Branching and Uncertainty

Simon Saunders[1] and David Wallace[2]

June 2007

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 sentences uttered in branching worlds have much the same truth conditions as they do 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 Metaphysics

Lewis's answer to Parfit's central puzzle on identity in the face of fission, or *branching* as we shall call it, has won few supporters. But this may be because his argument (Lewis 1976) took a wrong turn, rather than that it was inherently wrong-headed.

In the face of branching, Lewis proposed that we say two persons are present throughout, even prior to branching. This is equivalent to the stipulation, physicalistically, that 'persons' and 'continuants' are non-branching but possibly overlapping spacetime worms. There are plenty of homely analogies: the Chester A. Arthur Parkway, he observed, merges with Route 137 for a brief stretch, but still there are two roads.

But as the spatial example shows, there are times when we also say that there is only *one* thing present – for example, when saying how many roads (to someone with an infirmity) a man has to cross. Even more so in the temporal case: surely we want to say, prior to branching, that there is only *one* person present. The answer is that indeed sometimes we want to say the one thing, sometimes the other, depending on context. 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, we should say there are two. 'Local', spatially, means we count the roads as identical-at-x iff they

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

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 a temporal part at time t. The latter he called 'tensed identity'. With all this we are agreed.

But Lewis went on to attribute thoughts (and we must suppose utterances) at time t to the common temporal parts of speakers at time t. Why so? In the absence of branching, utterances are naturally attributed to persons, which, from a four-dimensional point of view, are either maximal continuants or stages. To attribute them to stages that are not persons would seem to be perverse.

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 themselves. 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.

Lewis ruled out this semantics peremptorily. 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, if the referent of 'I myself', thought or uttered at time t (at temporal part S) is a continuant, 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 they say: '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 can know the truth value of his utterance, for neither can know which of C_1 and C_2 he is.

That makes the branching of persons curiously similar to mere divergence, where the common segment of two branching persons is rather counted as two, numerically distinct segments, of disjoint worlds. Lewis was at pains in *On the Plurality of Worlds* to distinguish branching from divergence, whereas on

our proposal we are driving them together. Of course they are quite different taken as processes within a single, non-branching world: the challenge is to say how they differ when it is worlds that branch or diverge as wholes. As Lewis recognized, when it comes to the initial segments of diverging worlds that from the point of view of branching count as one, the similarity relation ex hypothesi holds exactly; ' exact similarity' is a transitive relation; in this limit, the formal discrepancy between similarity and identity disappears (Lewis 1986 p.209). That observation effectively neutralizes the intuition that branching is different from divergence because it yields de re transworld identities: replace identity with exact similarity, and one has a relation which functions exactly the same.

Lewis has two objections to branching worlds. The first is that it makes nonsense of ordinary beliefs about the future – that is the worry we have so far been dispelling – and the other is the 'the problem of intrinsic accidentals'. If we are right that there is little difference between branching and divergence on the semantics we are considering, the latter problem should arise equally for diverging worlds. The problem is this: if what is possible for C are the properties that C has in all possible worlds that contain C, then C has all its accidental intrinsic properties in all worlds possible for C. So it has all its accidental intrinsic properties necessarily. And indeed, the same problem arises for exact counterparthood as for identity, if what is possible for C are the properties that C has in all possible worlds in which it has an exact counterpart.

Lewis solves this problem by defining modal properties by an approximate counterpart relation, not the exact one that had the same formal structure as identity. No similar retreat is possible for branching worlds if de re identities play the fundamental role in the modalities. Does this sharply distinguish overlap from branching? Not really; as Lewis aknowledged, the de re identities do not have to play this fundamental role, no more than the exact counterpart relation does for diverging worlds. But anyway, on either framework there is a ready solution to the problem of intrinsic accidentals in which de re identity, respectively exact counterparthood, *does* play a fundamental role to modal properties: it is enough that there be identities among temporal parts. Thus: a thing C might have had property P insofar as it has a temporal part of a thing C' which has property P. Al Gore might have won the 2000 US presidential election, so long as he has a temporal part which is a part of a person who won – or, using counterpart theory, so long as he has a temporal part exactly similar to a part of a person who won.

