Macroscopic Ontology in Everettian Quantum Mechanics

Wilson, Alastair

Document Version
Peer reviewed version

Link to publication on Research at Birmingham portal

Macroscopic Ontology in Everettian Quantum Mechanics

Alastair Wilson
University College, Oxford

Abstract

Simon Saunders and David Wallace have proposed an appealing candidate semantics for interpreting linguistic communities embedded in an Everettian multiverse. It provides a charitable interpretation of our ordinary talk about the future, and allows us to retain a principle of bivalence for propositions and to retain the Law of Excluded Middle in the logic of propositions about the future. But difficulties arise when it comes to providing an appropriate account of the metaphysics of macroscopic objects and events. In this paper, I evaluate various metaphysical frameworks which might be combined with the Saunders-Wallace semantics. I conclude that the most appropriate metaphysics to underwrite the semantics renders Everettian quantum mechanics a theory of non-overlapping worlds.

1. Introduction

Simon Saunders and David Wallace [2008] have proposed a distinctive candidate semantics (for short: the SW semantics) for interpreting linguistic communities that are embedded in an Everettian multiverse[1]. On their view, further elucidated in Saunders [2010], all material objects and events, including linguistic tokens, are branch-bound - that is, each one is uniquely related to some particular complete branch of the quantum state. This account is semantically conservative, in that it allows for utterances to be true or false *simpliciter*, with no need for relativization of truth-value. Previous treatments of the semantics of talk about the future in Everettian quantum mechanics (EQM) have involved single utterances with complex semantic properties[2]; the SW semantics instead involves multiple utterances, each with simple semantic properties.

Saunders and Wallace [2008] do not give a complete formal breakdown of the SW semantics, but the following principles should give a general idea of how it works. In this reconstruction I've used a roughly Kaplanian framework.

> **T**: The primary bearers of truth-values — propositions — are sentences at contexts.
> **U**: Utterances express propositions; each utterance corresponds to a pairing of some concrete vocalization with some particular context.
> **C**: Contexts are centred complete branches of the quantum state.

T, U and C in combination entail that where there are multiple complete branches, there are multiple contexts and hence multiple distinct utterances. To reconcile such multiplicity with a charitable interpretation of ordinary thought and talk, we require that the quantifiers in utterances are restricted:

> **R**: Referring terms in utterances normally range over a domain containing only entities in the complete branch of the utterance's context.

Assume for simplicity that an agent has complete knowledge of the global properties of the quantum state; any remaining ignorance that the agent

---

[1] In this paper, I will presuppose along with Saunders and Wallace that a decoherence-based solution to the 'preferred basis problem' is successful, and that it provides an Everettian with the resources to solve the quantum measurement problem. For defence of this claim, see Saunders [1993, 1995].

[2] See, for example, Wallace [2003] and Tappenden [2008].

might have is cashed out according to the SW semantics as *self-locating* ignorance, ignorance about the agent's precise location within the quantum state. Much of our ordinary ignorance of contingent[3] fact can then be thought of as self-locating: for example, our ignorance about whether a sphere of gold a mile in diameter exists becomes ignorance about whether or not we belong to a mile-diametered-gold-sphere-containing branch. Similarly, an utterance of 'I exist' typically refers to exactly one person, located in the branch of the utterance's context. However, due to my incomplete knowledge of the properties of the complete centred branch in which I am located, I will typically be ignorant of many of the properties of the person to whom I refer using the term 'I'.

The purpose of this semantic framework is to recover a robust notion of self-locating ignorance, prior to a quantum interaction, which can underwrite an agent's assignment of objective probabilities to specific interaction outcomes even if that agent knows that quantum mechanics is true and that the quantum state evolves globally in accordance with the Schrödinger equation. This self-locating ignorance can then help provide a solution to the 'incoherence problem' (in the terminology of Wallace [2003]) of explaining how EQM makes room for non-trivial probabilities. In turn, a solution to the incoherence problem opens the way for a solution to the 'quantitative problem', of explaining why these probabilities should match those given by the Born Rule[4].

I think that the SW semantic proposal is promising: it allows for a charitable interpretation of our ordinary talk about the future on the supposition that we are living in an Everettian multiverse, and it allows the Everettian to retain a general principle of bivalence for propositions, and to retain the Law of Excluded Middle in the logic of propositions about the future. But difficulties arise with providing an appropriate account of the metaphysics of macroscopic objects and events which can underwrite the SW semantics. In this paper, I will evaluate various metaphysical pictures we might seek to combine with it.

The plan is as follows. In section 2, I discuss a view which I call 'Literal Fission'. Although it looks at first like a natural way for an Everettian to interpret the quantum formalism, I argue that Literal Fission in combination with the SW semantics cannot ground all the uncertainty-claims needed for a non-revisionary account of objective probability. Section 3 considers an

---

[3] It is commonly presumed that the laws of quantum mechanics, and the initial conditions of the universe, are contingent; uncertainty of these facts is not self-locating uncertainty according to the present proposal. Might this provide an Everettian with reason to think that such facts might after all be non-contingent?

[4] The Born Rule is the recipe for assigning probabilities to distinct quantum-mechanical outcomes. Much progress has been made on a decision-theoretic solution to the quantitative problem in recent years; see in particular Wallace [2010].

alternative view, 'Macroscopic Pairing', which has been proposed by Saunders and Wallace [2008] and by Saunders [forthcoming]. I raise some concerns about this picture, on the grounds that it requires a revisionary mereology and property-structure for macroscopic objects. In section 4 I outline a third account of macroscopic ontology for Everettians, which I call 'Space-time Point Pairing'. I argue that Space-time Point Pairing can underwrite the application of the SW semantics while preserving an orthodox mereology and property structure for macroscopic objects and events, and recommend it to the Everettian who takes EQM to be a branching theory. Section 5 is a discussion of some general objections which arise for both pairing proposals; section 6 takes a step back and examines the overall framework in which they are situated. The view which I eventually recommend to the Everettian in section 7 rejects branching of Everettian worlds in favour of divergence, and has a structural resemblance to Lewisian Modal Realism[5].

