The Relativity of Branching

Guido Bacciagaluppi

Conceptual Foundations of QFT

Vienna, 1 June 2026

https://philsci-archive.pitt.edu/25071/

Forthcoming in A. Ney (ed.), *Local Quantum Mechanics: Everett, Many Worlds and Reality* (OUP).

Everett theory:

(i) find stable structures at the level of components within the universal wavefunction (1957: patterns of relative states; ≥ 1990s: decoherent histories)

(ii) identify these as multiple ‘branches’ or ‘worlds’

Question: ′
How does branching fit with Minkowski spacetime (and thus QFT)?

Some authors think there is a problem of locality (Carroll and Sebens 2018): we need not go into details, but the argument presumably *reifies* branching as a physical process.

We instead will take branches to be *real patterns* (~ à la Wallace):

- they may be ‘in the eye of the beholder’
- but *objectively useful* for given purposes, e.g. prediction or explanation
- (with ‘objective usefulness’ grounded in physics)

Then, there may be *different* patterns that are all *real* and thus legitimate choices for ‘worlds’…

E.g. do quantum computers perform computations in parallel worlds (even if short-lived), or do we have a superposition in a single world?

(Different purposes: understanding the computation as a computation, or understanding the read-out as an interference effect.)

Main claim:

Natural proposals for relativistic branching may be different, but *equally real* in the above sense.

Differences arise through different choices of the *time ordering* of events, which are just different descriptions of the *same* invariant physics.

Plan of talk:

(A) Two natural proposals

(B) General formalism

(C) Locality and spacetime

.

(A) *Two natural proposals*

We start with *hypersurface-dependent branching* (Wayne Myrvold, private communication, 2003):

- branching events are associated with small spacetime regions
- the description to the future of a spacelike surface reflects the branching events to its past
- spacelike-related branchings commute, so the net result of spacelike-related branchings does not depend on their time ordering

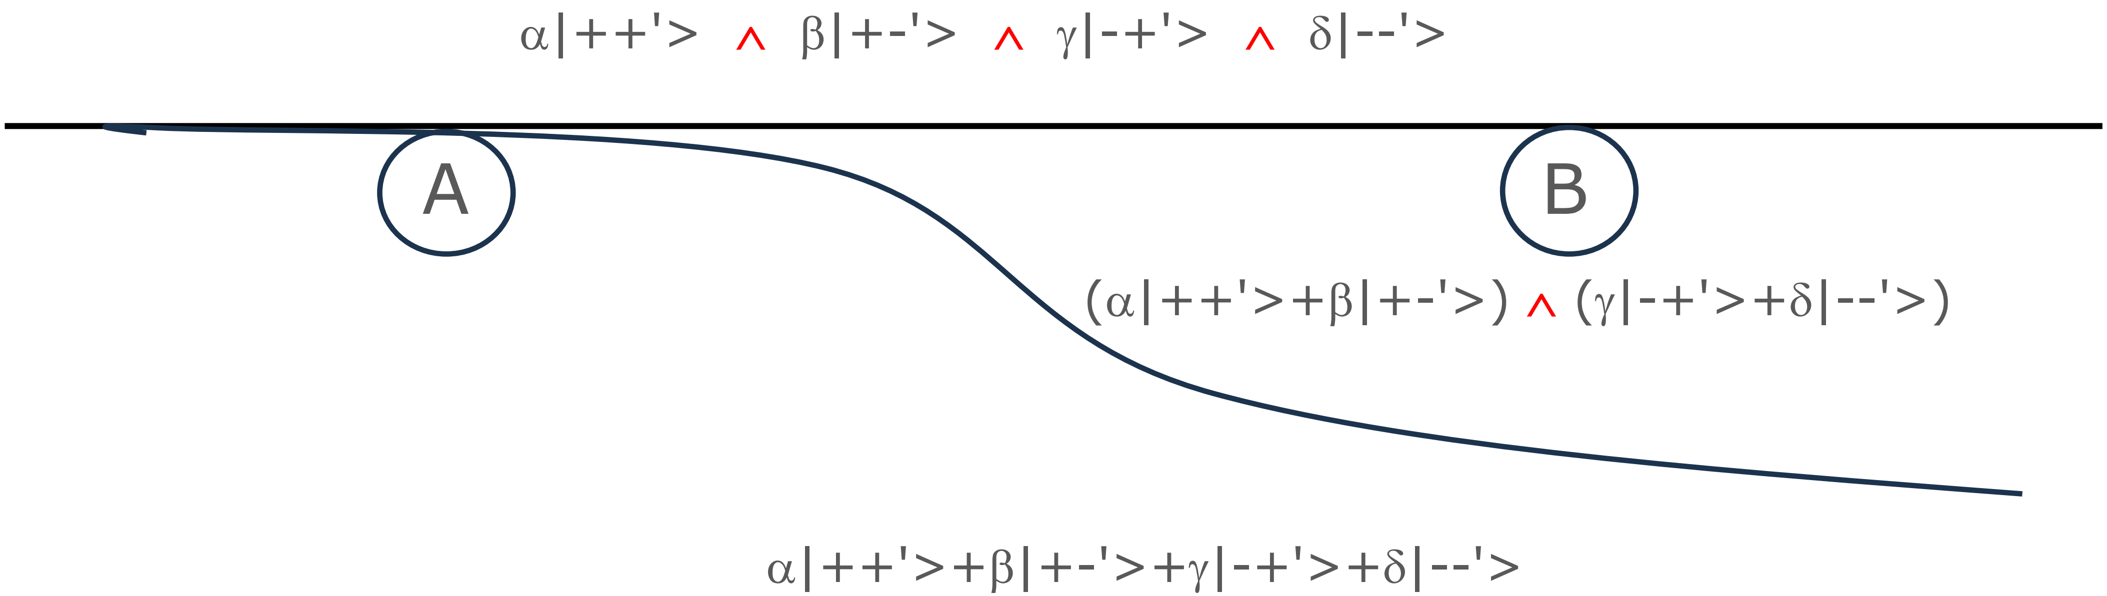
α|++'> ∧ β|+-'> ∧ γ|-+'> ∧ δ|--'>
A
B
(α|++'>+β|+-'>) ∧ (γ|-+'>+δ|--'>)
α|++'>+β|+-'>+γ|-+'>+δ|--'>