One can put the matter in terms of tensed identity instead: branching and overlapping worlds are only made to look the same by using tensed identity. Elsewhere Lewis wrote of tensed identity as a somewhat unnatural and ad hoc device (Lewis 1986 p.218-19), as applied to fission within a world. But it is not remotely ad hoc or unnatural at the level of branching worlds. If it is the world as a whole that is branching, there is never any question of observing the difference between branching as opposed to diverging worlds. One is, moreover, always concerned with just one world. If Al Gore had won the election he would have tried to save the Amazonian rainforest. Which rainforest did he want to

save? Why, *his* rainforest, the one (unique continuant, in a unique world) to which he (unique continuant, in a unique world) is spatiotemporally related. 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.

Whatever Lewis's reasons for insisting that thoughts are shared on overlap, and must have the same content, they went unstated. The claim is incompatible with ordinary language (which attributes thoughts to persons), and demands a strange attitude to personal survival (we should desire that 'at least one of us survives'), in conflict with the use of tensed identity (which he otherwise favoured). He gave no consideration of the more conservative and simpler semantics that we are proposing. Indeed, let thoughts be attributed to persons, as maximal continuants, and the semantics of branching and diverging worlds hardly differ.

We have deduced that on this semantics one *cannot say which person one is*, and hence that *one cannot say what one's future is*. Nor is there anything peculiar to the first person in such admissions: so long as everyone is subject to branching, I don't know which branch is *yours*, you don't know which branch is *theirs* (because we don't know which branch is *ours*). That is, on this semantics, branching can be understood as *objective uncertainty*. That is an idea that was never suggested by Lewis.

2 Physics

If not suggested by Lewis, it will hardly have been suggested by his critics. The important remaining distinction between branching worlds and diverging worlds, not so far canvassed, arises if only *one* of all the worlds is real. In that case perfect similarity cannot have the same properties as identity after all (the real world may be perfectly similar, as goes an initial segment, with an abstract or otherwise ersatz possible world; but it cannot be identical with it). In other words, branching as an alternative to divergence is only available to modal realists. And our conclusion, that a modal realism, with overlap, carries with it a notion of objective uncertainty, only opens up further questions. Can it be quantified as probability? If so it appears to be a single-case probability, that is as yet unrelated to statistics. In view of the unexpected similarity of branching with divergence, it is presumably a form of probability that can be analyzed in terms of the exact counterpart relation; but then it appears to depend on the whole space of possible worlds. And what about the converse of branching, recombination, or converging of possible worlds?

These are deep waters for metaphysicians. No wonder branching has been looked on with suspicion by all parties to the debate.

All parties, that is, save one. 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 claim to have a serviceable and 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 that it is ubiquitously a *branching* process, not a recombining or converging of worlds, at least at the macroscopic level.

The theory in question is Everettian quantum mechanics (EQM), one of the three realist theories (or classes of theory) of the microworld that extend smoothly to the macroscopic. The other two are hidden-variable theories, such as the de Broglie-Bohm theory, and dynamical state-reduction theories, such as the Ghirardi, Rimini, and Weber theory. But neither dBB or GRW theories have any satisfactory relativistic formulation, despite years – decades in the case of dBB – of attempts to define one. EQM is the only theory so far on the books that is empirically adequate to particle physics. That is because it is no more than an interpretation of the extant formalism, one which rests on no special feature of the non-relativistic theory. dBB and GRW theories, in contrast, supplement or modify the formalism, in a way that only makes sense (or has to date only made sense) in the non-relativistic context.

The key additional feature of EQM, in the broad brush of metaphysics, is that branches come with *weights* – quantities whose sum is preserved under branching. That, and the interpretation of branching in terms of objective uncertainty, is not quite enough to ensure that weights can be identified as probabilities, but it goes a long way in that direction (what more is needed is provided in EQM). But no condition 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.

That evidently transforms the nature of the argument, on both sides. EQM has long been criticized on the grounds that it can make no sense of probability. That it makes quantitative sense – that branch weights *should* be identified as probability and not some other quantity – is an argument that has now been spelt out in a series of papers (see Wallace 2005a, Saunders 2005, and references therein). There remains, however, the *incoherence* objection: the objection that the theory has no place for the notion of uncertainty (see e.g. Greaves 2004). But uncertainty, as we have seen, is the right epistemic attitude to branching, if thoughts and utterances are ascribed to persons. And on the other side, the known unknowns of a modal realism with branching worlds are solved or bypassed by the physics. We had uncertainty, which we had no idea how to quantify, and recombination or reconvergence of worlds, which we did not know how to conceptualize; but EQM provides weights, and the dynamics leads to branching (at the macroscopic level) taking place always with the same sense.

