The Relativity of Branching

Guido Bacciagaluppi*

15 April 2025

Abstract

I argue that different ways that branching fits within Minkowski spacetime are merely different descriptions of an invariant notion of branching and are due to the relativity of simultaneity. The argument fits in the wider framework of Everett branches as real patterns, and is both developed in the abstract setting of the (generalised) histories formalism, and discussed comparing the concrete examples of hypersurface-dependent branching and of branching along the forward lightcone. I formulate the latter in terms of branching spacetime, suggesting this is a way in which spacetime can emerge from the universal wavefunction, and I make tentative connections with causal set theory. The proposed view is compatible with both the Schrödinger and Heisenberg picture. [To appear in A. Ney (ed.), *Locality and the Many Worlds Interpretation of Quantum Mechanics* (Oxford: Oxford University Press).]

1 Introduction

A central question in the discussion of locality in Everett is how branching fits with Minkowski spacetime: whether it is global, local, instantaneous, along lightcones or what not. In this chapter I shall argue that apparent differences in these views are

---

*Descartes Centre for the History and Philosophy of the Sciences and the Humanities and Freudenthal Institute, Utrecht University; IHPST and SPHERE, Paris; Foundational Questions Institute, `https://fqxi.org/` (email: g.bacciagaluppi@uu.nl).

not substantial but arise merely through reifying frame- or better foliation-dependent descriptions of a fundamentally *invariant* notion of branching. As my title is intended to suggest, my claim is modelled on that of the relativity of simultaneity, but I shall be arguing for it in the more general context of thinking of worlds as *real patterns*.

Formally, and even restricting oneself (as I shall do) to Everettian worlds as described in some version of the histories formalism, there are a number of choices for making precise the notion of a world (as explained more extensively in the next section), and one can ask oneself which ones are substantial and which ones might be merely a matter of descriptive convenience. For instance:

(A) One can choose different stability conditions for histories: requiring lack of interference is captured by the condition of consistency, requiring the existence of generalised records is captured by the condition of decoherence, and requiring the existence of actual records implies that decoherence is induced by interaction with the environment.

(B) If one opts for environmental decoherence, the choice of exact times and projections in the corresponding finest-grained histories is merely constrained up to certain decoherence scales.

(C) Further, one may choose to identify worlds not with fine-grained histories but only with certain coarse-grainings.

(D) Finally, one needs to make choices as to the identity of the projections appearing in the histories that describe worlds: in particular, whether worlds are thought to *split* or *diverge* depends on whether one identifies the same projections across worlds or treats them as counterparts of each other.

Some cases are arguably clear-cut. In case (A), different choices lead to different notions of a world, and in the following I shall presuppose environmental decoherence (I shall not argue in detail for it, but consistency is mathematically ill-behaved and generalised records are presumably not robust enough). In case (B), a world as defined through environmental decoherence is a vague concept, and each choice of exact fine-grained histories is merely a choice of precisification describing the same vague world.[1]

---

[1] For more details see the next section. Vagueness of worlds has been discussed by Wallace (2012, Part I).

The other cases are not so obvious. Take different levels of coarse-graining (C). According to Deutsch, quantum computers perform parallel computations in different worlds (even though these worlds are short-lived because they are recombined at the read-out stage); but one can alternatively think of them as being in superpositions within a single world. Or take Vaidman's well-known example of a neutron following two different paths in an interferometer: does the neutron split and then recombine, or are we performing experiments in a single world on a neutron in a superposition? And in an EPR experiment, does Bob's electron split when *Alice* measures hers, or only when Bob measures his own electron? Take instead cases involving identity of projections (D). Do *we* split if the neutron in the interferometer does? Does *Bob* split when Alice performs her measurement, even if he has not (yet) performed a measurement on his electron? And (of course) do *worlds* split or do they diverge?[2]

To use a phrase by Simon Saunders, are we *choosing to talk about different things*? Or are we in fact choosing to use different (perhaps suggestive) descriptions of the *same* thing? To answer such questions we need to invoke metaphysics. For instance, our views about personal identity may impose a choice between splitting and divergence. In the following, I shall presuppose one specific metaphysical view about worlds or branches, namely that they are *real patterns* in the sense of Dennett and of Wallace.[3] That is, they are *patterns* insofar as they are 'in the eye of the beholder', but they are *real* because there is something that makes them objectively useful for prediction (and explanation) – where 'objectively useful' as used here will mean that the source of a pattern's usefulness lies in the physics itself.

My claim is that, from this point of view, while the choices under (C) are choices between different equally real patterns, the choices under (D) as well as different choices of relativistic branching structures involve no 'real differences' between patterns, and are merely different more or less convenient descriptions of the same real patterns.

This chapter is structured as follows. In Section 2 I discuss the non-relativistic

---

[2] Deutsch's view is well-known, but see e.g. Deutsch (1985). Vaidman's neutron appears in his classic Vaidman (1998). For views about local versus global branching, I refer to Ney (this volume) and references therein. For splitting and divergence as used here see Wilson (2020, Chap. 2, and this volume); see also Tappenden (2008) (and the references in these works).

[3] See Wallace (2012, Part I). Note that the term 'worlds' often carries strong global connotations ('worlds' should in some appropriate sense contain everything there is). For this reason I shall often use the more neutral term 'branches', which still suggests multiplicity but might be applicable also to subsystems or subregions of the universe.

case, sketching the histories formalism and applying real patterns to assess the examples under (C) and (D) above. In Section 3, I discuss the histories formalism in Minkowski spacetime: I present one manifestly invariant and one foliation-dependent generalisation of decoherence and of branching, and argue that despite *prima facie* differences in branching structure these generalisations are physically equivalent (thus the corresponding branching structures exhibit no 'real differences'). In Section 4, I make the argument more concrete, by describing what I take to be the two most natural approaches to relativistic branching in terms of the universal wavefunction. The first is *hypersurface-dependent branching*, as proposed (even though not in print) by Wayne Myrvold – which makes explicit the relativity of branching. The second is *branching along the forward lightcone*, as proposed among others by myself – which makes explicit the invariant nature of branching. In Section 5, I further elaborate the idea that it allows one to formulate the Everett theory in terms of *branching spacetime*. Finally, in Section 6, I conclude with a few comments on the possible ontologies suggested by this picture of branching.[4]

2 Non-relativistic case

The general strategy one pursues in the Everett theory is: (i) to find stable structures at the level of components within the universal wavefunction, and (ii) to identify these as multiple branches or worlds. Back in 1957, the stable structures considered by Everett were at the level of patterns of relative states.[5] Nowadays, especially since the work by Saunders (1993) and Wallace (2012), they are mostly described using

---

[4] I would like to thank first and foremost Alyssa Ney for offering me the opportunity to revisit an issue that has engaged me for many years, the other participants at her workshop for excellent discussions that have informed the writing-up of this chapter – especially Jacob Barandes, Charles Bédard, Sam Kuypers, Simon Saunders and Alastair Wilson – as well as Wayne Myrvold, Alyssa Ney, and Paul Tappenden for comments on the penultimate version of this chapter. I also have important historical debts: Simon Saunders first and then David Wallace convinced me that Everett provided a credible option for understanding quantum mechanics, while the newer ideas in this chapter have been largely inspired by Wayne Myrvold. In particular, the material of Section 4 was mainly worked out at the 2016 International Summer School in Philosophy of Physics at the University of Urbino, where both Wayne and I should have taught and where I did my best to fill the gap when Wayne had to cancel.

