Many Worlds as Anti-Conspiracy Theory: Locally and causally explaining a quantum world without finetuning

Siddharth Muthukrishnan

University of Pittsburgh
siddharth@pitt.edu

April 30 2025

To appear in *Everett and Locality*, editor: Alyssa Ney,
forthcoming with Oxford University Press.

Abstract

Why are quantum correlations so puzzling? A standard answer is that they seem to require either nonlocal influences or conspiratorial coincidences. This suggests that by embracing nonlocal influences we can avoid conspiratorial fine-tuning. But that's not entirely true. Recent work, leveraging the framework of graphical causal models, shows that even with nonlocal influences, a kind of fine-tuning is needed to recover quantum correlations. This fine-tuning arises because the world has to be just so as to disable the use of nonlocal influences to signal, as required by the no-signaling theorem. This places an extra burden on theories that posit nonlocal influences, such as Bohmian mechanics, of explaining why such influences are inaccessible to causal control. I argue that Everettian Quantum Mechanics suffers no such burden. Not only does it not posit nonlocal influences, it operates outside the causal models framework that was presupposed in raising the fine-tuning worry. Specifically, it represents subsystems with density matrices instead of random variables. This allows it to sidestep all the results (including EPR and Bell) that put quantum correlations in tension with causal models. However, this doesn't mean one must abandon causal reasoning altogether in a quantum world. After all, quantum systems can clearly stand in causal relations. When decoherence is rampant and there's no controlled entanglement, Everettian Quantum Mechanics licenses our continued use of standard causal models. When controlled entanglement is present—such as in Bell-type experiments—we can employ recently proposed quantum causal models that are consistent with Everettian Quantum Mechanics. We never need invoke any kind of nonlocal influence or any kind of fine-tuning.

1 Introduction

EPR/Bell correlations[1] (henceforth, simply *Bell correlations*) are puzzling. It seems as if, to explain them, we must invoke nonlocal influences or invoke delicately selected coincidences. Nonlocal influences are unpalatable since they conflict with relativity. Delicate coincidences are unpalatable on broad methodological grounds.

This is a standard way of phrasing the puzzle of quantum correlations, and it presents it as a dilemma: Either admit nonlocal influences (as Bohmians or collapse theorists do) or admit that experimental settings or outcomes are unavoidably fine-tuned (as superdeterminists and retrocausalists do). However, the recent work of Wood and Spekkens (2015) employs the framework of causal models (Spirtes, Glymour, and Scheines 2000; Pearl 2000) to show that this isn't really a dilemma, for we can't avoid the fine-tuning horn of the dilemma by accepting the nonlocality horn. Even if we admit nonlocal influences, some sort of fine-tuning persists. This suggests that the main puzzle posed by quantum correlations is that they require a kind of fine-tuning, no matter what.

How does Everettian Quantum Mechanics (a.k.a. the Many Worlds Interpretation) fit into this dialectic? Everettians argue they can explain Bell correlations without invoking nonlocality.[2] But what about the fine-tuning objection? While much has been written about how the Everett interpretation can avoid nonlocality (or indeed whether it does),[3] to the best of my knowledge nothing has been written about how or whether Everett avoids fine-tuning. This is important to engage with because if Everett falls prey to a fine-tuning objection despite requiring only local influences, then, all else equal, it is not clearly better than superdeterminist and retrocausalist views at explaining Bell correlations, which also preserve locality but admit fine-tuning.[4]

I will argue that Everettian quantum mechanics (EQM) avoids the fine-tuning challenge by rejecting the core principles of the causal modeling framework that lead to the fine-tuning challenge in the first place.

But is that too high a price to pay? The causal modeling framework is a powerful and valuable framework to represent, analyze, and discover causal structure. If EQM demands we jettison it, that would weaken EQM. But I argue

1. Einstein, Podolsky, and Rosen (1935) and Bell (1964).

2. So will QBists and Pragmatists (see Healey (2023)). But, in this essay, I will focus on realist approaches.

3. See, e.g., the papers in this volume.

4. Of course, all else is not equal, and there are other reasons one might favor or disfavor Everett over superdeterminism or retrocausalism.

that this isn't a worry for EQM. For one, I argue EQM does not require us to abandon the classical causal modeling framework in all contexts, but only in contexts with controlled entanglement, such as the contexts that lead to Bell correlations; decoherence licenses the use of classical causal models in most ordinary contexts, thus allowing us to retain that successful and well-tested framework.

For another, I argue when controlled entanglement is present, there is another framework that allows us to represent and analyze causal relations, namely the framework of quantum causal models recently developed by Allen et al. (2017) and Barrett, Lorenz, and Oreshkov (2021), which is compatible with EQM, by virtue of being compatible with pure unitary quantum mechanics.

