Replies to Peter Lewis "Uncertainty and Probability for Branching Selves":

Comments from APA session 2005

I should say first, that I thought that this was a stellar paper. I learned a lot from it. In fact, I was so convinced by one part when I finished reading it, that I didn't believe that I had ever thought otherwise. I thought to myself 'surely I didn't say there is uncertainty about the future in an Everett universe', but when I went back and looked at my own paper, sure enough, I said it. In fact, I wrote in the *abstract:* "I propose, in the context of Everett interpretations of quantum mechanics, a way of understanding how there can be genuine uncertainty about the future". To be fair, I also say

"The reason it is easier to get worked up about the problem of probabilities in a genuinely branching universe is that measurement events in that setting don't (in general) have unique outcomes: a single one has a different result at each post-measurement branch of the universe, and because each is symmetrically related to the event itself, if there is more than one, there is no way *before* branching of singling out any, that we may wonder of it, what its outcome will be. I can stand in front of my Stern-Gerlach apparatus, having fed my electron in, and wonder *de dicto* whether there will be a spin-up result, but I cannot wonder *de re*, of a particular outcome, whether it will be spin-up, because I have no independent means at my disposal – descriptive, indexical, or otherwise, from that perspective - of picking it out. It is only by looking back, after the fact, from a location in the universe that bears asymmetric relations to the various results (all equally real, all reflexively actual) that I can single out those to which a particular set of Born probabilities attach. It is only looking back that I can be surprised by the way things turned out in my branch, because it is only from my post-measurement position that I bear the outcomes in my branch the kinds of asymmetric relations that allow me to make identifying reference to them."

But I don't think I can claim the inconsistency as a credit to myself. In any case, he's quite right about that point, and I want to concede at the outset. There *isn't* genuine pre-measurement uncertainty in a literally branching world because there is nothing, pre-measurement, to be uncertain about. So what I should have said is that there can be genuine uncertainty in the sense of 'facts about my present state that I can't determine on the basis of complete knowledge of the entire history of the world including my own past, and complete knowledge of the physical laws.'[i],[ii]

That said, I'll spend some time talking about other points in the paper that I'm less sure that I agree with. I'll raise questions about some of the more important issues, and then turn the table over.

My view

Let me first say what my own project was; to find a physical interpretation for the Born probabilities, which is to say, to find something in the Everett Universe for the Born probabilities to represent. The proposed interpretation is this. In an Everett Universe, the Born probability of an up-result in, say, the spin measurement above is the probability of observing an up result here, where 'here' picks out a location in one of the post-measurement branches in the universe, and picks out a different one in the mouths of each of my post-measurement successors. It can't be interpreted as assigning expectation values to future events for the reasons we have seen, but it can (and I am proposing, *should*) be interpreted as assigning degrees of surprise to events from a situated post-measurement perspective. Degree of surprise is, after all, just the retrospective analogue of expectation. Quantitively, it is the probability of observing an up-result conditional on the pre-measurement state of the universe, it tells each observer how typical her situation is among counterparts that share her pre-measurement history. Whereas in a Copenhagen interpretation, we would say that the Born probability of the outcome of a measurement M being e is equal to the probability of e conditional on the pre-M state of the world (better, Pr(e/the contents of M's past light cone)), now we say the Born probability is the probability that an arbitrary post-M branch is an e-branch, or the probability that a post-M observer is an e-observer. Measurement results aren't any more predictable on the basis of purely historical information in an Everett Universe than they are in a Copenhagen one in the sense that the result observed in any branch is not determined by the global dynamical laws from the information contained in its own back light cone.

The probabilities arise not because of indeterministic dynamical evolution, but because the information that one would need in order to use the dynamical laws to predict the result that is showing in a given branch – *viz.*, information about the distribution of outcomes across other branches – is unavailable to inhabitants of that branch.[iii] They arise, that is to say, not because of the way information is spread back to front (as they do in a deterministic universe); but because of the way it is spread side to side.