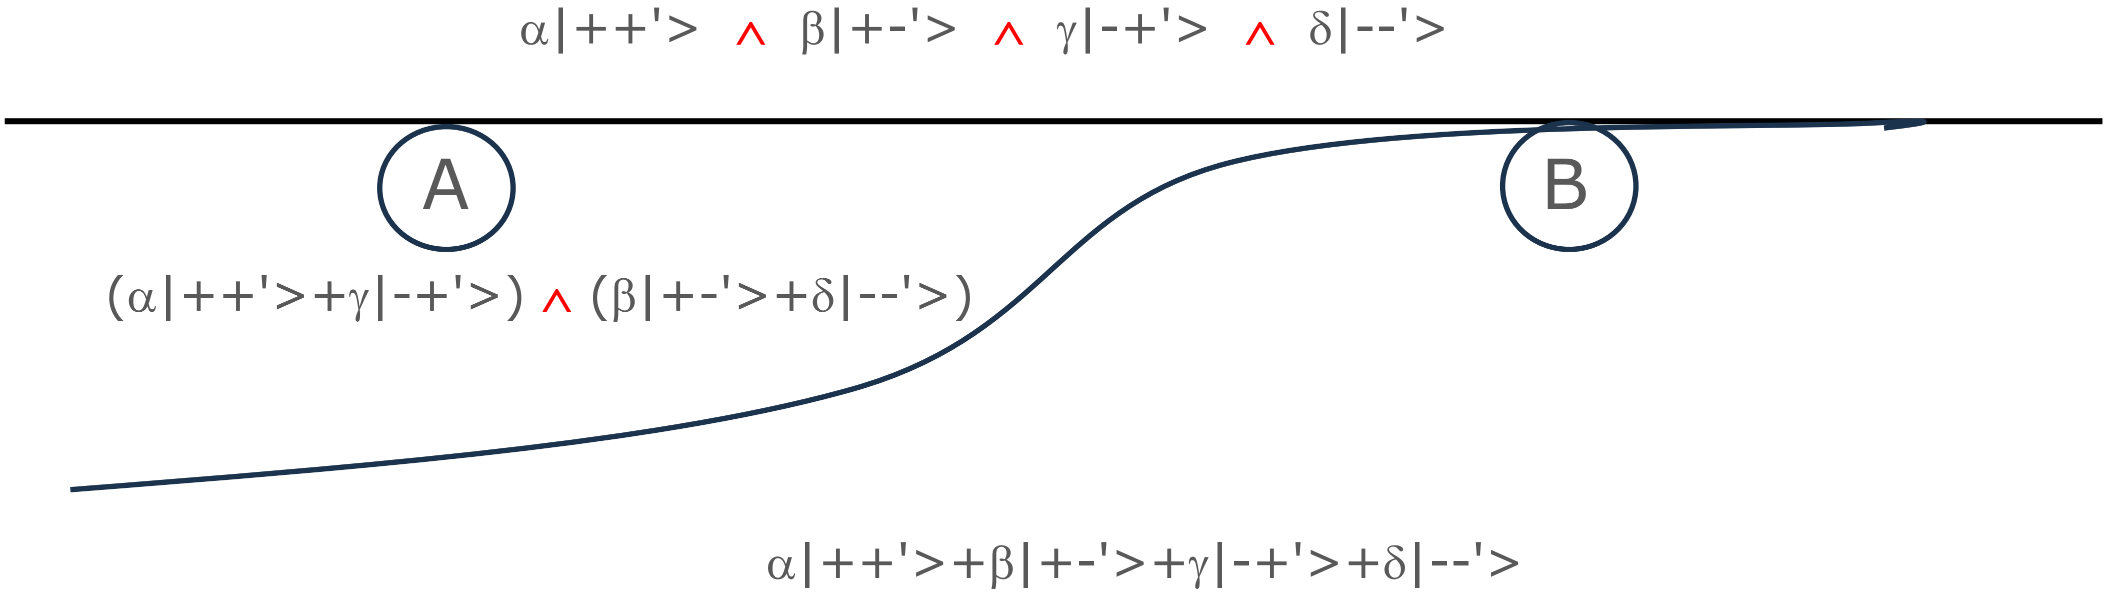
α|++'> ∧ β|+-'> ∧ γ|-+'> ∧ δ|--'>
A
B
(α|++'>+γ|-+'>) ∧ (β|+-'>+δ|--'>)
α|++'>+β|+-'>+γ|-+'>+δ|--'>