[5] Everett's collected works have been edited by Jeff Barrett and Peter Byrne (Everett 2012). For a reading of the historical Everett in terms of stable structures, see my review of that volume (Bacciagaluppi 2013).

the histories formalism, which I shall therefore very briefly sketch.[6]

A *history* is a (usually) finite time-ordered sequence of projections (often conveniently written in Heisenberg picture):

$$P_{i_1}(t_1), P_{i_2}(t_2), \ldots, P_{i_n}(t_n) \ . \tag{1}$$

We define the associated *history operator* as:

$$C_\alpha := P_{i_n}(t_n) \ldots P_{i_1}(t_1) \tag{2}$$

(where α stands for the multi-index $i_1 i_2 \ldots i_n$).

If for all times t_j the projections $P_{i_j}(t_j)$ are mutually orthogonal and sum to the identity, we have a *history space*. We obtain a *coarse-graining* of a history space by taking sums of (some of) the projections $P_{i_j}(t_j)$ at (some of) the times t_j (with *fine-graining* defined correspondingly).

For any two histories in a history space, and given a state $|\Psi\rangle$ (where we can also take a mixed state ρ), we define the *decoherence functional* as

$$\mathcal{D}(C_\alpha, C_{\alpha'}) := \mathrm{Tr}\Big(C_\alpha |\Psi\rangle\langle\Psi| C^*_{\alpha'}\Big) \ , \tag{3}$$

and the *weight* of a history as the positive number

$$\mathcal{D}(C_\alpha, C_\alpha) = \mathrm{Tr}\Big(C_\alpha |\Psi\rangle\langle\Psi| C^*_{\alpha}\Big) \ . \tag{4}$$

The right-hand side of (4) of course has the form of the (generalised) Born rule for a *fixed sequence* of measurements. Coarse-graining (or fine-graining) the history space corresponds to formally considering the Born rule for *different* sequences of measurements, so that probabilities interfere. Although the weights (4) obey the Kolmogorov axioms, in general the weights of disjunctions of histories will not be equal to the weights of the corresponding coarse-grained histories.

Requiring that they nevertheless be equal is equivalent to imposing that, for any two distinct histories in the history space, the real part of the decoherence functional should vanish:

$$\mathrm{Re}\mathcal{D}(C_\alpha, C_{\alpha'}) = 0 \ . \tag{5}$$

---

[6] I shall be talking about the universal wavefunction, but this should not be taken as commitment to a wavefunction ontology as fundamental (see also Sections 4 and 6 below).

This is known as the *consistency* condition (or weak decoherence). It corresponds to the case in which there happens not to be any interference even though no measurements need have been performed. In this case, notions such as 'joint weights' or 'transition weights' (between successive projections with non-zero joint weight) can be defined equivalently using disjunctions or coarse-grainings of histories. Consistency can be generalised to the case of mixed states (trivially), as well as to infinite histories (if all finite subhistories satisfy it). Consistency is also the minimal *stability condition* that has been proposed for history spaces in order for histories to qualify as Everettian worlds. It appears to be too weak, however, already on mathematical grounds. In particular, it generally fails to be preserved under composition of systems – for the simple reason that the real part of a product of two complex numbers is generally not the product of their real parts (Diósi 2004).

Another well-known and strictly stronger stability condition is the *decoherence* condition (or medium decoherence):

$$\mathcal{D}(C_\alpha, C_{\alpha'}) = 0 \tag{6}$$

for any two distinct histories. It can again be extended to mixed states and to infinite histories as above, and for the case of pure states it is equivalent to the existence of 'permanent records', meaning that for all times t_j and all history operators $C_{\alpha_j} = Pi_j(t_j) \dots Pi_1(t_1)$, there are mutually orthogonal and exhaustive projections $R_{\alpha_j}(t_j)$ such that

$$R_{\alpha_j}(t_j)|\Psi\rangle = C_{\alpha_j}|\Psi\rangle \ . \tag{7}$$

These projections can be used to fine-grain the original history space preserving decoherence, and it is easy to check that for all times t_j each $R_{\alpha_j}(t_j)$ is perfectly correlated with the corresponding previous history. In this sense, at all t_j there are 'records' of the projections at earlier times. These might just have the form

$$R_{\alpha_j}(t_j) = C_{\alpha_j}|\Psi\rangle\langle\Psi|C^*_{\alpha_j} \ , \tag{8}$$

in which case one talks of 'generalised records'. For instance take a system where energy is conserved, such as an isolated hydrogen atom: the projections onto the eigenstates of energy at any time are (trivially) records also of previous values of the energy ('conservation-based decoherence'). To my mind, this shows that also decoherence may be too weak as a stability condition, since such generalised records might be destroyed by small perturbations (quite apart from the restriction to pure states in order for records to exist).

The stability condition that I shall adopt in the following is the yet stronger condition that *actual records* of the projections in the histories should form in the environment. This can be thought schematically as the system interacting at each time t_k with some degree of freedom in the environment such that (in Schrödinger picture)

$$|\psi_i\rangle \otimes |e_0^k\rangle \to |\psi_i\rangle \otimes |e_i^k\rangle \ , \tag{9}$$

where each $|\psi_i\rangle$ is an eigenstate of a projection in a history and each such environmental degree of freedom then evolves separately to $|e_i^k(t_j)\rangle$. Each projection

$$E^1_{i_1}(t_j) \otimes E^2_{i_2}(t_j) \otimes \ldots \otimes E^j_{i_j}(t_j) \tag{10}$$

(with $E_i^k(t_j) := |e_i^k(t_j)\rangle\langle e_i^k(t_j)|$) is thus a permanent record in the sense of (7). As an example, take a hydrogen atom in interaction with the electromagnetic field: the field will 'spontaneously measure' the energy of the atom, resulting in spectroscopic records in the environment. Unlike the previous case, records exist even though transition probabilities are non-trivial (we have quantum jumps).[7]

We can now define branching. Take all pairs of successive projections with non-zero transition weights: the history space is *branching* if projections have unique predecessors but possibly different successors. Where branching is non-trivial, I shall talk of a *branching event.*

Finally, the *branching-decoherence* theorem (Griffiths 1993; Wallace 2012, App. A) states that a branching history space is automatically decoherent and that if one fine-grains a decoherent history space by including the permanent records of the projections in the histories, the resulting history space is branching.[8]

As mentioned, describing Everett worlds or branches in the histories formalism involves a number of choices: (A) of stability conditions, (B) of fine-grained projections, (C) of coarse-grainings, and (D) of the identity of the projections in the histories. If we understand branches as real patterns, I suggest that these choices can be read along the following lines.

The above differences between different stability conditions (A) arguably constitute objective differences in how useful these patterns are in describing worlds.

---

[7] For an explicit treatment of quantum jumps in the histories formalism, see Brun (2002).

[8] The proof is not difficult if one takes into account that the states (7) are mutually orthogonal and that the projections (8) are indeed permanent records. If the permanent records are actual records, the implicit restriction to pure states $|\Psi\rangle$ is not needed.

Instead, given that 'world' is a vague concept, the choices under (B) are objectively equally useful in making precise the notion of a world.

The examples of Deutsch, Vaidman, and Bob's electron under (C) are more interesting as they can all be understood as cases where *different patterns* in the wavefunction may be objectively useful for *different purposes*. For instance, there is an actual record of the polarisation of Vaidman's neutron in the neutron's path, or *vice versa* of the neutron's path in its polarisation, but these are microscopic records which are erased when we reinterfere the two paths. From the point of view of the neutron, it arguably makes sense to think of it not as bilocating but as splitting into two localised neutrons. From the point of view of the experimenter, instead, it may make more sense to think of it as in a superposition precisely because we then reinterfere it.[9] Parallel computations in different worlds arguably provide an explanation for how quantum computers work (at least those with a certain architecture); but, again, the necessity to reinterfere these computations at the read-out stage suggests one think of them as a superposition in a single world. Similarly, taking Bob's electron to split when Alice measures hers provides an explanation for Bob's later observation results (in particular the perfect correlations in the case of parallel measurements); but possibly no longer so if one is looking for explanations that are local in some appropriate sense. The choices in (C) may be controversial because we disagree as to what counts as useful, but we can perfectly well agree to disagree.[10]

The cases in (D) are even more controversial, I believe, because we disagree as to what counts as *objectively useful*. Take the case of Vaidman's neutron and the question of whether we split with it. In one case we have multiplicity at the global level of complete worlds, in the other we have multiplicity at the local level of individual systems.[11] It may be attractive to stick to the idea of complete worlds (maybe for

---

[9] If this fails to be intuitive for a neutron, please think of Wigner's friend.