2. Literal Fission

One picture that has often been associated with EQM identifies persisting macroscopic objects and events ('continuants') with non-branching parts of a branching structure of consistent histories[6]. Call the parts in question 'branch segments', and call complete big-bang-to-end-of-time consistent histories 'branches'. Each continuant then has branch segments in common with continuants in some other branches, and (on the assumption that each branch contains the big bang as an initial segment) each branch has branch segments in common with each other branch. I will refer to this part-sharing metaphysical picture as *Literal Fission*. ('Macroscopic', in this context, should not be read simply as 'large'. Macroscopic objects and events will be taken to be those for which quantum entanglement is negligible, and which can be described classically to a good approximation.)

Why think that Saunders and Wallace mean to endorse Literal Fission? Because it appears to be forced on us if we take literally their apparently unrestricted use of the term 'temporal part' and interpret this term (naturally enough) as referring to branch segments. For example, in discussion of how their proposal might account for *de re* modal claims, Saunders and Wallace consider the following suggestion:

> 'C might have had P' is true if and only if C has a temporal part that
> is a temporal part of a continuant C' that has P... Thus Al Gore

---

[5] See Lewis [1986].

[6] For an introduction to the consistent histories formalism for quantum mechanics, see Griffiths [2001].

> might have won the 2000 US presidential election, if he has a temporal
> part which is a part of a person who won.
>
> <div align="right">Saunders and Wallace [2008], p.298</div>

This use of temporal-part language suggests that branch segments which are shared between continuants in different branches will themselves be continuants located in *both* branches. It is hard to resist identifying such shared continuants with persons and things, contrary to the SW semantics. See Figure 1 for an illustration of this: continuants A and B share branch segments up to branching, and one of their shared branch segments, C, is located both in A's branch and in B's branch. Note that this proposal differs significantly from standard counterpart theory for *de re* identities: for example, it does not allow for *de re* identities

Figure 1

between individuals with no exactly similar temporal parts. I mention it not to defend it, but to highlight the use it makes of temporal-part language. (It should also be noted that Saunders and Wallace [2008] only tacitly endorsed this proposal, and that neither wish to defend it.)

In their discussion, Saunders and Wallace focus on utterances such as 'I will see spin-up', uttered when an agent is about to perform a quantum measurement with two possible outcomes: spin-up and spin-down. Here the combination of the SW semantics with the Literal Fission metaphysics gives unproblematic results. There are two continuant agents, one per branch, which share a common part before the measurement; there are two utterances, each attributed to one particular agent. One utterance is true, and one false, because one agent does in fact see spin-up, while the other in fact sees spin-down. At the time of the utterance, each agent is uncertain about which agent he is, and correspondingly uncertain about which utterance he has made – the true one or the false one. He is therefore uncertain of the truth-value of his utterance.

So far, so good. I will assume that no significant problems are generated by the extension of this approach to an agent's uncertainty about propositions such as 'the apparatus will read spin-up', which goes via the idea that an agent's quantifiers are restricted to range over that agent's branchmates. So where there are two agents, there are also two sets of apparatus, one of which will read spin-up and the other of which will read spin-down. An agent can be uncertain about which apparatus he is referring to, because he is uncertain about which one is a branch-mate of his. All of this is compatible with the Literal Fission metaphysics.

The problem I want to focus on is that Literal Fission does not generalize to give us an account of an agent's uncertainty about events occurring futurewards of that agent's death. Consider, for example, the claim 'at least one sea-battle will occur after my death' (call this claim SB)[7]. I take it that after my death, various chance events will occur which affect the objective probability of sea-battles occurring. Perhaps a year after my death crucial peace talks will occur between some feuding island nations. Hence, I take it that we ought to be able to recover uncertainty about SB at some time t prior to death, and to assign SB non-trivial objective chance at t.

However, uncertainty about SB cannot be accounted for by Literal Fission. For I, as a complete continuant, will appear both in branches containing sea-battles futurewards of my death, and in branches containing no sea-battles futurewards of my death. Both sea-battle branches and non-sea-battle branches contain the aggregate of branch segments which makes up my continuant. See Figure 2 for an illustration of this. I am wholly located in both branches, so how can I be uncertain about which branch is mine? Something has to give.

Figure 2

One potential way out would be to use the SW semantics to account for propositions about future events occurring before my death, and a different semantics to account for propositions about events occurring futurewards of my death. But this is an unattractive move: if there is an unproblematic semantics available for propositions about events futurewards of an agent's death, why not apply it more generally and avoid any dualism? The difference between these types of proposition doesn't seem to be linguistically marked; it would be better to have a single account to cover all cases.

Another option for the defender of Literal Fission might be to adopt an error theory of our claims about post-death uncertainty. Perhaps it simply does not make sense to talk about uncertainty about events occurring after our deaths, or probabilities for such events, and we have been mistaken to think that it does. According to this view, it is just impossible for me to wonder about what will happen after my own death, since there is no unique future for me to wonder about. I can still coherently wonder about whether certain events are open possibilities at the moment of my death; but I cannot coherently wonder about whether they *will occur*.

---

[7] I'm neglecting complications due to special relativity - but we could easily use instead the claim that a sea-battle occurs in the forwards light-cone of my death.