Note the analogy with Wayne’s understanding of *collapse* as depending on a choice of hypersurface (a Heisenberg cut).

Also in the case of branching, given a bounded region (e.g. A) and a spacelike branching event spacelike (e.g. in B) there is *no matter of fact* about whether the quantum state in A has branched, because it depends on the choice of hypersurface (a ‘Myrvold cut’).

But there are cases where branching is *invariant*,
namely when B is in the past lightcone of A
(A is in the future lightcone of B)

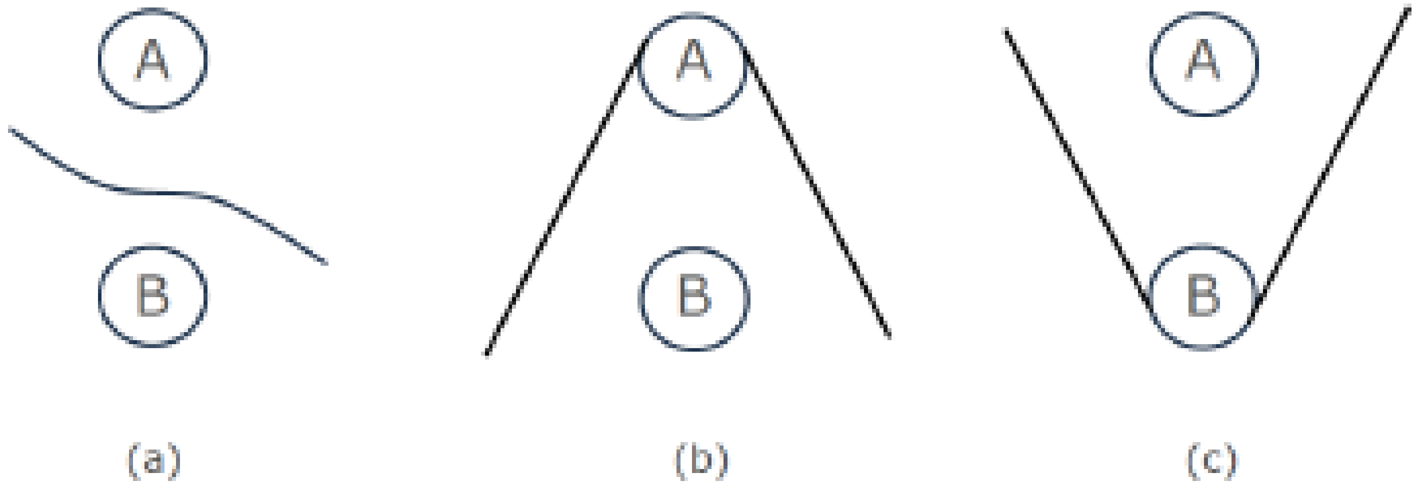

Figure 7: (a) region of branching events affecting A and region affected by branching events in B, given a Myrvold cut; (b) region of branching events invariantly affecting A; (c) region invariantly affected by branching events in B.