[10] A further choice involving levels of course-graining is whether to include along with a projection also its permanent records at later times or to coarse-grain over these records. This, too, will generally be a decision depending on the purpose at hand, e.g. we may be interested in macroscopic histories and not in their microscopic records, or – as we shall see in the next section – we may want to include permanent records to ensure that our history space is branching. (Coarse-graining over records also plays a crucial role when discussing issues of time-(a)symmetry – see Bacciagaluppi (2025a).)

[11] Another example we could use is the 'many-minds' approach as proposed by Zeh (1981): local systems like brains are decohered by their environments in such a way that the resulting dynamically independent components of their reduced states give rise to a multiplicity of mental states. Global notions such as 'many worlds' do not appear in this approach.

analyses of self-location), but it also introduces disadvantages (by pre-established harmony the two counterparts of the experimenter cooperate to reinterfere the neutron). Now, are these comparative advantages rooted in aspects of the physics? It seems to me that they have nothing to do with physics and everything with our notions of identity or perhaps causation. We have histories of the form $P_{\pm} \otimes Q_0$, where $P_{\pm}$ projects onto the upper or lower path in the interferometer, and Q_0 projects onto the state of the experimenter (which does not depend on the path or polarisation of the neutron). We can read the formalism as telling us that we have clearly two different histories, so we should not identify the projection Q_0 appearing in one with the projection Q_0 appearing in the other; or that we have clearly the same projection Q_0 in both histories, so we should not distinguish between its two instances. I take it these different readings are purely 'in the eye of the beholder'. In one case, we think of histories as non-overlapping, in the other as overlapping; but there are no 'real differences' in these descriptions. (The case of whether Bob splits ahead of his own measurement is entirely analogous.)

The cases discussed under (A), (B) and (C) all involve a choice of history space, characterised by different stability conditions, different fine-grained histories, or different levels of coarse-graining. There are different patterns in the universal wavefunction, and they turn out to be: (A) useful to different degrees for the same purpose, (B) useful to the same degree for the same purpose, or (C) useful to the same degree for different purposes. The cases under (D) instead involve different descriptions of the *same* history space, and their comparative usefulness appears to stem not from objective features in the physics but from features of the descriptions themselves.[12]

I suggest that this analysis applies also to the case of splitting versus divergence of worlds. Here the issue is about whether projections before a branching event should be taken as counterparts of each other in non-overlapping histories or as the same projection in overlapping ones. Divergence has some obvious advantages in the sense that it allows us to apply familiar notions of personal identity and of probability (persons do not split and probability can be interpreted in terms of ignorance). But it also has disadvantages: if we think that our choice of measurement settings is causing the wavefunction to branch in a certain basis, it seems that we have problems thinking of causes in terms of counterfactuals: different counterparts of the experimenter will

---

[12] Of course the fine-grained histories under (B) are both descriptions of different precise worlds as well as different descriptions of the same vague world, so one can treat the choice under (B) both as substantial or as merely descriptive. The main points I am making are unaffected.

have to counterfactually cooperate in order to set differently the spin direction to be measured. For my part, I again believe that these comparative advantages are not rooted in the physics, and that even the problems we wish them to be useful for are 'in the eye of the beholder'. Thus, as with the other examples just discussed, there are no 'real differences' between a pattern of splitting worlds and one of divergent worlds, and they are merely different descriptions of the same real patterns. This is not the main claim of this paper, and discussing these non-relativistic examples has the primary aim of illustrating the strategy I will use in arguing that the case of relativistic branching (which includes the case of whether Bob splits upon Alice's measurement) also involves no 'real differences'. But the results of the relativistic case will also strengthen the claim made here that the metaphysics of real patterns does not support a substantive distinction between splitting and divergence.

3 Relativistic case

In order to discuss the relativistic case we need first of all to generalise the notions of histories and branching, because we do not have a linear ordering of time.[13] As mentioned, I am assuming throughout that we have environmental decoherence, which by the localised nature of the interactions implies that our projections are associated with (small) bounded spatiotemporal regions. That is, instead of time-indexed projections

$$P_{i_j}(t_j) \ , \tag{11}$$

we consider projections indexed by spatiotemporal regions ω_j (of given small size):

$$P_{i_j}(\omega_j) \ . \tag{12}$$

I shall also assume that the regions ω_j are pairwise either timelike or spacelike related. (Worlds defined through environmental decoherence are vague and we have some latitude in making a precise choice of ω_j.) This assumption turns both the set Ω of regions ω_j and the set of the associated families of projections into *causal sets*.[14]

---

[13] The two options for doing so described in this section are both special cases of Hartle's generalisation of the histories formalism, which he calls 'generalised quantum mechanics' (Hartle 1995).

[14] A causal set is a partially ordered set such that there are only finitely many elements between any two elements. Causal sets are the fundamental structure used in *causal set theory* to provide a framework for a theory of quantum gravity along the lines of generalised quantum mechanics from which one may recover classical spacetime as a suitable continuous approximation (since relativistic spacetimes can be thought of as largely characterised by their causal structure). For a review of causal set theory, see Henson (2009). See also the remarks in Section 5 below.

One first natural option for generalising histories is thus to generalise them to *causal sets* of projections

$$\left\{P_{i_j}(\omega_j)\right\}_{\omega_j \in \Omega} . \tag{13}$$

If for each ω_j the associated projections are mutually orthogonal and sum to the identity, we further generalise history spaces to what we shall call *causal set spaces.* To each finite causal set of projections we can associate a *causal set operator*

$$C_\beta := \prod_{\omega_j \in \Omega} P_{i_j}(\omega_j) , \tag{14}$$

where the order of the product is any total order of the ω_j that respects their partial order (because spacelike related projections commute). We can finally generalise to finite causal sets also the *decoherence functional*:

$$\mathcal{D}(C_\beta, C_{\beta'}) := \mathrm{Tr}(C_\beta|\Psi\rangle\langle\Psi|C^*_{\beta'}) , \tag{15}$$

and accordingly the conditions of *consistency* and *decoherence*:

$$\mathrm{Re}\mathcal{D}(C_\beta, C_{\beta'}) = 0 \quad \text{and} \quad \mathcal{D}(C_\beta, C_{\beta'}) = 0 \tag{16}$$

(which will be said to apply also to infinite causal sets if all finite causal subsets satisfy them).

There is a second obvious option for generalising histories to a relativistic setting. For any total ordering of the ω_j that respects their partial order we can introduce (non-uniquely) a corresponding foliation that induces that ordering. We then define histories *with respect to the total order* or the corresponding foliation. But since the decoherence functional (15) depends only on the partial order of the ω_j and is independent of any choice of total order compatible with it, the resulting alternative definitions of consistency and decoherence are in fact foliation-independent and equivalent to the invariant ones.

The definition of branching given above was to take all pairs of successive projections with non-zero transition weights, and define a space as branching if projections have unique predecessors but possibly different successors. This definition can be taken over word for word even if the order that defines pairs of successive projections is partial. It is equivalent to saying that a causal set space is branching iff all its totally ordered causal subsets are branching histories.