Although at first glance this response seems to concede far too much, I do not think it can be dismissed as easily as the dualistic semantic proposal just discussed. In particular, it is hard to see what would count as empirical evidence against the error theory. All of our empirical evidence for quantum mechanics consists of past observed frequencies, and objective probabilities for past observed frequencies *can* be made sense of even on the Literal Fission picture. While it would require some revisionary metaphysics and involve potential conflicts with interpretative charity, I do think that the combination of the SW semantics with Literal Fission remains a live option for those who are prepared to give up on objective chances about events futurewards of their own deaths. However, I will say no more about this possibility.

A referee asks whether the Literal Fission account suffers from the 'Assertion problem' of Belnap et al [2001]; the problem within a branching-time theory of stating the truth-conditions (or correctness-conditions) for an utterance of (say) 'there will be a sea-battle tomorrow'. In the preferred framework of Belnap et al. this problem looks serious, as they take each individual utterance to be part of many branches. However, in the context of the SW semantics, the problem does not arise except in the case of post-death future contingents. If we restrict our attention to pre-death future contingents, we can say that it is correct to utter 'there will be a sea-battle tomorrow' only for agents in contexts (complete branches) which contain a sea-battle the following day. Pre-fission, there are multiple utterances: some correct, and some incorrect, depending on whether or not they are uttered by agents in sea-battle branches. As long as each agent and each utterance is a part of one complete branch only, the assertion problem does not arise. Of course, as I have just argued, once post-death future contingents are considered, the combination of Literal Fission and the SW semantics breaks down. This objection could perhaps be viewed as a special case of the assertion problem; but that terminology is not particularly germane, since the problem is not restricted to assertion in particular, but affects all cognitive attitudes one might have to post-death future contingents.

Note that even the application to pre-death future contingents of the SW semantics in combination with Literal Fission presupposed that the primary bearers of ordinary mental states are continuant persons, rather than the person-stages to which mental states are ascribed in many popular approaches to persistence. If the primary bearers of mental states are person-stages, then all future contingents would cause problems for Literal Fission, not just post-death future contingents: a person-stage could not worry about its unique future, since it would be a part of many distinct continuant persons. To give the account any chance at all of success in combination with Literal Fission, we must take the bearers of mental states to be continuant persons; this move is made explicitly by Saunders and Wallace [2008]. The question of the primary bearers of mental states is a delicate one; and the requirement that they be continuants rather

than stages is controversial. Fortunately, we will be able to relax it once we move beyond Literal Fission to more sophisticated accounts.

If ascribing mental states to continuant persons rather than to stages is an allowable move, might we not go the whole hog and instead take the primary bearers of mental states to be complete branches? On this picture, branches will typically be ascribed a variety of mental states, sets of which are strongly disconnected from one another; for example, our own branch would be ascribed both a network of 'authorish' mental states and a network of 'readerish' mental states. This radical suggestion would solve the problem of post-death branching; what I generally take to be uncertainty about whether a sea-battle would occur after my death would be reconceptualised as my branch being uncertain-authorishly about whether it includes a sea-battle after some particular time. I hope, though, that the reasons for wanting to avoid a view of this sort are clear. It is not too much of an intuitive stretch to claim that the primary bearers of mental states are continuant humans rather than slices of humans; it is prima facie much more implausible to claim that the primary bearers of mental states are entire universes. In any case, to adopt this view would be to significantly modify the SW semantics, and I will not consider it any further.

Saunders and Wallace did not intend to combine the SW semantics with the Literal Fission picture. As we have seen, it leads to serious problems for the application of their semantics to future contingents. The Literal Fission picture identifies macroscopic objects and events with non-branching aggregates of branch segments; but this fails to ensure that macroscopic objects and entities have unique futures. To underwrite the SW semantics, a metaphysical picture is required according to which macroscopic continuants are uniquely bound to complete branches. And it cannot be simply stipulated that macroscopic objects and events are branch-bound; saying does not make it so. We need some account of the constitution of macroscopic entities from which it follows that they are branch-bound. The next section considers Saunders' and Wallace's preferred way of underwriting the required link between macroscopic entities and branches.

3. Macroscopic Pairing

What metaphysical alternative to Literal Fission could meet the requirement that objects and events be branch-bound? Whatever the picture is, it must involve some relation in which each object or event stands to exactly one complete branch, or the problem discussed above with Literal Fission will recur. Call this requirement the *Branch Binding Criterion*.

The proposal made by Saunders and Wallace [2008] is to identify material objects and events not with continuants, but with ordered pairs of the following form: ¡maximal continuant, continuant¿. I will call this proposal *Macroscopic Pairing*, since its characteristic feature is to identify macroscopic

entities with pairs. It is expressed in rather compressed form in the following passage:

> As goes the incoherence problem of EQM, it is now rather clear, from section 2, of what we are ignorant: we don't know which world - which branch, big-bang to end-of-time - is ours. It is lack of knowledge de se, uncertainty of where we are located, not as a stage S but as a world-stage <W,S> or world-time <W,t>, among the branching worlds.
>
> Saunders and Wallace [2008], p.301

The proposal is, then, that macroscopic objects and events are to be identified with ordered pairs <branch p, continuant q>, where continuant q is a part of branch p in the sense employed by the Literal Fission proposal. I will use the term 'branch-continuant' for such pairs. This proposal has the consequence that you are identified with a pair: specifically, <your branch, your continuant>. Other objects and events in your branch are identified with branch-continuants such as <your branch, continuant r>.

Worlds, which are naturally thought of as maximal material objects or as maximal events, are identified with branch-continuants ¡branch, branch¿. Despite using the term 'world' in this way for convenience, I aim to remain neutral here on the way in which the entities in question relate to the possible worlds of modal metaphysics. Although I think the identification of Everett worlds with possible worlds is in fact a promising approach to modality, that's an issue for another day.