We thus get a second natural proposal, namely *branching along the forward lightcone* (GB 2002, Wallace 2012, Blackshaw *et al*. forthcoming):

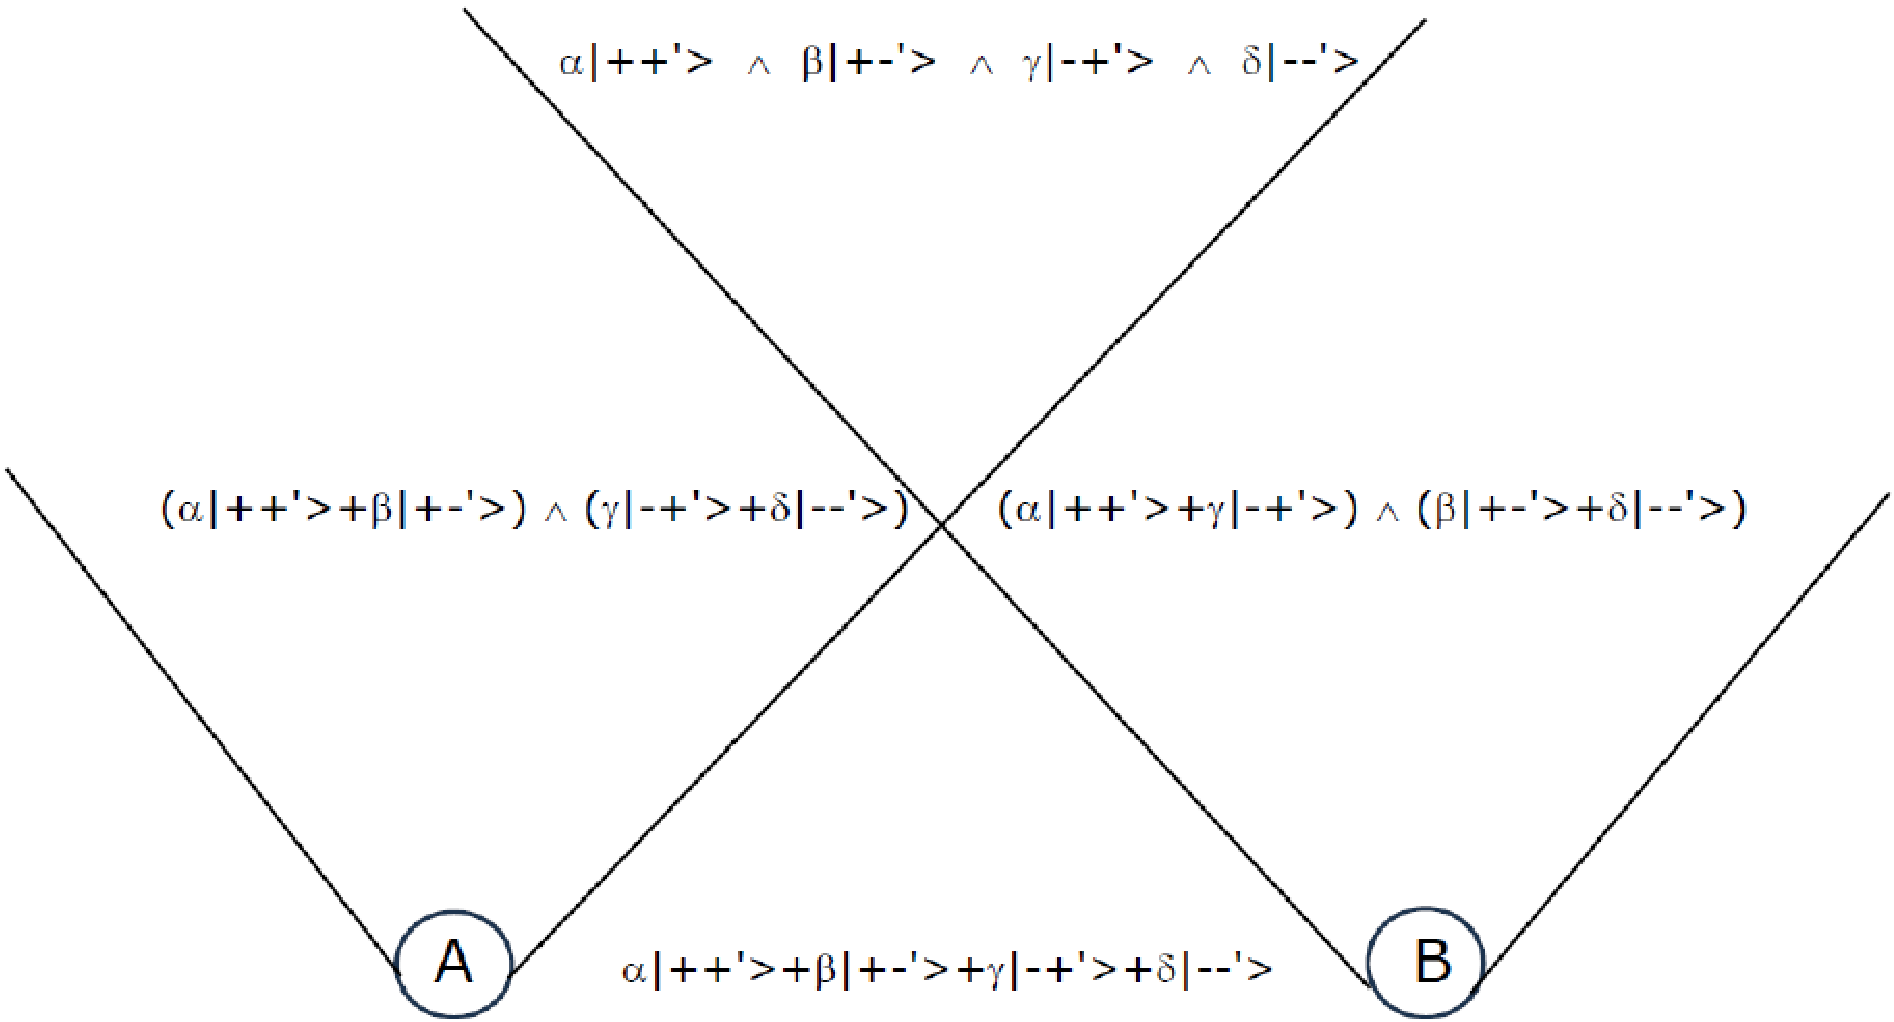

I wish to argue that hypersurface-dependent branching is just a *redescription* in different language of the *same* invariant physical structure: the two different branching patterns are *equally real*.

Below I sketch the formal tools to make this precise, but here is the informal argument...

(1) In the future lightcone of a branching event we have permanent records of it in the environment.

(2) These permanent records are what we use to label the branches in the forward lightcone.

(3) If we choose a foliation or a frame, we treat as *already present* also records of *distant* branchings.

(4) Thus we can use these records to label branches also to the future of spacelike hypersurfaces.

(5) These two branching patterns are grounded in the same physical records of the same events, and are thus *equally real*.

(B) *General formalism*

Recall the non-relativistic formalism of *histories*.

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

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

We define the associated *history operator* as:

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

If for all times the projections are mutually orthogonal and exhaustive, we have a *history space*.

For any two histories in a history space we define the *decoherence functional* as

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

and the positive number

$$\mathcal{D}(C_\alpha, C_\alpha)$$

as the *weight* of a history. Analogously as with probabilities, we can define ‘transition weights’ between two successive projections in a history.

We now define *stability conditions* for history spaces (extendable to denumerably infinite histories if satisfied by all finite subhistories).

Consistency (= lack of interference):

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