Alternatively, we can totally order the projections and define branching with respect to such a total order. But this yieds a more restrictive definition because

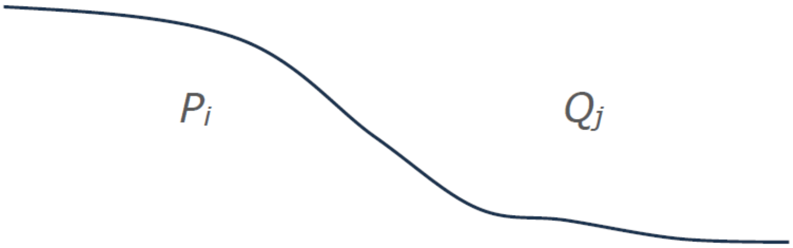

Figure 1: Causal set not branching with respect to the total order.

now there are in general *more pairs* of successive projections. E.g. if Alice and Bob measure spins at spacelike separation and we consider the causal set space consisting of the histories $\{P_i, Q_j\}$ of Alice's and Bob's spins ($i, j = \pm$), we trivially have branching with respect to the partial order, but not with respect to a total order, because in general their outcomes are not perfectly correlated (Fig. 1). Still, the two definitions are essentially equivalent, because we can fine-grain the causal set to include also the records of Alice's measurement at the time of Bob's measurement (Fig. 2). We easily see that the fine-graining that consists of the histories

$$\Big\{P_i, R_i Q_j\Big\} \tag{17}$$

is branching also with respect to the partial order (as it should by the branching-decoherence theorem). We can fine-grain the causal set to include also the records of Bob's measurement, so that it is branching with respect to a total order that reverses the two measurements. We can similarly fine-grain any finite causal set to include appropriate records for all the total orderings compatible with its partial order.

We see that if we include enough records, not only will a causal set be decoherent irrespectively of whether we define decoherence with respect to the partial order or with respect to a total order, but it will also be branching with respect to both definitions.[15]

---

[15] We also see as follows that we can equivalently express decoherence of a causal set by considering all totally ordered subsets of Ω and the associated history spaces, and requiring decoherence of all such history subspaces. In one direction, decoherence defined with respect to the partial order

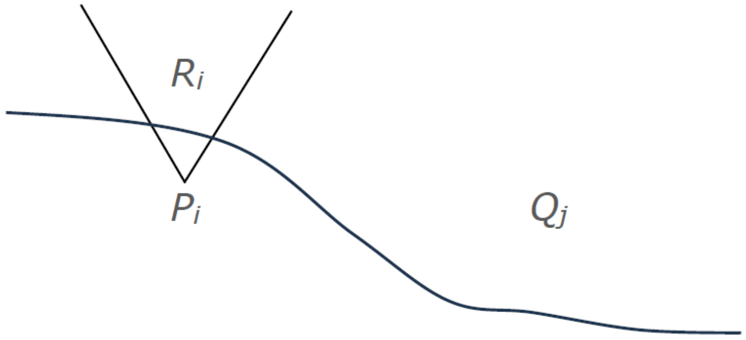

Figure 2: Finer-grained causal set branching both with respect to the partial order and the total order.

This equivalence is crucial for my claim. Including permanent records in a causal set is physically justified ('objectively useful') because permanent records are what makes branches distinct, and if we do include them the definition of branching in terms of the partial order and the one in terms of foliations arguably latch onto the *same* branching structure in the universal wavefunction. By introducing appropriate hypersurfaces, we can describe the record of a distant branching as already present, and so describe the distant branching event as having taken place already, but we are referring to the *same* branching event as in a manifestly invariant description. Insofar as frame- or foliation-dependence is not an objective feature of the physics, there are thus no 'real differences' between the branching pattern described in terms of the partial order and that described in terms of a total order.

---

clearly implies decoherence of all such subhistories. In the other direction, note that (directly from the branching-decoherence theorem) decoherent subhistories can be fine-grained by inclusion of records to obtain branching subhistories, but (as mentioned) branching with respect to subhistories is equivalent to branching with respect to the partial order, which (as in the branching-decoherence theorem) implies also decoherence with respect to the partial order. (In the case of consistency there is no analogous equivalence, as is evident using the example of composition of systems from Diósi (2004).)

4 Two natural approaches

I shall now discuss what I think are the two most natural approaches to relativistic branching in terms of the universal wavefunction, namely *hypersurface-dependent branching* and *branching along the forward lightcone*. The first of these two proposals explicitly embraces the relativity of branching and is due to Wayne Myrvold, who communicated it to me privately in Vancouver in April 2003. It is very close to his well-known proposal for hypersurface-dependent collapse (Myrvold 2000, 2002).[16] The second has been championed explicitly in print by myself (Bacciagaluppi 2002), Wallace (2012, Chap. 8) and Blackshaw, Huggett and Ladyman (this volume), and is a manifestly invariant proposal.[17]

4.1 Hypersurface-dependent branching

Recall Myrvold's ideas about relativistic collapse: collapse events are associated with (Kraus) operators localised in (small) bounded spacetime regions; a state is associated with (the future of) a spacelike hypersurface and determines the probabilities of collapses to its future given the collapses to its past; and since spacelike-related collapses commute, the net result of spacelike-related collapses does not depend on their time ordering.

For illustration, let us apply these ideas to the case of bipartite spin measurements on a pair of electrons in some arbitrary entangled state

$$\alpha|++'\rangle + \beta|+-'\rangle + \gamma|-+'\rangle + \delta|--'\rangle \ , \tag{18}$$

where Alice measures along the unprimed direction and Bob along the primed one (see Figs. 3 and 4). This will be the state to the past of both measurements. If we introduce hypersurfaces such that one measurement is to the past of the hypersurface

---

[16] I personally prefer the more detailed preprint version (Myrvold 2000) to the published paper (Myrvold 2002).

[17] I am indebted to Paul Tappenden for informing me that a version of branching along the forward lightcone (indeed, of branching spacetime) had been proposed already in the 1970s and 1980s in a series of papers by the noted French logician Roland Fraïssé (Fraïssé 1974, 1980, 1982, 1987). As for relativistic collapse, the idea of hypersurface-dependent collapse was first introduced by Albert and Aharonov (1984); Fleming (e.g. 1996) has long championed a form of hyperplane-dependent collapse; and collapse along lightcones was first introduced by Hellwig and Kraus (1970) (who however opted for collapse along the backward lightcone).

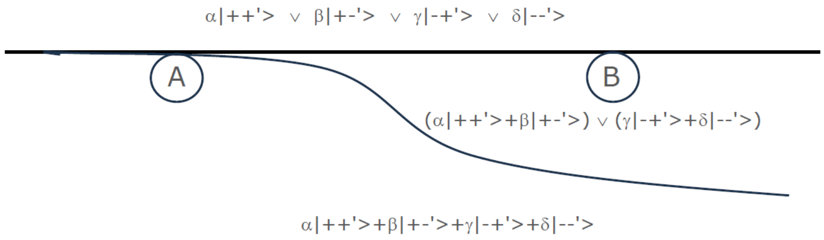

Figure 3: Collapse along one hypersurface.

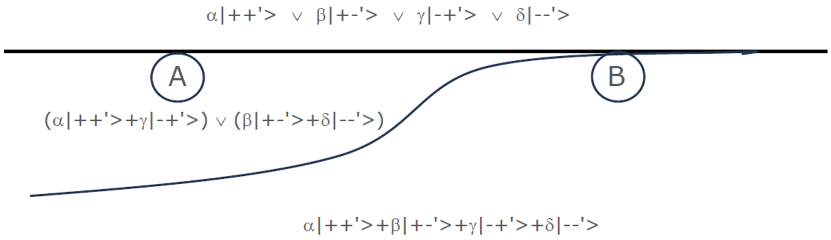

Figure 4: Collapse along another hypersurface.

and one to the future, the states to the future of these hypersurfaces will depend on the time order of the measurements. For instance, up to normalisation, if Alice's measurement takes place first, the state in the intermediate region will be one of the states

$$\alpha|++'\rangle + \beta|+-'\rangle \quad \text{or} \quad \gamma|-+'\rangle + \delta|--'\rangle \ , \tag{19}$$

and if Bob's measurement takes place first, it will be one of the states

$$\alpha|++'\rangle + \gamma|-+'\rangle \quad \text{or} \quad \beta|+-'\rangle + \delta|--'\rangle \ . \tag{20}$$

To the future of both measurements the state will equal one of the four terms in (18) (again up to normalisation).