Macroscopic Pairing straightforwardly meets the Branch Binding Criterion. The relation between each branch-continuant and some branch is just inclusion − each branch-continuant includes only one branch. Two branch-continuants will be parts of the same world if they include the same branch. And a branch-continuant x belongs to a world w if and only if x has a first element identical to each element of w. Note that we can now, if we like, drop the assumption (required for the combination of the SW semantics with Literal Fission to get off the ground) that the primary bearers of mental states are complete continuant persons rather than person-stages. Branch-continuants, of whatever length, are always branch-bound and could coherently be ascribed thoughts about their own, unique, branch.

At first glance, the proposal is likely to strike the reader as pretty counter-intuitive. In section 5 I will discuss some objections which question the motivation for the view. Before doing so, however, I want to discuss two issues of detail, which I think provide us with some motivation to modify Macroscopic Pairing.

Consider first the following question: what are the parts of a branch-continuant such as <my branch, aggregate of my branch segments>? Every answer to this question turns out to involve a serious cost.

One approach to mereology for pairs has it that pairs have no proper parts: they are mereological simples. This is not a happy position for a proponent of Macroscopic Pairing to take. Once objects and events are identified with pairs, the result would be that material objects and events have no proper parts at all. This is obviously in serious conflict with our ordinary mereological judgments, according to which my foot is a proper part of my body, and a battle is a proper part of a war.

An alternative approach, associated with David Lewis in *Parts of Classes* (Lewis [1991]), has it that pairs (thought of as sets) have all and only their subsets as proper parts. But this view is no more favourable to Macroscopic Pairing. My only proper parts, on this combination of views, would be the singleton set of the branch that I am on, and the singleton set of my aggregate of branch segments. Neither of these singletons are themselves macroscopic objects, on the Macroscopic Pairing proposal. So once again, macroscopic objects and events have no proper parts that are themselves macroscopic objects or events – my foot is not a part of me, and the battle is not a part of the war. Again, this does serious violence to ordinary thought about parthood.

To interpret ordinary thought and talk about parthood more charitably, we'd have to adopt a highly unorthodox mereology for branch-continuants. If I am identical with the pair <branch B, aggregate A>, then my foot will be identical to a pair containing B and an aggregate which is a proper part of aggregate A. This proposal thus requires a mereology whereby pairs like ¡a,b¿ have as parts not <a> or <b> but pairs like <a,c>, where c is a part of b. Of course, such a mereology would give bizarre results if applied to pairs other than branch-continuants. The pair <Bill, Ben> would have as parts only pairs like <Bill, Ben's arm> and <Bill, Ben's leg>, without having as parts pairs like <Bill's arm, Ben> or <Bill's leg, Ben>. This kind of violation of mereological symmetry seems extremely problematic. In addition, none of the following would be part of <Bill, Ben>: Bill, Ben, Bill's singleton, Ben's singleton. By itself, this consequence may seem pretty counterintuitive.

To avoid absurd results of this sort, we would need to restrict the modified mereology to apply only to branch-continuants, and say that when talking about parts of branch-continuants, we are not strictly speaking applying the parthood relation. Rather, we are applying some other relation – parthood* – which satisfies the principles outlined in the previous paragraph. But this is an

unwelcome conclusion. When we use the word 'part' in ordinary language, it seems that we are applying the relation of parthood itself, not some other relation defined in terms of parthood. Parthood is frequently taken to be a universal relation, applying in exactly the same sense to very different domains – to sets, to material objects, to events, to works of art, and so on. The universality is sometimes accounted for by saying that mereology is 'quasi-logical'; sometimes by saying that parthood is a highly natural relation; and sometimes by saying that ascriptions of parthood are not context-dependent. I think something of this sort is probably correct; but Macroscopic Pairing sits uncomfortably with this line of thought.

Macroscopic Pairing will also lead to complications involving the structure of properties had by macroscopic objects and events. Most of the ordinary everyday properties that we attribute to macroscopic objects will depend only on the second element of a branch-continuant pair. This introduces a need for an additional layer of properties which are possessed by branch-continuants in virtue of certain other properties possessed by the elements of the branch-continuants. Such an additional layer of properties comprises an ugly complication in our account of qualitative variation amongst material objects and events.

An example helps to illustrate this. Consider the lifespan of some particular mouse. The property of lifespan had by the branch-continuant mouse is had in virtue of some other property (call it lifespan*) which is possessed by the second element of the pair, the mousy continuant. Similarly for properties had by branch-continuants at times – a branch-continuant has a mass at a time in virtue of its second element having mass* at a time. Thus Macroscopic Pairing requires a significant extra commitment to properties: where we might have thought there was a single property, Macroscopic Pairing requires us to admit two properties, one of which is had by a pair in virtue of the other being had by the second element of the pair. I will call properties such as mass* and lifespan* 'dogleg properties'[8]. The properties typically ascribed to material objects in our ordinary discourse, according to Macroscopic Pairing, are dogleg properties.

Like the consequences of Macroscopic Pairing for mereology, its consequences for macroscopic property-ascription involve added ideological

---

[8] Thanks to an anonymous referee for this terminology. It is worth noting that certain other metaphysical views also require such additional layers of properties. For example, a Lewis-style four-dimensionalism identifies people with space-time worms while taking person-stages to be the fundamental property-bearers; so for Lewis, persons have properties like lifespan in virtue of properties had only by their stages. See the postscript to Lewis [1976].

complexity and are *prima facie* costs of the view. This need for modifications of the structure of macroscopic properties and of macroscopic mereology may not force us to give up Macroscopic Pairing. But I do think that they give us reason to examine other frameworks for implementing the SW semantics. In the next section, I outline one promising alternative.