Taken together, I argue that EQM provides non-fine-tuned explanations of quantum correlations, traffics only in local interactions, licenses the use of the standard causal modeling framework in most classical contexts, and fits well with a quantum causal modeling framework when controlled entanglement is present. This is a constellation of virtues that EQM enjoys. While I don't argue in this paper that rival interpretations of QM don't or can't enjoy the same constellation of virtues, it is hard to see how they will be able to. At any rate, my goal in this paper is primarily to highlight the virtues of EQM with regard to causal explanation and non-fine-tuning. I leave to future work the question of comparing EQM with rival views.

Before we get there, though, I'll need to motivate why we should at all care about how well the causal modeling framework fits with quantum mechanics. So, after briefly introducing the framework of causal models in the next section (Sec. 2), I'll show in the two following sections (Sec. 3 and 4), how the famous arguments of EPR and Bell can be quite naturally phrased as arguments for the empirical invalidity of physically plausible causal models. This will show that applying the framework of causal models to try and explain quantum phenomena isn't new or alien—it is a long tradition among philosophers and physicists going back at least to Einstein. This provides the motivation for us to take seriously the argument from causal modeling for the necessity for fine-tuning quantum mechanics (Sec. 5). Further, given the problems faced by attempts to satisfactorily causally model quantum correlations, we have greater motivation to consider truly quantum causal models.

2 A very brief introduction to graphical causal models

Graphical causal models (Spirtes, Glymour, and Scheines 2000; Pearl 2000) are a powerful framework to represent, analyze, and understand causal relations. They provide a clear mathematical representation of causal relations between variables, support an interventionist semantics (Woodward 2003), and greatly aid in discovery of causal explanations (see, e.g., Malinsky and Danks (2018)). So, if there's a phenomenon (such as quantum correlations) which we are trying to explain causally, then it makes sense to try to represent that phenomenon using a causal model.

The key structure used to represent causal relations in this framework is a *directed acyclic graph*. The variables in the graph represent the different systems or degrees of freedom that may be causally related to each other. An arrow represents a causal influence between the variables. We focus on *probabilistic* causal models, where we associate a probability distribution over the values the variables can take. That is, we take the variables representing systems to be *random variables*. This is a key assumption to which we will return later.

If the vertices of the graph are $V = \{X_1, X_2, X_3, \dots\}$, then we have a joint distribution over the variables $P(X_1 X_2 X_3 \dots)$. From the joint distribution we may obtain marginals and correlations. The structure of causal influences—represented by the arrows—constrains conditional probability relations between variables. Specifically, there are two central constraints: the Causal Markov Condition and Faithfulness.

Causal Markov condition (CMC).—Any variable X in a causal model is independent of all variables not descended from X, conditional on X's parents (i.e., direct causal ancestors).[5] Formally,

$$P(X|\text{parents}(X)\&\text{nondescendants}(X)) = P(X|\text{parents}(X)). \tag{1}$$

Equivalently, the CMC states (*factorization* form):

$$P(X_1, X_2, \dots, X_n) = \prod_i P(X_i|\text{parents}(X_i)). \tag{2}$$

Intuitively, the CMC states that causal influences on a given variable always flow through its immediate parents.

5. This is the *screening off* formulation.

Why believe the CMC? It seems core to our notion of "cause". Several arguments support it: (i) **Generalization of Common-Cause Principle.** It generalizes, and hence derives plausibility from, Reichenbach's Principle (Reichenbach 1956). (ii) **Determinisn + Noise.** Pearl (2000) argues CMC follows if causal relations between variables arise from deterministic relations plus noise. (iii) **Manipulability.** Hausman and Woodward (1999, 2004) argue causes being useful to *manipulate* effects supports the CMC. (iv) **Track record.** It's central to the successful framework pioneered by Spirtes, Glymour, and Scheines (2000) and Pearl (2000). (v) **Physical features.** But most pertinent to us, one can motivate the CMC by appeal to specific physical features, such as spatial separation and connection to random sources—what Weinberger, Williams, and Woodward (2024) call "worldly infrastructure".

Let's turn now to the other central constraint that the causal graph places on the distribution over variables.

Faithfulness.—Faithfulness is the 'dual' of the CMC. The CMC says that there are no more *correlations* than expected from the causal diagram; faithfulness says there are no more *independences*. E.g., faithfulness applied to Fig. 1 says X and Y are not independent because they have an arrow between them. Thus, if the distribution over XYZ is to be faithful, then X and Y cannot be independent, conditional or otherwise.

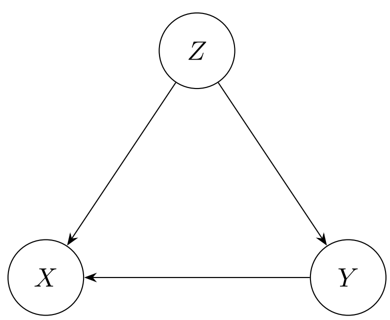

Figure 1: A causal model that is unfaithful to—i.e., fine-tuned for—a distribution in which X is probabilistically independent of Y.