These ideas are most easily understood from a Heisenbergian perspective, where we think of the state as a global object used to define probabilities for appropriate 'events'. A spacelike hypersurface is like a *Heisenberg cut*. We consider that on one side of the cut an event has actually happened (e.g. we have used a piece of apparatus to perform a measurement), so we update the global state we use to calculate quantum probabilities by conditionalising on that knowledge. But there is *no* matter of fact about the state having 'collapsed'. We can choose to work with the 'initial' Heisenberg state, which determines the joint probabilities for all events throughout the whole of spacetime. Or we can take a new global state conditional on events that are to the past of a hypersurface, and use this new 'collapsed' state to calculate the joint probabilities for all events in the entire spacetime region to the future of that hypersurface. And we can freely move the hypersurface anywhere further to the future and obtain again the same probabilities.[18]

For each global state that applies to a region of spacetime and for any subregion A of that region, there is a reduced state on A (i.e. the restriction of that global state to observables localised in A). And in the presence of collapse not only the global state applying to A but also the *reduced state* on A will generally depend on the hypersurface considered – specifically if in the passage from the global state associated with one hypersurface to that associated with another there is a change in the degree of entanglement between A and B. This is just the familiar phenomenon that a measurement collapses the reduced state of a distant particle. But (at least in the Heisenberg picture) one can see that this is just the 'relativity of entanglement' – to use Myrvold's lovely phrase. As pointed out by Shimony (1986), a view of quantum theory in terms of (Heisenbergian) events and their probabilities is manifestly invariant, and it is only a view in terms of (Schrödingerian) processes that violates relativity.[19],

---

[18] The historical Heisenberg understood events as observables taking values upon measurement (where a 'measurement' takes place when we have decided to use an interaction for the purpose of a measurement, rather than to reinterfere two branches of an experiment). He also understood the probabilities of quantum mechanics as *transition probabilities*, and regarded Schrödinger's 'state' simply as a mathematical artifact for calculating them. Any state that gave the correct transition probabilities was equally good, hence the 'collapse' of the state is not a physical process but a conceptual one. For the development of the 'statistical interpretation' from Born through Heisenberg to von Neumann (who gives us quantum theory essentially as we know it), see Bacciagaluppi (2022).

[19] Formally of course one can write down hypersurface-dependent states also in the Schrödinger picture or, as is often done, in the Tomonaga–Schwinger interaction picture. These will then evolve from hypersurface to hypersurface not just unitarily but also through the action of collapse operators; see e.g. Myrvold (2017) and references therein (where Myrvold in fact proves a no-go theorem showing that all such theories will have infinite energy production, unless one constructs a theory based on a 'non-standard field', or possibly unless one drops the assumption of Markovianness

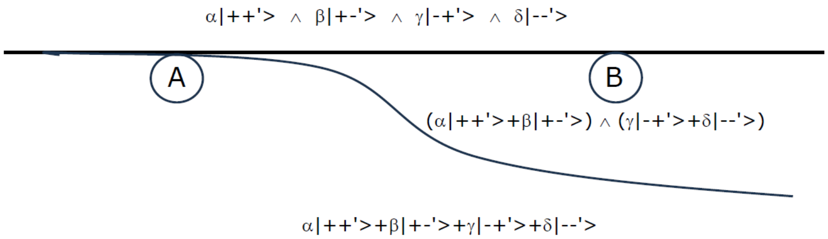

Figure 5: ‘Or’s become ‘and’s: branching along one hypersurface.

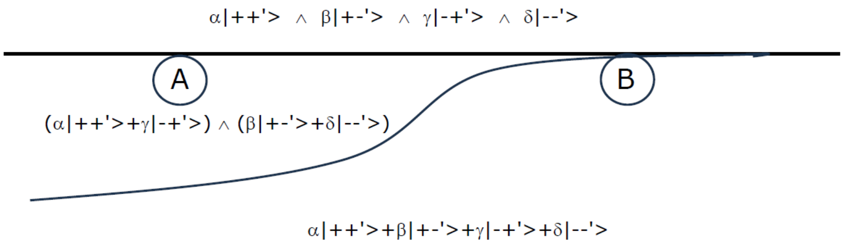

Figure 6: ‘Or’s become ‘and’s: branching along another hypersurface.

Now, Myrvold’s proposal for relativistic branching is *entirely analogous*: branching events are associated with small spacetime regions; the quantum state to the future of a spacelike surface describes the branches resulting from branching events to the past of that surface; and since spacelike-related branchings commute, the net result of spacelike-related branchings does not depend on their time ordering. The illustration using spin measurements is exactly the same as before, with the difference that the ‘or’s become ‘and’s: different states are not alternative descriptions of single worlds but coexisting descriptions of multiple ones (Figs. 5 and 6).

As opposed to the case of collapse, hypersurface-dependent branching is truly

---

or that of Minkowski spacetime). But in the presence of collapse, reduced states are ill-defined and *a fortiori* cannot be constructed from a well-defined Schrödinger or Tomonaga–Schwinger state as an (albeit non-separable) invariant object on spacetime.

neutral with respect to the choice of a Heisenberg view or a Schrödinger view, because if we consider the evolution of the Schrödinger state, branching is not reified as a physical process (unlike collapse). By introducing an arbitrary foliation, we choose to describe the Schrödinger state as evolving from hypersurface to hypersurface, and the resulting hypersurface-dependent branches are merely those components of the Schrödinger state that (because of environmental decoherence) are dynamically independent of each other given this particular description of the dynamics. Thus, the relativity of branching is due to our choices in the description of unitary dynamics.[20]

4.2 Branching along the forward lightcone

We have seen that, on a hypersurface-dependent reading of branching, given a bounded region (e.g. A) and a spacelike branching event (e.g. in B) there is *no matter of fact* about whether the quantum state in A has branched (Figs. 5 and 6). But if A and B are timelike related, say with A to the future of B, then *no matter what hypersurface* we trace separating A and B, the corresponding state in A will have branched with respect to the event in B (Fig. 7(a)). Indeed, any branching event in the past lightcone of A will affect the state on A (Fig. 7(b)). And any region in the future lightcone of B will be affected by a branching event in B. Therefore, the manifestly invariant core of hypersurface-dependent branching is *branching along the forward lightcone* (Fig. 7(c)).[21,22] In the case of bipartite spin measurements, in

---

[20] This account is quite close to that given by Ney (this volume): (environmental) decoherence is a physical process, but (hypersurface-dependent) branching is best seen as a *pseudo-process*, so that the 'spooky action at a distance' that it seems to involve is merely apparent.

[21] Note that Fig. 7(a) is showing a global 'Myrvold cut': all regions A to the future of the hypersurface are affected by branching events in all regions B to its past. Instead, Fig. 7(b) shows how the local region A is affected by branching events in all regions B to its past, and Fig. 7(c) how branching in the local region B affects all regions A to its future. To obtain a picture of what happens across spacetime if branching is defined invariantly, we need to look at all regions A and see which branching events affect each such region as in Fig. 7(b), or we need to look at all branching events and consider all the regions A they affect as in Fig. 7(c). Fig. 8 below shows the latter, and as such differs subtly from Figs. 5 or 6: in both cases we have multiple states applying to the future of certain spacelike or piecewise lightlike hypersurfaces, and thus an association between spacetime regions and the corresponding set of branches. But Figs. 5 or 6 also illustrate how the state to the future of a hypersurface reflects all the branching events to its past, while Fig. 8 no longer indicates this.