4. Space-time Point Pairing

The non-standard mereology and property structure required for macroscopic objects and events are *prima facie* costs of the Macroscopic Pairing proposal. But there is an alternative account available which avoids these complications. It also goes some way towards reducing the intuitive weirdness of Macroscopic Pairing. Instead of identifying all material objects and events with pairs, we could limit such identifications to space-time points. A space-time point would be a pair <b,a>, where a is a pointlike part of a branch[9], and b is a complete branch. Call this proposal *Space-time Point Pairing*.

Once we adopt Space-time Point Pairing, familiar macroscopic ontology can be recovered from the picture in a number of ways. My preference is for a supersubstantivalist view which identifies material objects and events directly with spatio-temporal region, and then identifies spatio-temporal regions with fusions of space-time points. Fusions of ordered pairs might seem rather strange entities — but all that we need say about them for present purposes is that the principles of classical mereology apply to them unrestrictedly. For ease of exposition, I will restrict myself to discussing this version of the view, though Space-time Point Pairing is in principle compatible with other accounts of the relation between space-time, material objects, and events, for example with various accounts which posit some kind of fundamental occupation relation.

The combination of Space-time Point Pairing with the SW semantics avoids the problems with mereology which dogged the previous proposal. The mereology of space-time regions and of macroscopic objects and events works exactly as it ordinarily does: the parthood relation which relates my foot to me and the battle to the war is characterized by the usual principles of classical mereology. The Branch Binding Criterion is satisfied; objects and events are identical to fusions of space-time points, and each of the space-time points which is a part of such a fusion is branch-bound in virtue of including a branch as an element. Hence macroscopic objects and events (and all of their extended parts) are indirectly bound to particular branches, in virtue of all of their pointlike

---

[9] What does it mean to say that a is a pointlike part of a branch? Writers on Everettian QM, such as Saunders and Wallace [2008] and Tappenden [2008], freely make use of the 'temporal part' terminology, and apply it to parts of branches. As already noted, this is an idealization; the current proposal merely extends this idealization to pointlike parts.

parts being directly bound to particular branches. And as with Macroscopic Pairing, there is no need for the troublesome assumption that complete continuants (rather than stages) are the primary bearers of mental states.

Have the problems with the mereology of macroscopic objects and events on the Macroscopic Pairing proposal merely been shifted to problems with the mereology of the space-time points themselves? In a sense, yes; but in the context of space-time points the conflict with ordinary usage is far less of a problem. If we take it that pairs are mereological simples, then space-time points will have no parts, just as is usually supposed: this is my preferred option. If we take it (following Lewis) that pairs are sets, and that sets have their subsets as parts, then each space-time point has as parts the singleton set containing an entire branch, and the singleton set containing a pointlike part of that branch. This is perhaps a little counterintuitive, but is not in serious tension with ordinary usage; our ordinary thought and talk about parthood does not adjudicate clearly either way on the question of whether space-time points have (non-spatio-temporal) parts, whereas it does adjudicate clearly on the question of whether macroscopic objects and events have macroscopic parts.

If we grant the plausible assumption that the properties of macroscopic objects and events supervene on the properties of their pointlike parts (without limiting the supervenience base to local or to intrinsic properties of pointlike parts), Space-time Point Pairing also removes the need for modification of the metaphysics of macroscopic properties. Pointlike parts of macroscopic objects and events will still have dogleg properties, which are had in virtue of properties had by their second elements; however, ascribing properties to macroscopic objects and events themselves will not need an underlying layer of macroscopic dogleg properties. Property-ascription for macroscopic objects works in the Space-time Point Pairing picture just as it does in a standard one-world theory. We need only be revisionary about the metaphysics of space-time points: we can be as conservative as we like about the metaphysics of macroscopic objects and events which are constructed out of or located at sets of those points.

An interesting consequence of this proposal is that the 'space-time' of the quantum mechanics and quantum field theory formalism, in terms of which branches are defined, is not the same as the 'space-times' of macroscopic worlds. The former 'space-time' is a single entity common to multiple branches, while each of the latter 'space-times' is tied to a particular macroscopic course of events. It seems likely that charity arguments in the style of Wallace [2005] will tend to lead us to interpret our ordinary term 'space-time' as referring to the latter kind. And this conception of space-times as linked to branches could have independent motivation; having a distinct space-time for each branch would help to make conceptual room for the matter-distribution-dependence of the metrical structure of space-time which appears to be needed for quantum gravity. So I'd

suggest that the Space-time Point Pairer should reserve the term 'space-time' for the emergent entity linked to individual branches, and something like 'pre-space-time' for the fundamental entity common to branches.

Unlike Macroscopic Pairing, Space-time Point Pairing does not require a modified mereology for macroscopic objects and events, but can incorporate familiar systems of mereology unchanged. Nor does it require a network of macroscopic dogleg properties underlying all ordinary attributions of properties to macroscopic objects and events. It therefore spares us two fragments of rather arcane and unlovely metaphysics. And in common with Macroscopic Pairing, it has the advantage over Literal Fission that it can give a non-revisionary account of our ignorance of post-death future contingents, and hence can provide for a fully general account of objective probability.

5. Objections to the pairing proposals

In this section I will discuss some objections which arise for both the Macroscopic Pairing and Space-time Point Pairing proposals. But first I wish to record my suspicion that the most common response to the proposals will be that they are simply absurd, no matter whether or not they recover what we take to be the correct truth-values for our utterances. This sort of response was referred to by Lewis [1986, section 2.8] as 'the incredulous stare'; like Lewis, I will not attempt to refute it.