Distributions that respect CMC but violate faithfulness must be *fine-tuned*. Consider the parameters governing the relation between variables (i.e., between the values of variables or between probabilities over those values). Faithfulness-violating parameter settings will need to be very specially selected.[6], and small

6. See Spirtes, Glymour, and Scheines (2000, pp. 41-42) and Weinberger (2018).

deviations away from these settings will almost always restore faithfulness.

In sum, the CMC and faithfulness demand that the distribution over variables *respect the structure of the causal diagram*, containing all and only those correlations implied by the causal graph.

We will now apply the graphical causal models framework to quantum correlations. We will see, in the following two sections, that Einstein's and Bell's arguments rule out certain classes of causal models as explanations of quantum correlations. This will motivate applying the framework to quantum correlations. We can then turn (in Sec. 5) to the result of Wood and Spekkens (2015), and see how attempts to explain quantum correlations using this framework invariably violate the faithfulness condition.

3 Einstein/EPR and the Causal Markov Condition

In 1935, EPR argued that quantum mechanics (QM) is incomplete (Einstein, Podolsky, and Rosen 1935). The essential argument was developed by Einstein in 1927 (see Howard (1985) and Harrigan and Spekkens (2010)) and can be presented as follows. Suppose we have two spin-$\frac{1}{2}$ particles in the entangled singlet state, i.e., $\frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)$. Measuring the first particle's spin and seeing that it comes "↑" (or "↓"), I immediately know that a measurement of the second particle will yield "↓" (or "↑"). This will be true even if the particles are far enough from each other that no light-speed signals can get from one particle to the other within the time-frame of the measurements.

But just *look* at our state representation $\frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)$: we see nothing which tells whether a measurement of either particle will yield "↑" or "↓"; the representation is symmetric between those two outcomes. In measuring, of what feature of the world have we acquired knowledge? If the answer is what the state of the second particle was all along, then we have conceded that the singlet state does not successfully represent all relevant physical features. Consequently, quantum mechanics's representation of the world would be incomplete.

Could measuring the first particle *cause* the second to change from being in a symmetric state to a definite state? That would imply superluminal—indeed, instantaneous—influences between the particles, since the *immediacy* of our knowledge of the other particle remains unaffected by distance. Disallowing those, it seems our measurement of the first particle must reveal a pre-existing property of the second particle. But this property isn't represented by the quantum formalism. Consequently, quantum mechanics must be incomplete. This was the essence of Einstein's argument.

Now let us phrase this in causal modeling terms.[7] We can think of Einstein as arguing that a certain kind of causal model is insufficient for explaining quantum phenomena.

Consider the causal graph represented in Fig. 2. In this figure, A and B represent the possible measurement outcomes on the two particles, taking values in $\{\uparrow, \downarrow\}$. S and T represent measurement settings. Above, we considered only one possible measurement setting; hence, S and T are fixed at "measure z-spin". λ_A and λ_B are local variables that help fix, perhaps only probabilistically, the values that A and B take; these variables can be seen as encoding whatever local property the quantum state attributes to the particles.

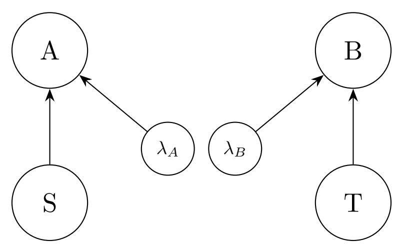

Figure 2: Einstein's argument rules out the possibility of this kind of causal model reproducing the predictions of quantum mechanics.

Applying the screening-off formulation of the CMC [Eq. (1)] to this graph, we get the independences:

(i) A is independent of B, λ_B, and T conditional on S and λ_A; and B is independent of A, λ_A, and S conditional on T and λ_B.

(ii) S, T, λ_A, and λ_B are all independent of each other.

We can see the key premises of Einstein's *reductio* as motivations for the CMC applied to this graph. The first key premise of Einstein's argument is that the two wings of the experiment cannot be causally connected since the two wings can be arbitrarily separated, and no signal can travel between them during

7. Phrasing EPR-style arguments in causal modeling terms isn't new; see, e.g., Van Fraassen (1982) and Hausman (1999) for classic treatments, and see, e.g., Suárez and San Pedro (2010) and Näger (2016) for more recent treatments. While I make no great claim to originality, my approach isn't exactly the same as other authors'.

the experiments. This motivates the structure of the causal graph (lack of arrows connecting wings) and the CMC applied to it (the conditional independences specified in statement (i)), because the lack of causal influence between the two wings suggests that whatever happens on one side should be sufficient to causally explain what happens there.

Einstein's second key premise is that the entangled quantum state is indifferent between the two outcomes and representationally complete. To respect this premise, whatever determines outcomes on either wing should be indifferent between $\uparrow$ and $\downarrow$; so, λ_A and λ_B should assign equal probability to those outcomes. And since λ_A and λ_B are causally disconnected, they must be uncorrelated, which is what the CMC requires.[8]

Formally, applying the CMC (as given in Eq. (2)), we get

$$P(ABST\lambda_A\lambda_B) = P(A|S\lambda_A)P(S)P(\lambda_A)P(B|T\lambda_B)P(T)P(\lambda_B). \tag{3}$$

Summing over S, T, λ_A, and λ_B, we get:

$$P(AB) = P(A)P(B), \tag{4}$$

i.e., A and B are uncorrelated. This contradicts quantum experiments.