[22] There is another invariant notion of branching that one can introduce, namely by considering a branching in A due to an event in a region B not with respect to *all* hypersurfaces but with respect to *at least one* hypersurface. Thus, A is affected by branching events in all regions outside of its

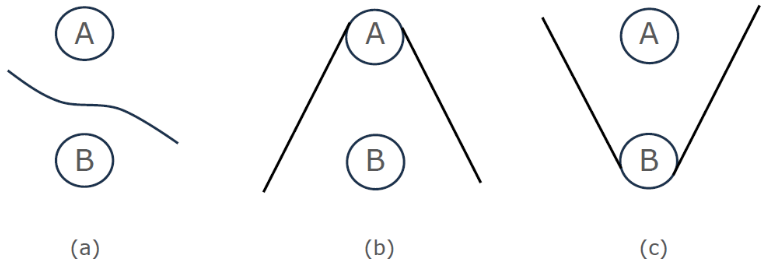

Figure 7: (a) region of branching events affecting A and region affected by branching events in B, given a Myrvold cut; (b) region of branching events invariantly affecting A; (c) region invariantly affected by branching events in B.

the forward lightcone of Alice's measurement the state will have branched such that different branches contain different outcomes of her measurement, and similarly for the future lightcone of Bob's measurement; thus, in the intersection of the two forward lightcones the wavefunction will have branched such that different branches contain different *pairs* of outcomes of Alice's and Bob's measurements.

If we think of a branching global Heisenberg state, we shall have the situation depicted in Fig. 8: to the future of A, we distinguish the branches

$$\alpha|++'\rangle + \beta|+-'\rangle \quad \text{and} \quad \gamma|-+'\rangle + \delta|--'\rangle \ , \tag{21}$$

and to the future of B we distinguish the branches

$$\alpha|++'\rangle + \gamma|-+'\rangle \quad \text{and} \quad \beta|+-'\rangle + \delta|--'\rangle \ . \tag{22}$$

In general these are incompatible ways of distinguishing branches in the global state, in the sense that each branch on Alice's side is described by a superposition of components belonging to different branches on Bob's side and *vice versa.* But when the two branches on Alice's side split again in the intersection of the two forward lightcones, the resulting branches that we distinguish are given by the four terms in (18), and are the *same* four components of the global state that we distinguish as

---

forward lightcone, and branching events in a region B affect all regions outside its *past* lightcone. We shall not discuss this 'branching along the past lightcone' any further (but, as already mentioned, see Hellwig and Kraus (1970) for collapse along the backwards lightcone).

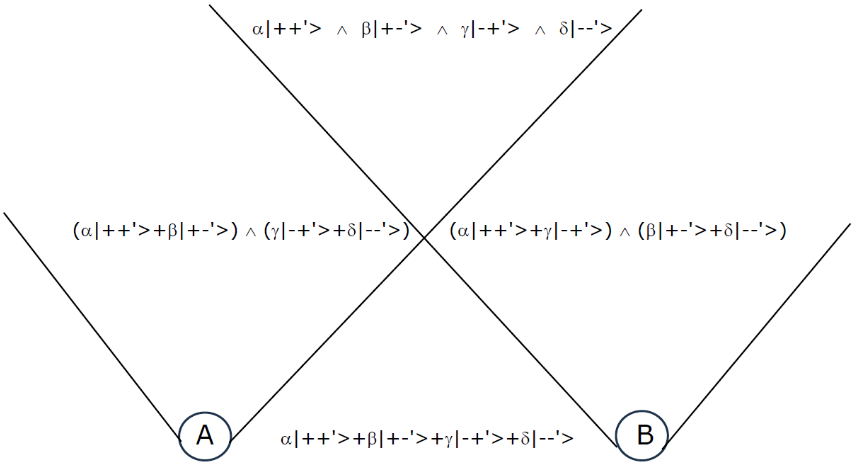

Figure 8: Branching along the forward lightcone.

branches when the two branches on Bob's side split again.[23]

Similarly, if we consider the Schrödinger picture, the global state is evolving unitarily on spacetime, and the reduced states of bounded regions are well-defined. We can then distinguish as local branches the dynamically independent components of the reduced states in the future lightcones of A and B. Again, the two branchings are in general unrelated, but when the reduced states branch again in the intersection of the two future lightcones, we have the same four components resulting from the two successive branchings.

This account of branching is arguably local in the sense that different branches are defined in the future lightcone of branching events, but note that it is the entanglement in the global state (in either Heisenberg or Schrödinger picture) that allows us to identify the branches in the common future of the two measurements regardless of whether we consider a history running from the preparation, through Alice's measurement, to the common future of the two measurements, or one running instead through Bob's measurement. Any separable state (other than a mixture of the four components in (18)) would lead to *more* branches in the intersection of the two future

---

[23] In the special case in which Alice and Bob both measure a Schmidt basis on their side we have, say, $\beta = \gamma = 0$, and we see that (21) and (22) coincide: the global state splits into the same two components already with Alice's and Bob's original measurements.

lightcones.[24] It is thus more perspicuous to speak of *nested locality*: the branching structure due to the branching event in Alice's or in Bob's lab is describable locally using the reduced states on that side, but the combined branching structure due to the two branching events in Alice's and Bob's labs can be described locally only using the reduced state on a larger region of spacetime,[25] and not simply combining the separate local accounts in the future lightcones of the two measurements. It is this combination of locality and non-separability that ensures the 'peaceful coexistence' of quantum mechanics and relativity, to borrow another concept from Shimony (1978, 1986).[26]

Comparing with the more abstract discussion of Section 3, hypersurface-dependent branching as described in this section corresponds to the notion of branching in terms of a total order. Instead, branching along the forward lightcone corresponds to the notion of branching in terms of the partial order. Both approaches agree on what branching events there are, and (since I am assuming environmental decoherence) both distinguish branches by the formation of records. But when a record is deemed to have formed depends on *what we describe as past*, so the two approaches differ precisely in that they adopt a description of past and future in terms of lightcones or in terms of foliations. And this is not a 'real difference'.[27]

Recall also my remarks on splitting and divergence in Section 2. If we apply the picture of divergence to the relativistic case, then worlds are different across the whole of spacetime, and the distinction between hypersurface-dependent branching and branching along the forward lightcone becomes (if anything even more clearly) just one between different descriptions of when worlds are deemed to diverge. If we

---

[24] Note incidentally that the picture of diverging worlds is analogous to the case of such a mixture, and that this in turn is analogous to simple measurement-dependent hidden variables models, as emphasised by Bacciagaluppi, Hermens and Leegwater (2025).

[25] I believe this analysis is neutral with respect to a choice of Schrödinger or Heisenberg state, and thus also compatible with the Heisenberg-picture analysis provided by Kuypers (this volume), although ontologically leaner. For contrasting views on Heisenberg-picture ontology, see Bédard (this volume) and Timpson and Wallace (this volume).

[26] This is precisely the case also with collapse along the forward lightcone. I take the opportunity to note that the latter was not yet my understanding in my (Bacciagaluppi 2002, Sect. 6.3), where I was thinking of collapse along the forward lightcone as failing to reproduce distant correlations because local collapses would be independent of each other, and I did not yet have the concept of nested locality even for the case of branching. I shall argue elsewhere that nested locality is the correct formulation of 'local causality' – removing the alleged incompatibility between relativistic locality and Bell nonlocality (Bacciagaluppi 2025b).