1.1.1   Objection 1 – Phenomenology

It might be objected that we don't *feel* like pairs, or like fusions of pairs; that our actual phenomenology could not be available for a pair. Similarly, we might worry that we have *a priori* knowledge (or certainty) that we ourselves exist, but not that pairs exist or that fusions of pairs exist.

A quick response to this sort of objection goes as follows: how would you know what it would feel like to be a pair or a fusion of pairs? Without knowing that, how could you know that it would match our actual phenomenology? Perhaps, then, our actual phenomenology reflects just how it in fact feels to be a branch-continuant; a defender of the pairing proposals could likewise say that our knowledge/certainty that we ourselves exist just is the knowledge/certainty that certain pairs exist, since that's what we've turned out to be.

The burden of proof here is with the objector. What is needed to make the objection from phenomenology stick is an argument that it couldn't feel like anything at all to be a branch-continuant or a fusion of point pairs, or at least that it couldn't feel anything remotely like *this*; and absent a positive argument of this sort, the identification of macroscopic objects and events with branch-

continuants or with fusions of point pairs appears legitimate. Compare the situation with materialism in the philosophy of mind; if materialism is right, human beings are just identical with their bodies, and bodies are lumps of meat. The complaint 'I don't feel like a lump of meat' isn't likely to cut much ice in the materialism debate; what is wanted is a convincing argument that it couldn't feel like anything to be a lump of meat. Similarly, we need to distinguish the worry that it seems surprising or counterintuitive that we are pairs or fusions of pairs from the worry that we know what it would be like to be a pair, or a fusion of pairs, and that that we know our experience is not of this sort.

1.1.2   Objection 2 – Branch-continuants are unnatural

Another possible line of objection is that branch-continuants are in some sense gerrymandered, and hence unnatural candidates for the referents of our terms. In the Lewisian picture, naturalness is fundamentally a property of properties; so the objection put in this setting is that *being a branch-continuant* is less natural than *being a continuant*. Therefore, so the objection goes, we should reject a semantics according to which branch-continuants are taken to be the referents of our ordinary object-terms and event-terms.

The response that a friend of Macroscopic Pairing should make to this objection is that candidacy for reference is determined through a charity-style procedure, in which naturalness is only one factor. Naturalness must be balanced by the need to interpret as much as possible of our discourse in such a way that it comes out true (or as knowledge, if we prefer the knowledge-maximizing principle of charity defended by Williamson [2007]). The arguments of section 2 were intended to establish that Literal Fission falsified too much of our common-sense thought about the future to meet any reasonable requirement of charity; and it seems plausible that the lower degree of naturalness of the property *being a branch-continuant* compared to *being a continuant* is not enough to counterbalance this uncharitability.

In sum, if Everettian QM is correct, then there are no *better* candidates as the referents of our ordinary terms for material objects and events than branch-continuants or fusions of point pairs; other candidates would falsify too much of what we take to be true (or to be known).

1.1.3   Objection 3 – Pairs are abstract; we are concrete

According to the Macroscopic Pairing proposal, the concrete objects are a species of the abstract objects: the abstract pair ¡my branch, aggregate of my branch segments¿ is identified with a concrete object, viz., me. It follows that some abstracta are concreta, and that, for example, some abstracta are causally active and have spatio-temporal location. And the Space-time Point Pairing proposal as I have developed it has it that concreta are fusions of abstracta.

These views accordingly involve some revisionism about the abstract-concrete distinction. But is this a fatal difficulty? Our grasp of the abstract-concrete distinction is nothing like as clear as we would like it to be. There are already various well-known borderline cases: examples might include the global economy, facts, or heat. This seems like a case in which our metaphysical intuitions might be negotiable, if giving them up results in gains elsewhere.

The objection from the abstract/concrete distinction is particularly unconvincing when applied to Space-time Point Pairing[10]. Space-time points are well-known to be a case where the abstract/concrete distinction is at its most problematic; our judgements about abstractness and concreteness are on particularly weak ground when it comes to points (hence the debate about whether Hartry Field's [1980] nominalism really deserves the name). If any entities are good candidates for occupying a middle-ground between abstractness and concreteness, space-time points are.

Indeed, there is a familiar precedent for the identification of objects and events with fusions of abstracta; Quine [1965] proposed a sparse fundamental ontology which reduced objects and events to space-time regions, which reduced space-time regions to sets of space-time points, and which reduced space-time points to ordered quadruples of real numbers. The present proposals are less radical than that: the pairs in question have as elements parts of branches of the quantum state, entities which have by many past Everettians been identified directly with macroscopic objects and events. In contrast, Quine's proposal reduces objects to fusions of sets of quadruples of numbers; nothing at all recognisably 'concrete' enters in at the metaphysical ground floor.

It should also be borne in mind that adopting Space-time Point Pairing does not commit us to the supersubstantivalist move discussed in section 4. If we choose not to make this move, then we don't even have to identify material objects and events with fusions of abstracta. All we have to say is that material objects and events are located at fusions of abstracta, which may be less controversial still.

1.1.4   Objection 4 – Pairs are arbitrary

Mathematical logic gives us more than one way of defining an ordered pair[11]. Which of these entities are you and I identical to? The objection runs

---

[10] This fact might provide a further reason for preferring Space-time Point Pairing to Macroscopic Pairing, over and above the arguments of section 4.

[11] http://en.wikipedia.org/wiki/Ordered_pair currently lists nine, eight of which are set-theoretic.

exactly parallel to Benacerraf's well-known objection to the combination of Platonism with set-theoretic constructions of numbers[12].