Faced with this contradiction, it seems one of Einstein's two key premises must be rejected. Given relativity's success, Einstein/EPR think that it's untenable to reject the premise that the two wings are causally disconnected. Thus, they reject the representational completeness of the quantum state.

4 Bell and the Causal Markov Condition

If QM is local but incomplete, then perhaps we can come up with a local *completion* of QM—one containing elements representing those properties of systems left out by QM while still trafficking entirely in local influences. These elements are usually called *hidden variables*.

The language of causal models allows us to frame this precisely. We saw that Einstein argued against the causal model in Fig. 2 with the CMC applied to it. There, the lack of correlation between λ_A and λ_B was justified by appeal to (a) the spatial separation between the wings and (b) the quantum state's indifference between the two outcomes. However, one might think that because the particles

8. Independence of λ_A and λ_B with S and T is trivial since the latter are single-valued.

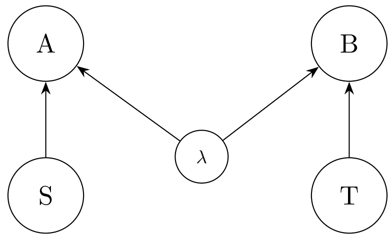

Figure 3: The violation of the Bell inequality rules out the possibility of this kind of causal model reproducing the predictions of quantum mechanics.

started out together (perhaps they are electrons from the same atom) we should not assume that the two wings have entirely causally disconnected variables. This leads us naturally to consider the causal diagram of Fig. 3, and ask whether it can explain quantum correlations.

Bell's (Bell 1964; Bell 1975) famous results answer this in the negative. Bell proved an inequality that the measurement statistics of *any* theory conforming to Fig. 3 must satisfy. These inequalities are *violated* by the predictions of QM. These violations have been repeatedly empirically verified by exceptionally careful experiments.[9] Thus, Einstein's hope of a local completion of QM is a dead end.

We will now see how the core assumption that Bell made in proving his theorem is closely connected to the CMC applied to Fig. 3. Much like the previous case, the CMC will be justified here by the physics of the experimental situation.

Let's unpack the causal graph of Fig. 3. S and T are random variables representing measurement settings; A and B are random variables representing measurement outcomes; and λ represents whatever common element that is shared between the two wings and which may influence the measurement outcomes. Standard presentations of Bell's result take S, T, A, and B as binary-valued. Note also that λ need not only encode whatever properties the quantum state represents the system as having. After all, this model is being considered after accepting Einstein's argument that quantum mechanics is incomplete.

The key premise in Bell's theorem is what is often called *factorizability*:[10]

$$\Pr(AB|ST\lambda) = \Pr(A|S\lambda)\Pr(B|T\lambda). \tag{5}$$

9. See, e.g., Hensen et al. (2015), Giustina et al. (2013), and Shalm et al. (2015).
10. See, e.g., Myrvold, Genovese, and Shimony (2021, Sec. 3.1) for more on this.

From this, it is straightforward to derive a version of the Bell inequality.[11]

Factorizability follows from the CMC applied to Fig. 3. To see this, first note Eq. (1) entails that S, T, and λ are mutually independent. Therefore,

$$\Pr(ST\lambda) = \Pr(S)\Pr(T)\Pr(\lambda). \tag{6}$$

Now, from Eq. (2), we have:

$$\Pr(ABST\lambda) = \Pr(A|S\lambda)\Pr(B|T\lambda)\Pr(S)\Pr(T)\Pr(\lambda). \tag{7}$$

Writing the LHS as $\Pr(AB|ST\lambda)\Pr(ST\lambda)$, then canceling $\Pr(S)\Pr(T)\Pr(\lambda)$ on both sides by using Eq. (6), we get the factorizability condition, Eq. (5).

As in the EPR-style case, we can justify the CMC here by appealing to physical facts. The CMC here amounts to the following independences, obtained from Eq. (1): $A \perp\!\!\!\perp BT|S\lambda$, $B \perp\!\!\!\perp AS|T\lambda$, $S \perp\!\!\!\perp BT\lambda$, $T \perp\!\!\!\perp AS\lambda$, and $\lambda \perp\!\!\!\perp ST$.[12]

The first two independences (i.e., $A \perp\!\!\!\perp BT|S\lambda$, $B \perp\!\!\!\perp AS|T\lambda$) can be justified by the lack of direct influences between the two wings. In Fig. 3, this is reflected in the absence of causal pathways between the left wing and the right wing. This is enforced by spacelike separation of the measurements. This leads us to expect that the settings and outcomes in one wing are independent from the settings and outcomes on the other, *except* for those correlations accounted for by the shared features (encoded in λ) that might arise from past interactions between the systems.