Decoherence (= existence of 'generalised records'):

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

Environmental Decoherence: when *actual records* form in the environment.

Note that if any of these are satisfied, we have *real patterns* in the wavefunction, but environmental decoherence is arguably the *most useful*.

Finally, we define *branching*:

- take all pairs of successive projections with non-zero transition weights
- the space is *branching* if projections have unique predecessors but possibly different successors

Whenever a projection does have different successors, we talk of a *branching event*.

In Minkowski spacetime we do not have a linear ordering of time. What to do?

Assuming environmental decoherence, projections are associated with *localised* spatiotemporal regions.

That is, instead of time-indexed projections

$$P_{i_j}(t_j)$$

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

$$P_{i_j}(\omega_j)$$

Everett worlds are vague, so we have some latitude in defining the ω_j. We shall conveniently choose them so that they are pairwise either timelike or spacelike related.

This turns the set Ω of regions ω_j and the set of the associated families of projections into *causal sets* (partially ordered sets such that there are only finitely many elements between any two elements).

(N.B. We are thereby *bypassing the impossible measurements problem*!)

A first option for generalising histories and history spaces is thus to generalise them to *causal sets* of projections

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

(where for each ω_j the associated projections are mutually orthogonal and sum to the identity).

To each finite causal set of projections we can associate a *causal set operator*

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

where the order of the product can be any total order of the ω_j that respects their partial order, because spacelike related projections commute.