Empirical Incoherence

Lewis' central claim is that unless an interpretation of the Born probabilities can be provided in the context of an Everett universe, the view is empirically incoherent because it is the agreement between the observed frequencies of various measurement results and the probabilistic predictions of quantum mechanics that gives us reason to think the quantum mechanics is true. This is a common charge against Everett, and of course, in one sense, it is correct. We need to give some physical significance to the Born measure that ties it to frequencies in a way that allows it to be confirmed by those frequencies. In that sense, the requirement is satisfied by the proposed interpretation. The theory constrains long-term frequencies of measurement results in the universe as a whole and, indeed, along each branch (given the customary assumptions about sampling) in the sense that appreciable divergence from the Born probabilities occurs in a vanishingly small set of branches. If we write down two-dimensional models of the Everett Universe and the Copenhagen universes side by side using solid lines to represent actual history and dashed lines to represent dynamically possible, but *non*-actual histories we find first, that the only difference is that the Everett universe replaces dashed lines with solid ones. And second, that they come with symmetries that guarantee that long term frequencies of results of measurements of type M in typical branches approaches the distribution of results across any lateral cross-section of the universe just after an M-measurement has been carried out.[iv] And this means that the Born probabilities, interpreted as per above will relate in the exactly the same way to frequencies of results observed along each branch. That's not an accident, of course, it was this insight that inspired the interpretation, and the empirical adequacy of Everett's interpretation depends on it.

There are, of course, both technical questions and problems of other kinds (most seriously, of course, the problem of justifying the Born measure and the preferred basis problem). But provided that all of that can be worked out, I don't see why the theory can't be confirmed, and confirmed in the ordinary way by information gathered by situated observers about the frequencies of results of various kinds along their branches. Long term frequencies that diverge appreciably from the Born Probabilities would disconfirm (lend negative empirical support to) Everett's theory as surely as they would disconfirm a Copenhagen interpretation. The ordinary mathematical framework relating probabilities to frequencies applies no matter how the probabilities are interpreted. It just happens that, in the Everett universe, the information observers gather along their own branch will tell them a lot about the distribution of results along a lateral cross section of their universe and that will underwrite their degrees of surprise in observing the particular results that they observe.

Lewis interprets the requirement that a theory has to generate predictions to be confirmable in a different way than I have. I have interpreted it as meaning simply that the theory has to have testable empirical consequences, but Lewis gives the notion of prediction a temporal spin. If I understand him, he grants that Born probabilities can be regarded as a quantitative measure of degrees of surprise of post-measurement observers,[v] but he thinks that is not enough to defuse the charge of empirical incoherence because, as he says,

"… not just any account of probability will do … the probability measure must be one that I can use *before a measurement* to ascribe probabilities to the possible outcomes. Otherwise, the probabilities can't be used to confirm the theory."[vi]

Here, I confess, I simply don't understand the source of the requirement. A theory generates predictions if it has empirical consequences and the theory is confirmed to the degree that the empirical consequences are fulfilled.[vii] If a theory has to make predictions about events before they occur in order to be confirmable what do we say about theories with a purely historical subject matter? What about theories of continental drift, cosmological theories about the early history of the universe, or Evolution? I'm quite sure that there must be something that I'm missing here. In any case, the theory does make predictions, even in Lewis' sense, if the Copenhagen interpretation does. The points of contact between theory and evidence on the two theories are precisely the same.[viii]

At this point, one might well say, "that's all fine and good, but are the born probabilities, interpreted as degrees of retrospective surprise, *probabilities*?" I see no reason to deny them that title. There are various points at which Lewis seems to suggest that the term probability is only properly applied to 'ontic chances', by which he means the outcomes of genuinely indeterministic processes. He's free, of course, to adopt the terminological stipulation, but then he needs a new term to refer to the kinds of probability we have, for example, in statistical mechanics (which, not incidentally, bear a strong resemblance to Everett probabilities on my proposal), and then he owes us an argument for the claim *"if there is no place for probability within Everett's theory, then it cannot make predictions about the relative frequencies of measurement results."*

In sum

So the bottom line is that I concede there is no pre-measurement suspense, no well-defined question to which an Everettian measurer can assign a Born probability in advance. That is an important, and appreciated, correction. I propose instead that we interpret the Born probabilities as retrospective degrees of surprise, and dispute the charge of empirical incoherence.