The other independences ($S \perp\!\!\!\perp BT\lambda$, $T \perp\!\!\!\perp AS\lambda$, and $\lambda \perp\!\!\!\perp ST$) are justified by the measurement settings on each wing being determined by processes disconnected from the other and from the history and physics of the measured particles. In a recent experiment, this was enforced by choosing the measurement settings based on photons arriving from parts of the universe causally disconnected for billions of years (Rauch et al. 2018), ensuring not only that the two measurement settings are entirely uncorrelated with each other but also uncorrelated with any system variable (such as λ). Of course, we needn't go to such lengths to satisfy the independence condition. Choosing measurement settings randomly might suffice, say by coin tosses. As might whatever processes—biological, mental, personal—that drive experimenters' choices. For, no plausible physics links such

11. See, e.g., Brunner et al. (2014, p. 3) for details.

12. Notation: $\alpha \perp\!\!\!\perp \beta\gamma\delta\ldots|\chi\psi\omega\ldots$ means that α is independent of any subset of $\{\beta, \gamma, \delta, \ldots\}$ conditional on *all* of $\chi\psi\omega\ldots$. If there's no third entry, then it's an unconditional independence between α and any subset of $\{\beta, \gamma, \delta, \ldots\}$.

processes with the system under study. Indeed, it is hard to articulate what physics could lead to a violation of these independences outside of implausible stories involving conspiracies (Shimony, Horne, and Clauser 1976). (We will see in the next section what theories violating these independences will look like within the framework of causal models.)

Thus, the experimental violation of Bell inequalities amounts to a crisis for causal explanations of quantum phenomena. The CMC applied to Fig. 3 seems to encode locality and independence of measurement settings from the particles, enforced in our experiments by known physics. Thus, it seems as though if we want a causal explanation of Bell-inequality violations, we must abandon one of these assumptions, even if it conflicts with physics we think we know. This is the standard articulation of the puzzle of quantum correlations, which is as a dilemma: Either abandon locality or accept that there are conspiratorial coincidences.

As we will see next, however, this is a false dilemma: even if we abandon locality, our explanations of quantum correlations may still require some conspiratorial fine-tuning, threatening the possibility of satisfactory causal explanations.

5 Wood and Spekkens and Faithfulness

The violation of Bell inequalities suggests that to causally explain quantum correlations, we must edit the causal model of Fig. 3, even if such edits require jettisoning some cherished physical principle.

How might we make such edits? Suppose, as the Bohmian and spontaneous collapse theories do, we allow for superluminal influences.[13] This allows arrows from (say) the left wing to the right wing, as in Fig. 4. Such arrows clearly violate the empirically well-supported principle that physical influences cannot be superluminal. However, some have concluded that empirical Bell-inequality violations falsify this principle.[14]

A different editing strategy is that of superdeterminists, who embrace the idea that measurement settings are not truly freely chosen—i.e., the settings induce the relevant correlations between outcomes despite every effort to ensure free selection of settings. This leads to a model like Fig. 5, with influences from λ to A, B, S, and T.

13. See, e.g., Maudlin (2019) for an overview of these theories.
14. See, e.g., Albert and Galchen (2009) and Maudlin (2014).

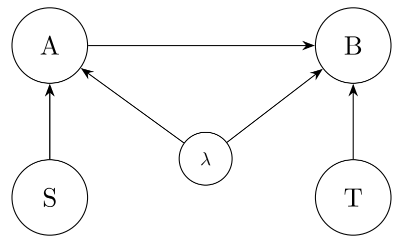

Figure 4: A causal diagram allowing for superluminal influences. These kinds of causal models have to be fine-tuned (i.e., violate faithfulness) if they are to preserve the no-signaling criterion.

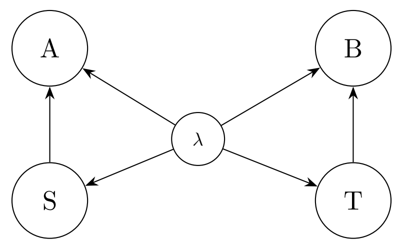

Figure 5: A causal diagram with superdeterminism, which prevents free choice of measurement settings. These kinds of causal models have to be fine-tuned (violate faithfulness) to preserve no-signaling.

Such edits, however, don't provide a fully satisfactory causal explanation of quantum correlations because they're *fine-tuned*—they violate *faithfulness* (Sec. 2). Models allowing nonlocal influences (Fig. 4) must reproduce the no-signaling condition: measurement outcomes on one side must be probabilistically independent of measurement settings on the other. But given superluminal influences, the physics must be fine-tuned so as to prohibit using these influences to signal. This fine-tuning is why faithfulness is violated in such models.

Specifically, from Fig. 4, we see that the probabilistic independence of the measurement settings on one wing and outcomes on the other wing will be unfaithful. There is a causal pathway from S to B, via A. If the distribution were faithful to this graph, then S and B would not be independent. But no-signaling says we must have this independence; otherwise, we signal by manipulating S to influence B. Such an independence is also empirically supported. Hence, models with nonlocal influences must be unfaithful.