It should be noted that the pairs appealed to by the pairing proposals don't need to be ordered pairs. One element of any branch-continuant pair will always be a (perhaps improper) part of the other. So we have the option of replacing ordered pairs with non-ordered pairs, and replacing talk of the first element or the second element with talk of the smaller element and the larger element. In the case of worlds, where the two elements are the same size, it makes no difference which element we take to be the continuant and which we take to be the branch.

If we follow this route, we avoid the problem of choosing one definition of an ordered pair over another. It also has the advantage that we don't need to make an arbitrary choice about whether our ordered pairs are <continuant, branch> pairs or <branch, continuant> pairs. Just the unordered set with the two elements will do. But even if there is some simplest natural candidate (such as the unordered set of the two objects, perhaps) this wouldn't make the problem go away altogether. There would still be many distinct set-theoretic constructions with all the structure we require. And it's difficult to see what sort of evidence could rule out the claim that we are, for example, Kuratowski ordered pairs instead of unordered sets.

To this concern, we can respond as many have wanted to respond to Benacerraf's problem, and say that there is no fact of the matter about which of the various set-theoretic constructions are really identical to macroscopic objects and events. This indeterminacy can then be cashed out in terms of your favourite bivalence-preserving theory of vagueness[13]; appealing options might include epistemicism (Williamson [1994]), contextualism (Fara [2000]), semantic indeterminacy (in the sense of Dorr [2003]), and primitive metaphysical indeterminacy (Barnes and Williams [2010]). Indeterminacy of this general sort in the reference of our terms is in fact already fairly widely accepted; it has often been embraced in order to solve the Problem of the Many.

---

[12] See Benacerraf [1965].

[13] Supervaluationists characteristically distinguish between preserving the Law of Excluded Middle and preserving bivalence, and choose to preserve the former but not the latter. Since the SW semantic proposal is in part motivated by the desire to preserve bivalence, the appropriate theory of indeterminacy for our purposes ought to be bivalence-preserving.

6. Overview of the pairing proposals

At this stage it is worth clarifying the structure of the ontology behind the Space-time Point Pairing and Macroscopic Pairing versions of EQM. The pictures are complicated, and it is easy to lose track of the commitments that they really involve.

At the fundamental level, the ontology is monistic - there is just one single highly-structured object, the universal quantum state. Call the language which describes reality in these terms the Universal State language. By use of decoherence theory to pick out privileged structure from the universal state, we can construct a language (call it the Pluriverse language) in which we quantify over structural features of the universal state: branches, branch segments, and so on.

The Pluriverse language can be thought of as the working language of metaphysical theorizing for Everettians. It is flexible enough to enable us to refer to everything we want to speak about, whether inside or outside the metaphysics room. But it is not a charitable account of our ordinary thought and talk about the macroscopic world, as the problems with Literal Fission show. We therefore need a third language - call it the Universe language − in order to account for our uncertainty of future contingents. The pairing proposals comprise alternative accounts of the relationship between the Pluriverse language and the Universe language[14]. Both involve the idea that our Universe language quantifiers are generally restricted to range over a very special kind of entities − ordered pairs, or fusions of ordered pairs.

According to the point pairing proposals, no macroscopic world overlaps any other macroscopic world. What is left of the intuitive thought that EQM is a branching theory? Both Space-time Point Pairing and Macroscopic Pairing involve no branching at the level of the Universe language; however, they do still incorporate branching at an underlying level, the level of the Pluriverse Language. The structure produced by decoherence, which provides the 'raw material' for the elements of the pairs identified with space-time points, *is* taken to be a branching structure; different branches genuinely do have parts in common.

The pairing proposals essentially specify procedures for constructing non-overlapping worlds out of a set of overlapping branches, which raises the question of whether the branching structure might be in principle dispensable. An alternative approach to EQM would interpret consistent histories *directly* as

---

[14] The terms 'Pluriverse Language' and 'Universe Language' are intended to recall the two languages employed by Lewis [1968] to help make sense of modal operators in ordinary language.

representing non-overlapping worlds. No branching structure then need enter into the picture at any stage, and the SW semantics could be applied without appeal to Macroscopic Pairing or Space-time Point Pairing. I suspect that this may in the end be the most natural way to understand EQM, and I will discuss it in more detail in the next section. But the discussion so far has aimed only to clarify various metaphysical pictures which do take it that decoherence provides us with a branching structure: Literal Fission, Macroscopic Pairing, and Space-time Point Pairing. All three of these views give us the resources to explain probability and uncertainty about events in our own future, while preserving a bivalent semantics. Like Macroscopic Pairing, Space-time Point Pairing does better than Literal Fission by making sense of post-death uncertainty. Unlike Macroscopic Pairing, Space-time Point Pairing does not achieve this at the cost of an implausible mereology and property structure, and it deals better with the objection from the abstract/concrete distinction. I therefore recommend Space-time Point Pairing to those Everettians who think that decoherence gives us a structure of overlapping branches.

7. Everettian Quantum Mechanics without branching

Macroscopic Pairing and Space-time Point Pairing both involve branching at the level of the Pluriverse language, while involving no branching at the level of the Universe language. In the previous section I argued that Space-time Point Pairing is the best way of constructing non-overlapping macroscopic worlds, objects and events from the materials provided by a structure of overlapping branches. But this construction itself gives some cause for concern. The distinction between branches and worlds is an unwelcome one, since these terms have often been used interchangeably; and the pairing proposals commit us to a number of odd-sounding claims which, while perhaps not directly contradicting our everyday beliefs, certainly do not appear among them. For example, both views commit us to the claim that space-time points have entire branches as elements. In the light of this, there is some motivation to consider an alternative framework, which dispenses altogether with overlapping branches and removes the need to identify concrete objects and events with ordered pairs. In this final section, I will consider such an approach to EQM, which interprets consistent histories as diverging worlds, in the sense of Lewis [1986, p.206]. Diverging worlds are exactly alike up to the time of divergence, but - unlike branching worlds - they have no parts in common.