[27] One could frame in terms of real patterns the very discussion about the relativity of simultaneity, but I shall leave this point largely implicit.

apply the picture of splitting, then we observe something even more interesting. The records that form upon a branching event are localised only in the forward lightcone of the branching event, therefore if we adopt hypersurface-dependent branching, two branches will be locally distinguishable only in that forward lightcone, and locally indistinguishable outside. That means that even though two branches have *split* along a spacelike hypersurface, they still *diverge* outside of the future lightcone of the corresponding branching event (which makes it perhaps less natural to combine splitting with hypersurface-dependent branching).

I have argued that hypersurface-dependent branching and branching along the forward lightcone are two sides of the same coin. Still, I find branching along the forward lightcone (combined with splitting) to be the most natural description of relativistic branching, not merely because it is manifestly invariant but because it aligns with the physical process of formation of records. Its most interesting feature, however, will arguably be that it offers a perspective on branching *in terms of spacetime*, which we discuss in the next section.[28]

5 Branching spacetime

We saw that in the case of hypersurface-dependent branching, the entire region to the future of a hypersurface (including e.g. the contents of Bob's lab) is taken to branch with respect to the branching events to the past of the hypersurface. This is regardless of any interactions with the regions containing the branching events (e.g. Alice's measurement). That a branching event affects the entire region to its future is assumed also in the case of branching along the forward lightcone (Fig. 7(c)), but in this case physical interaction can occur, and one typically expects a cascade of records to form in the forward lightcone of a decoherence event, so that the assumption that the whole of the future lightcone bears records of a branching event, and that different branches extend throughout the future lightcone, is physically justified (at least as an idealisation).

Treating branching thus, as applying universally to the entire future lightcone of a branching event, suggests that we describe it not as branching of the *material content* of the lightcone, but as branching *of the lightcone itself* into two or more 'leaves' of spacetime with different material content. In fact, this must be so if one sees a

---

[28] This section is based and further elaborates on (Bacciagaluppi 2002).

relativistic spacetime as characterised by the structure of *causal relations* between events. In the Everett theory, different local branches are to the causal future of their branching event, but (barring recoherence) they are also causally inaccessible to each other, so that they must indeed correspond to *separate regions* of spacetime. This gives a picture of a *branching spacetime* – which locally branches along the future lightcones of branching events, and globally is characterised by how the leaves in the common future of any two branching events are pasted together. Thus we can literally think of different copies of Alice and Bob following *different paths* in spacetime (but each copy of Alice meeting a copy of Bob where the spacetime leaves are pasted together again).[29]

The foremost advantage of understanding branching as branching of spacetime is that it provides us with a framework for *how spacetime can emerge from the universal wavefunction*. We should distinguish between *background spacetime* and *concrete spacetime*. Background spacetime is simply the Minkowski spacetime that we use to represent relativistic quantum theory when we talk about observables or states associated with spacetime regions, and I believe it is to be understood as encoding the universal dynamical symmetries of the theory (i.e. the Poincaré group) along the lines notably proposed by Brown (2005). Arguably, quantum theory provides an especially clear-cut case for this 'dynamical' view of spacetime, because in Hilbert space one cannot make the case that spatial notions are somehow presupposed. Concrete spacetime by contrast is a collection of 'concrete events' such as fingers snapping, fire-crackers exploding, Geiger counters clicking, or vials of poison being smashed. It is the classical spacetime we know and love, and it is not the same as the background spacetime, because due to the measurement problem and related issues such concrete events are difficult to find in quantum theory. Background spacetime has thus approximately the same status as configuration space or momentum space in non-relativistic quantum mechanics: configuration space is also a background space we use to represent quantum mechanics, but that representation leaves open the question 'configurations of what?'. Decoherence is a familiar tool in trying to solve the problem of recovering familiar phenomena from quantum theory, whether it be

---

[29] There is a long tradition of understanding relativistic spacetimes in terms of causal structure, which has become especially important since the work of Malament (1977). Here I understand 'causal structure' in the appropriately thin sense used in this literature of functional relations between physical quantities. Should someone object that quantum theory allegedly includes spacelike causation, my answer is that the spacelike correlations in quantum theory are relevant to the global causal structure of spacetime by way of how spacetime leaves get pasted together in the common future of branching events, but that all direct causal relations are and remain timelike. See again Bacciagaluppi (2025b).

the classical world around us or even the quantum phenomena in it.[30] Recovering concrete events is part of this problem, and I propose to identify concrete spacetime events with *branching events*, which I have indeed assumed are decoherence events. Crucially, the spacetime that we recover is a *branching spacetime*, unlike the background spacetime.[31]

6 Conclusion

Everett himself was an empiricist, but if one takes a realist view of the Everett theory, presumably the default position is to take the universal wavefunction as the fundamental object of the theory.[32] In its most radical form, as explicitly advocated e.g. by Carroll (2022), this view takes the Everett theory to be about an abstract vector evolving unitarily in Hilbert space. The question that immediately arises is whether this is sufficient to explain the detailed structure of worlds as we see them, with alternatives proposed e.g. by Wallace and Timpson (2010) and by Bédard (this volume). The proposal described in Bacciagaluppi (2002) and in the last section is explicitly meant to explain how spacetime as we know it can emerge from the universal wavefunction, with more familiar applications of decoherence doing the

---

[30] In a completely isolated atom, *nothing ever happens*, and it is only in the presence of environmental decoherence that we have *quantum jumps*.

[31] It is also a *discrete* spacetime, but if decoherence scales are sufficiently small I believe one can adapt standard results from causal set theory for recovering continuous approximations of a causal set: not only do branching events form a causal set, but presumably the analogy with relativistic collapse extends also to an effectively random distribution of the locations of the branching events with respect to the background spacetime, so that these satisfy the condition of 'sprinkling' from causal set theory. Other possible connections with causal set theory ought to be explored. For instance, in order to have enough branching events to justify a continuous approximation it may be necessary to hypothesise that the relevant decoherence interactions arise from a quantised theory of gravity. This would also justify the idealisation that branching is universal, since gravitation is. The 'dynamics' of the branching causal set in the Everett theory is 'covariant' in the sense of causal set theory, because it is independent of the order in which spacelike branchings are applied. Everettian branching could be a way of understanding where causal sets come from, and furthermore suggests that there is no need to reject unitary dynamics in causal set theory. Finally, it is suggestive for generalisations to quantum gravity that (on the relevant scale) in an Everettian branching spacetime the causal structure depends on the correlational structure in the quantum state.

[32] Everett's thesis was titled *Theory of the Universal Wave Function*, and the wavefunction is the central theoretical object in his original formulation. But Everett had an explicit empiricist view of theory: see Appendix II of his thesis, 'Remarks on the Role of Theoretical Physics' (Everett 2012, pp. 168–171).

rest (Saunders 1993; Wallace 2012, Part I).

On the other hand, as I have remarked in various places in this chapter, the approach described here is also meant to be neutral with respect to questions of ontology in the Everett theory. In describing branching along the forward lightcone and its reading in terms of branching spacetime, I have adopted the language of the universal wavefunction, with branching structure understood as a real pattern arising dynamically from the wavefunction. But one can see the *branching events* themselves as the fundamental beables of the theory, and the wavefunction as playing the nomological role of determining their dynamics. This is analogous to the *flash ontology* in collapse theories.[33] Flash ontology, however, is often presented as one of *structureless* point events (Allori *et al.* 2008). Shimony's (1986) distinction between (Schrödinger) process-based ontology and (Heisenberg) event-based ontology is not committed to events being so austere. Heisenberg himself understood events as being observables taking on certain values, and a branching event is not characterised merely by its location in the background spacetime, but also by the associated projections. Distinguishing between *different* branches seems to require a richer notion of event than a mere 'flash'. Such a richer notion of event will also likely be needed if one wishes to truly follow Heisenberg's intuition and formulate the dynamics purely in terms of *transition* probabilities between events, rather than defining it in terms of a quantum state (even one reduced to a mere Heisenberg state). Branching would then become the central notion in an Everett theory without quantum states(!).