Turning now to superdeterminist theories, which require fine-tuning because microphysics would have to be in an extraordinarily specific state so that just the right certain measurement settings occur in all the various experiments and data. This is reflected in the failure of faithfulness for superdeterministic causal models. Consider Fig. 5. In experiments, S is certainly independent of T, but this is not an independence delivered by the CMC, revealing a faithfulness violation.

To be clear, I'm not saying superdeterministic theories are on a par with theories with nonlocal influences. Bohmian mechanics and spontaneous collapse theories are much more precisely formulated than extant superdeterministic proposals, and so enjoy theoretical virtues that are independent of fine-tuning. Even on the count of fine-tuning, Bohmian mechanics fares better than a generic superdeterministic theory, because Bohmians can argue that the Born-rule distribution, though needing to be specially selected, is natural and straightforwardly specifiable within QM.[15] Meanwhile, it is unclear if any superdeterministic theory can provide a simple, physically plausible specification of the joint distribution over variables.

In sum, any empirically adequate causal model of quantum correlations—whether it violates relativistic locality or independence of systems and settings—must be fine-tuned. Could more sophisticated models avoid fine-tuning? No. Wood and Spekkens (2015) prove that no faithful classical causal model can reproduce quantum correlations. Even if we allow novel physics, an uncompromising causal explanation of quantum correlations remains out of reach.

6 Where do we go from here?

The arguments above dim the prospects for causally explaining quantum correlations. Bell correlations resist being cast into the well-motivated framework of causal models.

What then? Some take EPR and Bell as undermining the CMC,[16] while others see them as evidence for the inapplicability of the CMC in quantum contexts.[17] Such responses are plausibly in line with pragmatist/anti-realist views of quantum mechanics.[18] For, if one does not take the quantum state as representing the physical state, then the physical arguments for adopting the CMC in the EPR/Bell-type cases are far less persuasive.

15. See, e.g., Goldstein and Struyve (2007).
16. See, e.g., Van Fraassen (1982), who takes Bell as refuting Reichenbach's principle of common cause.
17. See, e.g., Hausman and Woodward (1999).
18. As suggested by Van Fraassen, but not Hausman or Woodward.

Another response, popular among realists about QM, is to retain the CMC and explain away violations of faithfulness as unproblematic on their preferred view of QM (such as Bohmian mechanics, collapse, or retrocausality).[19]

Without evaluating the overall persuasiveness of these realist replies, let me highlight how unusually strong faithfulness violations in causal models of quantum phenomena have to be when compared to faithfulness violations in causal models of more ordinary phenomena.[20] The temperature of a room with a well-functioning thermostat and a heating system is usually independent of outside temperature, because the thermostat detects when the room temperature is starting to drop or rise and the heating system compensates with hotter or colder air. This would be a faithfulness violation in some appropriate causal model because the outer and inner temperatures are independent despite the causal processes of heat exchange between the inside and outside. However, this independence will never be perfect or exceptionless. Imperfect because at a sufficiently fine grain (in precision, time, space) the room's temperature will be weakly correlated with the outside. Not exceptionless because at least sometimes (a cold-snap, a faulty thermostat) the independence will break. However, quantum no-signaling appears perfect and exceptionless. Consequently, any attempted explanation of faithfulness violations in quantum mechanics carries a far heavier burden than what explanations of faithfulness violations in ordinary contexts carry.

I offer a different response, one with the following commitments.

1. It is realist in its view of quantum mechanics.

2. It rejects the blanket applicability of the standard framework of causal modeling in the context of quantum mechanics (unlike the views that retain the framework but excuse faithfulness).

3. It explains why the apparatus of causal modeling works well in most classical contexts, failing only in certain quantum contexts.

4. It offers a way to *quantum* causally model quantum experiments.

The view I'm proposing is the combination of Everettian quantum mechanics (EQM) and a framework for quantum causal modeling, developed by Allen et al. (2017). This view is realist because EQM is realist. It rejects the applicability

19. See, e.g., Egg and Esfeld (2014), Näger (2016), and Evans (2021).
20. See, e.g., Andersen (2013) for instances of the latter.

of the standard framework of causal models to quantum phenomena because it represents subsystems using density matrices instead of random variables. It explains the usual validity of classical causal modeling in a quantum world via appeal to decoherence: widespread decoherence licenses the use of random variables. However, in the presence of controlled entanglement, this assumption fails, rendering standard causal modeling invalid. In such cases Allen et al. (2017) offers a causal modeling framework accommodating the way QM represents subsystems—namely, via density matrices. Because this framework is compatible with unitary quantum mechanics, it is also compatible with EQM. This yields a satisfying way to causally model quantum experiments, avoiding the pitfalls faced by previous attempts.