Nothing in the formalism of quantum mechanics compels us to think of consistent histories as branching rather than diverging. In the consistent histories formalism, a branch is a consistent history; a string of quantum-mechanical projection operators obeying certain decoherence conditions. The point at issue between the diverging and branching conceptions is whether the entities represented by some particular projection operator appearing in two different consistent histories are numerically identical, or whether they are

(numerically distinct) qualitative duplicates. Numerically identical entities give us branching worlds; qualitative duplicates give us diverging worlds.

A diverging interpretation of the consistent histories formalism has a great deal in common with Lewisian modal realism[15] (with the exception, of course, that the laws of quantum mechanics hold in every Everettian world). Both views posit a set of complete histories, many of which have initial segments which are qualitatively identical with one another, but none of which have any parts in common; both views locate each history in a distinct space-time; and both views interpret much of our ordinary uncertainty of contingent fact as self-locating uncertainty within the set of histories.

Although Saunders and Wallace [2008] tacitly assumed an underlying branching structure, Saunders no longer considers this to be warranted by the quantum formalism; Saunders [2010] argues that from a mathematical standpoint one can well defend the picture of Everettian histories as diverging, consistent with the Saunders-Wallace semantics. I agree with him that the mathematical structure of quantum mechanics does not itself decide between branching and divergence. Nor, as far as I can see, does any element of physicists' own practice in talking about worlds in the Everett interpretation. The question of whether the entities represented by the projection operators concerned are numerically identical or just qualitatively identical has a distinctively metaphysical flavour; so we might do well to look to metaphysics when answering it.

Lewis [1986, p.207-8] famously argued that a branching version of modal realism presents problems for making sense of our thought and talk about the future. If there are multiple futures branching from this very moment, then it seems that an expression like 'the future' must suffer from reference failure; the definite description presupposes uniqueness, but this presupposition is false. It should be clear that Lewis' considerations apply to EQM in exactly the same way as they apply to modal realism; the argument against Literal Fission in section 2 is a close relative of the Lewisian argument against a branching version of modal realism. It is therefore surprising that while this argument has been highly influential in the metaphysics of alethic modality, it has until recently had little impact on the interpretation of EQM.

A diverging interpretation of the consistent histories formalism permits a straightforward bivalent semantics, without involving us in the rather arcane metaphysics of Macroscopic Point Pairing or Space-time Point Pairing. I

---

[15] Lewis [1986].

conclude that Everettians would do well to abandon the problematic metaphor of branching in favour of the more germane metaphor of parallel worlds[16].

References

Barnes, E. & Williams, R. [2010]: 'A theory of metaphysical indeterminacy', forthcoming in *Oxford Studies in Metaphysics vol. 6* (Oxford University Press).

Belnap et al. [2001]: *Facing the Future: Agents and Choices in Our Indeterministic World* (Oxford University Press).

Benacerraf, P. [1965]: 'What numbers could not be', *Philosophical Review* 74, pp. 47-73.

Dorr, C. [2003]: 'Vagueness without ignorance', *Philosophical Perspectives* 17:1, pp. 83–113.

Field, H. [1980]: *Science Without Numbers* (Princeton University Press).

Fara, Delia Graff [2000]: 'Shifting Sands: An Interest-Relative Theory of Vagueness', P*hilosophical Topics* 28, pp. 45-81.

Griffiths, R. [2001]: *Consistent Quantum Theory* (Cambridge University Press).

Lewis, D. [1968]: 'Counterpart theory and quantified modal logic', *Journal of Philosophy* 65(5), pp. 113-126.

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

Lewis, D. [1979]: 'Attitudes de dicto and de se', *Philosophical Review* 88:4, pp. 513-543.

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

Lewis, D. [1991]: *Parts of Classes* (Oxford: Blackwell).

Quine, W.V.O. [1965]: 'Propositional Objects', reprinted in Quine, W.V.O. [1969]: *Ontological Relativity and Other Essays* (Columbia University Press).

Saunders, S. [1993]: 'Decoherence, Relative States, and Evolutionary Adaptation', *Foundations of Physics* 23, pp. 1553-1585.

Saunders, S. [1995]: 'Time, Quantum Mechanics, and Decoherence', *Synthese*, 102, pp. 235-66.

Saunders. S. [1998]: 'Time, Quantum Mechanics, and Probability', *Synthese*, 114, pp. 405-44.

Saunders, S. [2010]: 'Chance in the Everett Interpretation', in *Many Worlds? Realism, Everett, and Quantum Mechanics*, eds. S. Saunders, J. Barrett, A. Kent, D. Wallace (Oxford: OUP).

Saunders, S. & Wallace, D. [2008]: 'Branching & Uncertainty', *British Journal for the Philosophy of Science* 59:3, pp. 293-305.

Tappenden, P. [2008]: 'Identity and uncertainty in Everett's multiverse', http://philsci-archive.pitt.edu/archive/00003896/ .

Wallace, D. [2003] 'Everettian rationality: defending Deutsch's approach to probability in the Everett interpretation', *Studies in History and Philosophy of Modern Physics* 34:3, pp. 415-439.

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

Wallace, D. [2010]: 'How to Prove the Born Rule', in *Many Worlds? Realism, Everett, and Quantum Mechanics*, eds. S. Saunders, J. Barrett, A. Kent, D. Wallace (Oxford: OUP).

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

Williamson, T. [2007]: *The Philosophy of Philosophy* (London: Routledge).