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. Or that it is contrary to some other that is well-accepted – for how better to account for the fact that we have had such difficulty in understanding quantum mechanics correctly? Still less that there may be some other, alternative semantics available: perhaps there is, but we are not looking for any deep metaphysical truth in our choice of semantics. We are looking for serviceability.

The criticism can only be that the semantics we are proposing is *unintelligible*, or *inconsistent*, or that it is contrary to such basic assumptions in philosophical logic or philosophy of language as to call into question too much else that we take to be true. But that charge is clearly unfounded. The semantics as given is simplicity itself: it is implausible that it harbours any inconsistency. Nor is it contrary to basic assumptions in philosophical logic: it preserves bivalence, it preserves the platitude that utterances are attributed to persons, and it preserves the standard form of four-dimensionalism, that persons are maximal continuants. As for basic assumptions of philosophy of language, the view we are proposing is the one positively *enforced* by them, in that meanings are determined by use (Wallace 2005b): if EQM is true, words like 'uncertainty' are used just where there is ignorance of the ordinary sort, or where there is branching – essentially, self-locating ignorance – so that is what those words mean.

Granted all of this, there remains a question. If there is so little to choose between branching and divergence, why not settle for divergence? Given which: why not throw away all of the worlds save one? Doesn't this open the possibility of a one-world solution to the problem of measurement, which like EQM preserves quantum mechanics unchanged?

We should first confirm the suspicion that there is indeed little to choose between branching and divergence at the level of the representation of worlds. The standard device for the latter is to define a *history space* - a collection of histories, each of them a unique continuant, a maximal sequence of world-stages, where each stage is represented mathematically by a projection operator. Fixing the set of projection operators, all possible histories constructible from them are considered: the universal state then dictates their relative weights. But from that one routinely passes to a representation of the universal state in terms of a tree-like, dendritic structure, by collecting together histories all with identical initial sequences of projection operators, and passing from relative weights of histories to relative conditional weights of histories, or relative weights of branches. In applications of the formalism of quantum mechanics one passes back and forth between the two representations without any ambiguity or conceptual difficulty.

But that doesn't mean all the histories save one can be discarded. For the key provision of EQM is that this branching structure to the state, or representation of the state in terms of a history space, is only *one* (albeit important) structure to the universal wave-function, among various others – and that it is an emergent, approximate structure, arrived at in methodologically similar ways to emergent ontology across the board in the special sciences and in the specialisms of physics. These 'quasi-classical histories', so-called, are emergent just like proteins in chemistry are emergent: they do not have a separate axiomatization, a unique and exact description, separated off from the larger dynamical story of which they are a part. Hence neither does any single one of them. The larger dynamical story of which quasi-classical worlds are parts is the wave-function of the universe, the fundamental object of the theory, and the standard against which other descriptions can be judged as approximations.

To conclude: if – as Lewis claims – in cases of branching there are two persons present even before the branch, then 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. 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

Greaves, H (2004), 'Understanding Deutsch's probability in a deterministic multiverse', *Studies in the History and Philosophy of Modern Physics,* **35***,* 423-56*;* available on-line at http://philsci-archive.pitt.edu/archive/00001742/.

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 University Press, 1983.

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

Saunders, S. (2005), 'What is probability?', in *Quo Vadis Quantum Mechanics*, A. Elitzur, S. Dolev, and N. Kolenda, eds., Springer. Available on-line at http://philsci-archive.pitt.edu/archive/00002011/.

Wallace, D.(2005a), 'Quantum probability from subjective likelihood: improving on Deutsch's proof of the probability rule', forthcoming. Available on-line at http://philsci-archive.pitt.edu/archive/00002302/.

Wallace, D. (2005b), 'Language use in a branching universe', forthcoming. Available on-line at http://philsci-archive.pitt.edu/archive/00002554/.