7 The Everettian explanation of Bell violations

Everettian quantum mechanics (EQM) treats all systems (measurement devices, humans) as constituted by quantum mechanical parts. Thus, they can all enter into superpositions. When a quantum system in superposition interacts with another quantum system with many uncontrolled degrees of freedom (e.g., a measurement device or the environment), the latter enters into a superposition too, governed by the Schrödinger equation. Within the different branches in this large superposition, there emerge multiple quasiclassical *worlds*, stable states of affairs exhibiting approximately classical behavior.[21]

Thus, measurements result in multiple branches; each branch corresponding to a definite outcome. Branches come with a *branch weight*, their mod-squared amplitudes. Branch weights supply a measure over branches, determining the probabilities of measurement outcomes.[22]

Branching which grounds probabilistically distributed stable outcomes only obtains when the system is *decohered*. Only then may we describe the system using random variables. That we can describe systems using random variables is a central assumption of the causal modeling framework, but this is not uniformly permitted by quantum physics. (This point is technically independent of EQM. However, it fits neatly with EQM, for decoherence is crucial within EQM for explaining the emergence of randomness and stability.)

21. See (Wallace 2012) for a detailed development and defense.

22. How this measure connects to the observed probabilities is the most contentious question concerning EQM (see, e.g., some of the papers in (Saunders et al. 2010) for a good discussion). Here, I assume *some* account answers this question. Needless to say, if EQM cannot explain observed probabilities, then it's not viable.

This observation by itself isn't enough to tell us that the causal modeling framework will fail in explaining quantum correlations. After all, in EPR/Bell-type experiments, we have measurement statistics, and so there is enough decoherence to license the use of *some* random variables. That is, post-measurement, in each wing of a Bell-type experiment, we can associate to each system a decohered density matrix that is almost exactly diagonal in the measurement basis, with diagonal entries equal to the Born probabilities of various possible outcomes.

The tension between EQM's picture and the causal-models' picture arises because the latter picture also assumes that the *joint* distribution over all the variables is always well-defined. We can see this in the way in which constraints like the CMC and faithfulness are formulated: they constrain the joint probability distribution over all the variables based on the structure of the causal graph. However, in a quantum world, because a random variable description is only licensed by a decohered density matrix, we can't always assume a global decohered density matrix is well-defined that licenses describing the systems with a joint probability distribution.

The upshot of this in the context of Bell-type experiments is that even though we have local decohered structures, these structures don't *immediately* ramify up to a global one because there are no instantaneous interactions. Thus, there won't be a globally decohered density matrix simply because we have local decohered density matrices. Consequently, we won't have the license to use random variables to represent the *joint* state of the two wings. This means that we cannot talk about the correlations between the two wings of experiments until and unless the two wings interact in a way that establishes a globally decohered structure—which typically happens when the future lightcones of the two experiments intersect. In such a structure, the overwhelming branch weight will be on branches with correlations that violate the Bell inequality. Consequently, observers in such branches will be unable to explain their observations with a causal model because of the arguments considered in the previous sections. Thus, on an Everettian picture, the violation of the Bell inequality signals the absence of a globally decohered structure at earlier times, not nonlocality or finetuning.

We typically get failures of classical causal modeling when the system we are modeling contains entanglement that is *accessible* or *controllable*. To see this, suppose we are building a causal model of some situation that doesn't require the use of quantum theory. Say the relevant variables are the amount of precipitation in a forest and the population of beavers in that forest. While the quantum degrees of freedom *constituting* or *grounding* these variables are entangled with each other in all sorts of complex ways, that entanglement is not accessible or controllable. This then means we can continue using the standard framework of classical causal models, because we do have a globally decohered structure along with locally decohered structures. This is not the case for systems whose entanglement is maintained carefully, such as in Bell-type experiments. Outside of such cases, we may safely employ the classical causal modeling framework.

Let me consider an objection. One might deny that only a globally decohered structure licenses the use of random variables. Random variables are cheap; you can define one whenever you want. As soon as the measurements are performed in the two wings, it is determined that the EPR/Bell correlations will be observed in the future. Consequently, we could define the correlations as obtaining even before the future lightcones intersect. Yes, the grounds for these correlations are in the future, but that doesn't logically stop us from defining a joint distribution now.

But this objection doesn't affect the main thrust of my argument. These kinds of "cheaply defined" random variables cannot be taken as *representing* the physical goings-on. My key point is that the causal models framework assumes that the physics of systems can be *represented* by random variables, and that is the assumption violated in Bell-type experiments.

To better understand the importance of random variables *representing* the physical goings-on for successful causal explanation, let us go back to EPR. EPR's argument posed a problem for causal explanation only because we thought that a measurement in one wing *immediately* gave us knowledge of a measurement outcome on the other wing, suggesting the need to posit nonlocal influences or hidden common causes. That is, if we had a random variable that encoded these correlations, then such a random variable creates a problem for a causal explanation only if we think of it as representing the physical goings-on at the time of measurements. Otherwise, there'd be no puzzle.

We can see this point more clearly by trying to build a causal model for the Everettian story of what is going on in Bell-type experiments. For instance, one might think that a causal model of the sort in Fig. 6 might be made to work. Here, the variable C is assumed to represent the values that the correlations between A and B will take. However, this isn't empirically adequate for the same reasons that the causal model of Fig. 3 isn't empirically adequate: within the standard causal modeling framework, the Bell-violating correlations between A and B are established before we get to C, and so we would need nonlocal influences or finetuning.