Sleeping beauty

Some remarks about sleeping beauty, if there's time. Very briefly, I'm not at all tempted to appeal to a reflection principle or intertemporal consistency constraints to generate pre-measurement uncertainty. Such constraints apply only in the case that one has a unique continuant. Indeed, if you know beforehand that you will have multiple successors that will form correct, but mutually incompatible beliefs, applying a reflection principle to form your current opinion would quickly generate inconsistency. As to the challenge at the end of the paper; the relevant question isn't whether sleeping beauty should change her probabilities upon discovering what the scientists plan to do for what will happen to her on Tuesday. If she fully expects to attend a concert on Tuesday, she will continue to do so, but she will be less than certain upon waking whether she should expect to attend a concert *today*.

SB & Everett; in SB there are two questions whose answers diverge: what is the chance that the coin will come up heads, and what is the probability that a post-toss waking event was preceded by a <heads>? The first is a single case probability, the second is an indefinite probability, both available before measurement. The pre-measurement indefinite probability than an arbitrary post-waking event was preceded by a <heads> toss determines the post-measurement single case probability that *this* event was preceded by a <heads> toss. Likewise in Everett, except that there's no divergence between the chances and the pre-measurement indefinite probabilities. That is a feature of the SB cases that hinges on the way the thing is set up. In the Everett case, there is also an indefinite probability of obtaining a positive result in any post-measurement branch of the universe that turns into a definite post-measurement probability of having obtained a positive result in *this* branch. In the Everett case, one might want to deny that there is a pre-measurement indefinite probability, but I'm not sure on what grounds; indefinite probabilities pertain to event types (the indefinite probability of A/B is approached by the relative frequency of A's among B's in large typical ensembles) and there is no problem with identifying event types, both pre- and post-measurement in the Everett universe. The difficulty is with identifying particulars to which to attach a definite probability pre-measurement; it is these that are ill-defined from a pre-measurement perspective.

---

[i] This is true only in the Everett universe that *literally* branches. There is pre-measurement uncertainty in the merely metaphorically branching universe because 'here', understood as a branch indexical, is well-defined in that context.

[ii] There's a twist here; if one holds that branching occurs at the moment that one becomes conscious of a measurement result, there is no uncertainty post-measurement, either. The answer to the question arises is the one at which the answer becomes available

[iii] Notice that there's going to be a residual probability unless we insist that the frequencies perfectly match the probabilities. If we do that, and we do know everything about the results across other branches, then we have to know the result in the branch that we occupy. Questions about measurement results in an Everett universe are pure self-location, where selves are as short-lived as the times between measurements.

[iv] with increasing probability, to an increasing degree of approximation, as size of ensembles increases, granted the customary assumptions about sampling. Cf., Finkelstein's quantum mechanical version of Bernoulli's Theorem.

[v] "I do not doubt that there is genuine uncertainty here… My concern is whether such a measure could be predictive."

[vi] p. 3, italics mine.

[vii] Some people think there is an asymmetry in the confirmational value of evidence acquired before and after a theory is proposed. Such people owe an explanation of why this is so because confirmational relations, as usually understood, have to do with the degree of fit between evidence and theory and this is indifferent to the way the evidence is spread through time. It is often suggested that theories that are cooked up after the fact to account for evidence in hand are more likely to be *ad hoc* or unprincipled, but this can't provide a foundation for the general requirement that Lewis imposes.

[viii] David Lewis, in his last paper, argues that if the Everett interpretation is correct, it should be possible to obtain evidence that discriminates it from a Copenhagen interpretation. I am unconvinced by the argument, but even if (David) Lewis is correct, his conclusion does nothing to support (Peter) Lewis. To claim that observers in an Everett universe can come to possess confirmatory evidence that favors Everett over Copenhagen is to deny that the theory suffers from empirical incoherence even more strongly than I have done.