Many-Worlds Interpretation of Quantum Theory. *International Studies in Philosophy of Science 12*, 245–261. Available online at http://arxiv.org/abs/quant-ph/9609006.

Wallace, D. (2002). Worlds in the Everett Interpretation. *Studies in the History and Philosopy of Modern Physics 33*, 637–661. Available online at http://arxiv.org/abs/quant-ph/0103092 or from http://philsci-archive.pitt.edu.

Wallace, D. (2003a). Everett and Structure. *Studies in the History and Philosophy of Modern Physics 34*, 87–105. Available online at http://arXiv.org/abs/quant-ph/0107144 or from http://philsci-archive.pitt.edu.

Wallace, D. (2003b). Quantum probability from subjective likelihood: Improving on Deutsch's proof of the probability rule. Forthcoming in *Studies in the History and Philosophy of Modern Physics*. Available online at http://arXiv.org/abs/quant-ph/0312157 or from http://philsci-archive.pitt.edu.

Wallace, D. (2005). Epistemology quantized: circumstances in which we should come to believe in the Everett interpretation. Forthcoming in *British Journal for the Philosophy of Science*. Available online from http://philsci-archive.pitt.edu.

Williamson, T. (1994). *Vagueness*. London: Routledge.