References

Albert, D. Z., and Aharonov, Y. (1984), 'Is the usual notion of time evolution adequate for quantum-mechanical systems? II. Relativistic considerations', *Physical Review D* **29**, 228–234.

Allori, V., Goldstein, S., Tumulka, R., and Zanghì, N. (2008), 'On the common structure of Bohmian mechanics and the Ghirardi–Rimini–Weber theory', *The British Journal for the Philosophy of Science* **59**(3), 353–389.

Bacciagaluppi, G. (2002), 'Remarks on space-time and locality in Everett's interpretation', in T. Placek and J. Butterfield (eds.), *Non-locality and Modality* (Dordrecht: Kluwer), pp. 105–122.

---

[33] Note also the analogy with 'growth models' of the dynamics in causal set theory.

Bacciagaluppi, G. (2013), ‘Book review of: *The Everett Interpretation of Quantum Mechanics. Collected Works 1955–1980 with Commentary* by Hugh Everett III’, *HOPOS* **3**, 348–352 (version with footnotes and references at `http://philsci-archive.pitt.edu/10768/`).

Bacciagaluppi, G. (2022), ‘The statistical interpretation: Born, Heisenberg and von Neumann, 1926–27’ in O. Freire *et al.* (eds.), *The Oxford Handbook of the History of Quantum Interpretations* (Oxford: Oxford University Press), pp. 203–231.

Bacciagaluppi, G. (2025a), ‘The misleading equivalence of decoherence and branching’ (working title), in preparation.

Bacciagaluppi, G. (2025b), ‘Against “local causality”’, in preparation.

Bacciagaluppi, G., Hermens, R., and Leegwater, G. (2025), ‘Extending Bell’s theorem: nonlocality via measurement dependence’ (working title), in preparation.

Bédard, C. (this volume), ‘The Heisenberg and the Schrödinger pictures are realistically inequivalent’.

Blackshaw, N., Huggett, N., and Ladyman, J. (this volume), ‘Everettian branching in the world and of the world’.

Brown, H. (2005), *Physical Relativity: Space-Time Structure from a Dynamical Perspective* (Oxford: Oxford University Press).

Brun, T. (2002), ‘A simple model of quantum trajectories’, *American Journal of Physics* **70**(7), 719–737.

Carroll, S. (2022), ‘Reality as a vector in Hilbert space’, in V. Allori (ed.), *Quantum Mechanics and Fundamentality: Naturalizing Quantum Theory between Scientific Realism and Ontological Indeterminacy* (Cham: Springer), pp. 211–224.

Deutsch, D. (1985), ‘Quantum theory, the Church–Turing principle and the universal quantum computer’, *Proceedings of the Royal Society of London A* **400**(1818), 97–117.

Diósi, L. (2004), ‘Anomalies of weakened decoherence criteria for quantum histories’, *Physical Review Letters* **92**, 170401/1–3.

Everett, H. III (2012), *The Everett Interpretation of Quantum Mechanics. Collected Works 1955–1980 with Commentary*, edited by J. A. Barrett and P. Byrne (Princeton: Princeton University Press).

Fleming, G. (1996), 'Just how radical is hyperplane dependence?', in R. Clifton (ed.), *Perspectives on Quantum Reality: Non-relativistic, Relativistic, and Field-theoretic* (Dordrecht: Kluwer), pp. 11–28.

Fraïssé, R. (1974), 'Essai sur la logique de l'indéterminisme et la ramification de l'espace-temps', *Synthese* **29**, 27–54. Reprinted in P. Suppes (ed.), *Logic and Probability in Quantum Mechanics* (Dordrecht: Springer, 1976), pp. 19–46.

Fraïssé, R. (1980), 'L'échec de l'inégalité de Bell, une présomption en faveur de la ramification d'Everett', *Epistemological Letters* **27**, 17–40.

Fraïssé, R. (1982), 'Quelques arguments en faveur de l'interprétation de la mécanique quantique par la ramification d'Everett', *Synthese* **50**, 325–357.

Fraïssé, R. (1987), 'La ramification selon Everett: une interprétation critiquée mais logiquement impeccable de la mécanique quantique', *History and Philosophy of Logic* **8**(2), 215–241.

Griffiths, R. (1993), 'Consistent interpretation of quantum mechanics using quantum trajectories', *Physical Review Letters* **70**, 2201–2204.

Hartle, J. B. (1995), 'Space-time quantum mechanics and the quantum mechanics of space-time', in J. Zinn-Justin and B. Julia (eds.), *Gravitation and Quantizations: Proceedings of the Les Houches Summer School, Session LVII, 5 July–1 August 1992* (Dordrecht: North-Holland), pp. 285–480.

Hellwig, K.-E., and Kraus, K. (1970), 'Formal description of measurements in local quantum field theory', *Physical Review D* **1**, 566–571.

Henson, J. (2009), 'The causal set approach to quantum gravity', in D. Oriti (ed.), *Approaches to Quantum Gravity: Toward a New Understanding of Space, Time and Matter* (Cambridge: Cambridge University Press), pp. 393–413.

Kuypers, S. (this volume), 'Restoring locality'.

Malament, D. B. (1977), 'The class of continuous timelike curves determines the topology of spacetime', *Journal of Mathematical Physics* **18**(7), 1399–1404.

Myrvold, W. (2000), 'Einstein's untimely burial', `http://philsci-archive.pitt.edu/222/` .

Myrvold, W. (2002), 'On peaceful coexistence: is the collapse postulate incompatible with relativity?', *Studies in History and Philosophy of Modern Physics* **33**, 435–466.

Myrvold, W. (2017), 'Relativistic Markovian dynamical collapse theories must employ nonstandard degrees of freedom', *Physical Review A* **96**(6), 062116/1–5.

Ney, A. (this volume), 'Branching (almost) everywhere and all at once'.

Saunders, S. (1993), 'Decoherence, relative states, and evolutionary adaptation', *Foundations of Physics* **23**(12), 1553–1585.

Shimony, A. (1978), 'Metaphysical problems in the foundations of quantum mechanics', *International Philosophical Quarterly* **18**, 3–17.

Shimony, A. (1986), 'Events and processes in the quantum world' , in R. Penrose and C. J. Isham (eds.), *Quantum Concepts in Space and Time* (Oxford: Clarendon Press), pp. 182–203.

Tappenden, P. (2008), 'Saunders and Wallace on Everett and Lewis', *The British Journal for the Philosophy of Science* **59**(3), 307–314.

Timpson, C., and Wallace, D. (this volume), 'Locality in Everett: dynamics vs ontology'.

Vaidman, L. (1998), 'On schizophrenic experiences of the neutron or why we should believe in the many-worlds interpretation of quantum theory', *International Studies in the Philosophy of Science* **12**(3), 245–261 (more detailed version at `https://arxiv.org/abs/quant-ph/9609006`).

Wallace, D. (2012), *The Emergent Multiverse* (Oxford: Oxford University Press).

Wallace, D., and Timpson, C. G. (2010), 'Quantum mechanics on spacetime I: spacetime state realism'*The British Journal for the Philosophy of Science* **61**(4), 697–727.

Wilson, A. (2020), *The Nature of Contingency: Quantum Physics as Modal Realism* (Oxford: Oxford University Press).

Wilson, A. (this volume), 'Explanations of and in worlds'.

Zeh, H.-D. (1981), 'The problem of quantum observation in quantum mechanical description', *Epistemological Letters* **30**, 1–12. (A slightly revised version with 'Addendum' was published in *Foundation of Physics Letters* **13**, 221–233 (2000).)