We can thus generalise to finite causal sets also the *decoherence functional*:

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

And correspondingly *consistency* and *decoherence*:

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

(also extendable to denumerably infinite causal sets).

Alternatively, for any total ordering of the ω_j that respects their partial ordering we can introduce (non-uniquely) a corresponding foliation that induces that ordering.

We can then define *histories* wrt the total ordering or the corresponding foliation.

But causal set operators are independent of the total order, so it follows that consistency and decoherence are foliation-independent and *equivalent* to the manifestly invariant notions.

The notion of *branching* is more complex, however.

We can define it as above, both if the order is partial (and thus invariantly) and if it is total:

- take all pairs of successive projections with non-zero transition weights
- the space is *branching* if projections have unique predecessors but possibly different successors

But in the second case, the definition is *more restrictive* because there are *more pairs* of successive projections.

E.g. if Alice and Bob measure spins at spacelike separation, we trivially have branching wrt the partial order, but not wrt a total order, because in general their outcomes are not perfectly correlated:

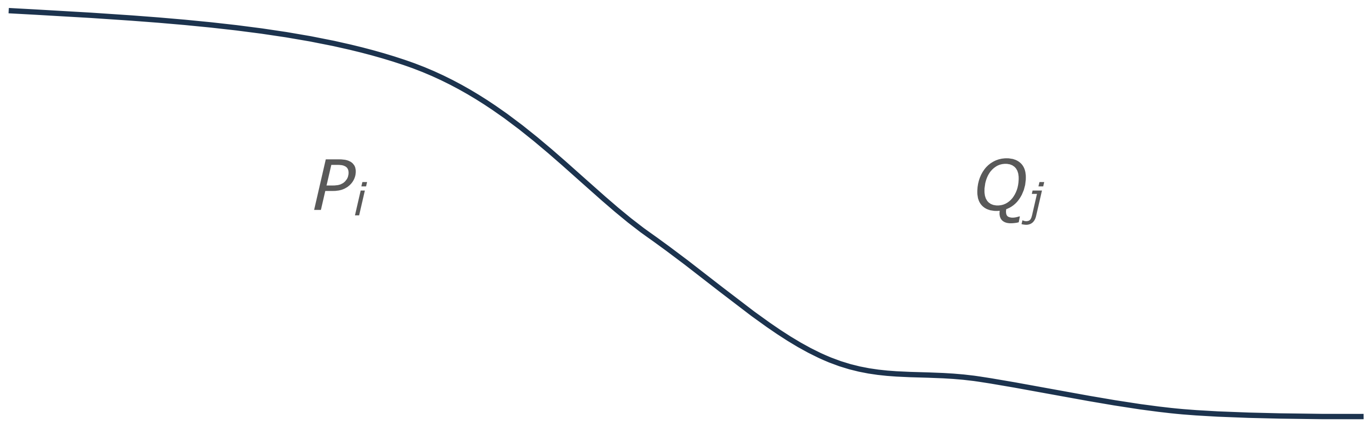

Still, the two definitions are essentially equivalent, because we can include records of previous measurements:

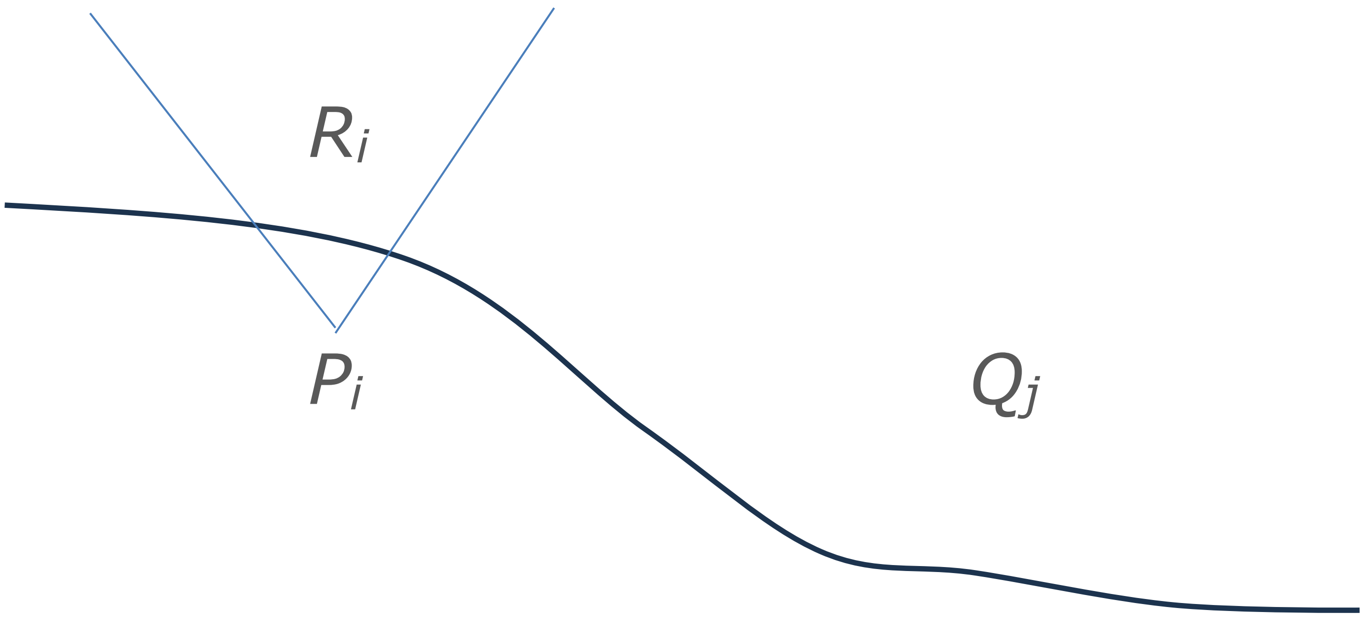

Indeed, the fine-graining

$$\Big\{P_i, R_i Q_j\Big\}_{i,j}$$

is branching also with respect to the total order, and including records of Bob's measurement yields branching wrt the opposite total order.

(N.B. Fine-graining a decoherent set of histories by adding permanent records is standardly used to prove the branching-decoherence theorem.)

We have thus described branching also formally.

We have seen that *how* a causal set branches generally depends on whether or not we choose to introduce a definite time ordering, and which one.

But we have also seen that if we include enough records, *whether* a causal set is branching is independent of whether we define branching wrt to the partial order or wrt to a total order, and is grounded in considerations of decoherence alone.

(C) *Locality and spacetime*

When Alice measures her electron, whether or not Bob’s electron also branches is not a substantive question: it depends purely on considerations of simultaneity, which are relative to choices of frame or foliation.

What remains is the local process of formation of records, which are objectively useful for labelling the branches in a branching pattern.

But even if we consider only branching along the forward lightcone, we still have choices we could make, e.g. whether branches *split* or *diverge*.

(Again, I take it these are different patterns that are equally real, because they are grounded in exactly the same physics.)

We can also choose whether to identify branching in the states of *all* systems in the future lightcone of a branching event, or only in those systems that *actually* interact to form records of the event.

Insofar as we conceptualise spacetime as encoding universal features of the dynamics (we love you, Harvey Brown!), the former allows us to think of branching not as branching of states of the material content of spacetime, but of *spacetime itself*.

This fits also with intuitions that a (relativistic) spacetime encodes possible ‘causal relations’: barring recoherence, different local branches are to the causal future of their branching event, but are causally inaccessible to each other, so they should indeed correspond to separate regions of spacetime.

We thus get a *branching spacetime* structure, and can thus literally thinking of (successors of) Alice and Bob following certain paths in spacetime and meeting up when those paths come together.

One might still worry whether there is some undesirable nonlocal aspect involved in how locally forming branches eventually match up, e.g. in the case of perfect correlations:

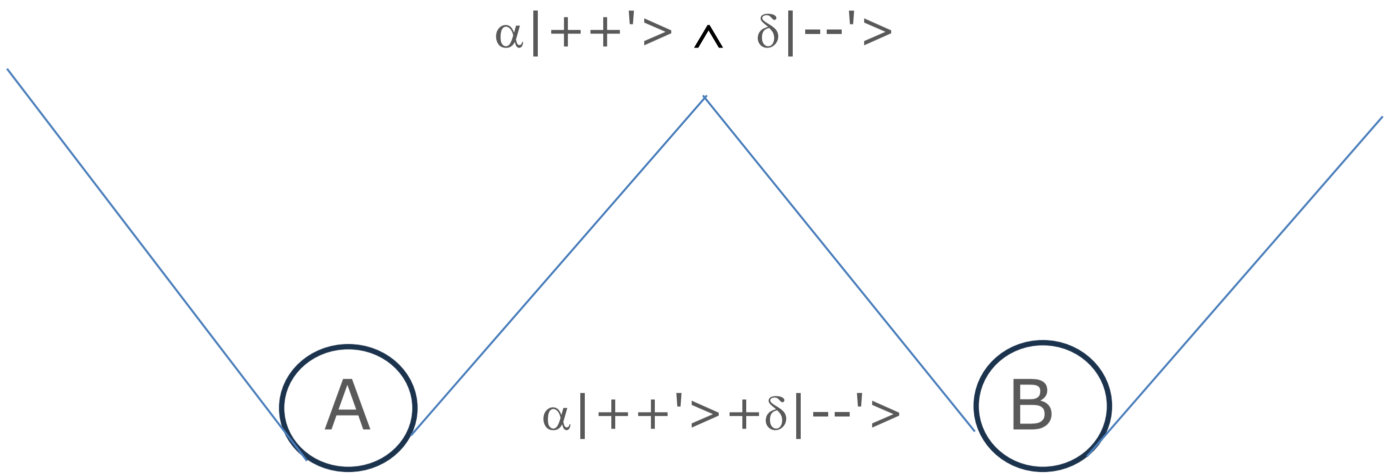

I suggest thinking of *nested locality*: each branching event is determined by the local states in A and B; but how branches match up later is determined by the local branch of the union of A and B:

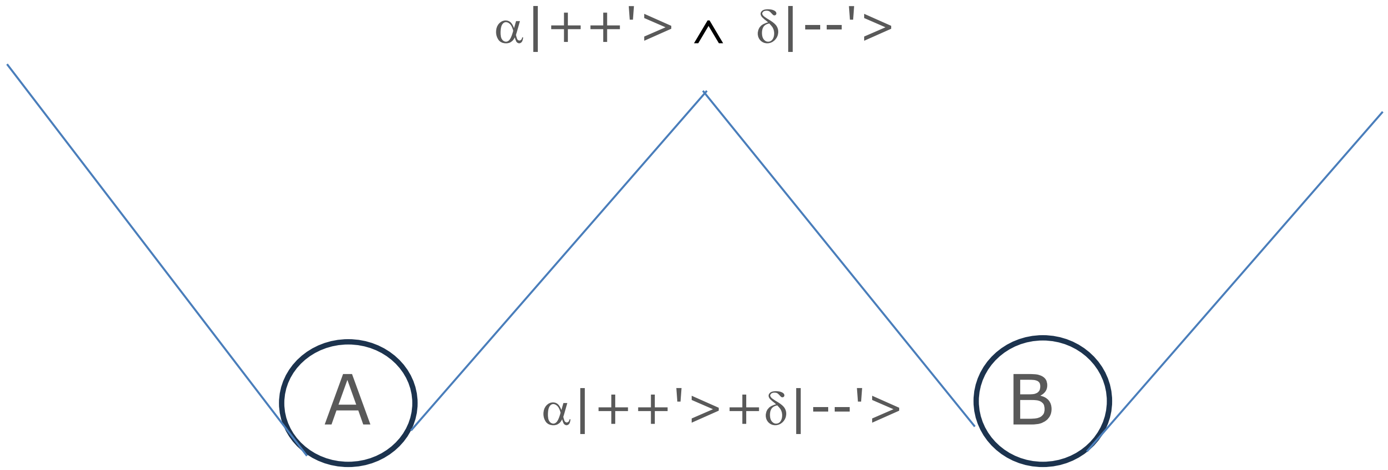

Finally, branching spacetime has the advantage of providing a framework in which to think how spacetime may *emerge* in Everett (GB 2002).

The Minkowski spacetime that we use as background to represent QFT encodes the universal dynamical symmetries of the theory (i.e. the Poincaré group), but it has approximately the same status as configuration space or momentum space in non-relativistic quantum mechanics.

Like configuration space (configurations of what?), the background Minkowski spacetime does not have an obvious interpretation as a space of events (which events?).

But we can identify concrete events as *decoherence* events: literally *nothing* happens unless *branching* happens.

Concrete spacetime is then simply the collection of concrete events.

Crucially, it is a *branching spacetime*, unlike the background spacetime.

It is also *discrete*, but if decoherence scales are sufficiently small, we recover Minkowskian structure on appropriate scales.

[This is analogous to the strategy for recovering spacetime in causal set theory. I am sketching emergence only for Minkowski spacetime, but this may be where causal sets come from in the first place – except that causal set theory does not consider branching.]

Finally, I have talked about emergence of spacetime, but have not said spacetime should emerge from the *universal wavefunction*.

In fact, I believe this picture is neutral between a Schrödinger picture and a Heisenberg picture: decoherence events are not structureless (different values are needed to label different branches), but they may provide a new ‘flash’ ontology for the Everett theory itself!

THANK YOU!