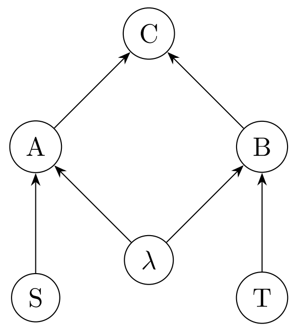

Figure 6: A causal diagram that one *might* think captures the Everettian story of Bell correlations. However, this can't work because it still assumes that correlations between A and B are always well-defined.

So far we have a picture that satisfies commitments 1, 2, and 3. The Everettian story about Bell correlations is realist, it rejects the blanket applicability of classical causal modeling since it only licenses the use of random variables in the presence of appropriately decohered structures, and it explains why classical causal modeling usually works well: because of the widespread presence of decoherence. So that leaves the question of how we might satisfy commitment 4, and we turn to that question next.

8 Quantum causality and Everett

So, can we causally explain quantum correlations? If what we mean by "causally explain" is captured by the standard causal modeling framework discussed in Secs. 2-5, then the answer is no—at least not entirely satisfactorily. Further, if we are Everettian, we saw that the causal modeling framework is the wrong kind of tool to model Bell-type experiments.

Should we abandon hope of causal talk in Bell-type or other quantum contexts? Perhaps we should only provide physical descriptions, and eschew causal descriptions. Just run the physics forward, so to speak. No. It's clear that quantum systems can stand in causal relations.[23] One can clearly cause changes in quantum states of systems and quantum systems can clearly be causes. This is explicit in quantum experiments, where we intervene on quantum systems so as to make a difference. For example, in a Bell-type experiment,

23. Shrapnel (2014) makes a related point.

we can choose to perform different measurements on either wing, or we can choose to destroy the system or isolate the system. These interventions will cause differences downstream in the statistics when we measure correlations. Such causal structure, available for manipulation, becomes particularly salient when considering implementations of quantum information theoretic protocols or quantum computational algorithms. It won't do to render such clearly available and frequently employed causal handles on quantum experiments unintelligible or unmodellable.

This is why the Everettian should seek some framework of causal modeling in which the relata of causal relations are density matrices instead of random variables. If such a framework were available, perhaps one analogous to Spirtes, Glymour, and Scheines (2000) and Pearl (2000), then that would be a strong start towards causally explaining quantum mechanical phenomena as well.

I will now argue that at least one recently developed framework, due to Allen et al. (2017), of *quantum* causal models fits the bill. It is a framework for causal reasoning about quantum phenomena that takes as its starting point the idea that we should represent systems with reduced density matrices. Even though it was developed independently of the Everettian worldview, it fits well with it because it is compatible with unitary quantum mechanics, and Everettian quantum mechanics is, at its core, just unitary quantum mechanics. Thus, once we have this framework in our toolkit, we have a story about Bell correlations that simultaneously satisfies commitments 1-4. Furthermore, given that this framework was developed with unitary compatibility as a core assumption, it is unclear how realist rivals to Everett can quite so easily help themselves to such a framework.

Let me emphasize that Allen et al. (2017) do not take their framework to be Everettian. They use the "church of the larger Hilbert space"[24] as a tool to derive their framework much like how Pearl (2000) uses the deterministic structural equations framework as a tool to derive his framework of probabilistic causal models. But the unitary-compatibility of the Allen et al. (2017) framework allows it to be interpreted in an Everettian way.

The best way to introduce this particular framework of quantum causal models for our purposes is by analogy with the deterministic structural equation modeling (SEM) framework and its relation with probabilistic classical causal

24. This is the view that we can always see impure states and non-unitary dynamics as resulting from subsystem states and dynamics in a larger Hilbert space, where the states are pure and the dynamics unitary.

models (see, e.g., Hitchcock 2020). Imagine *deterministic* functions (structural equations) govern the behavior of variables in a classical causal model. These relations are arranged in a directed acyclic graph (DAG). For a given SEM graph, suppose we consider a proper subgraph (which is also a DAG). Then we can get probabilities on the variables of the subgraph by averaging over those variables that aren't in this subgraph. This is similar to how we can see probabilistic behavior in statistical mechanical systems at an emergent level even if the underlying dynamics is deterministic. Mathematically, if we have a deterministic function that is of the form $Y = f(\lambda, X)$, then we can get a distribution $P(Y|X) = \sum_{\lambda} P(\lambda) f(\lambda, X)$, where $P(\lambda)$ is some distribution (which could arise due to a deterministic process) over the "hidden variable" λ. In Fig. 7, I show in diagrammatic form a deterministic SEM model and the probabilistic causal model derived from it. (The representation I employ here of the SEM and the probabilistic model derived from it is really the dual of their representation as a DAG in the sense that the lines represent variables and the boxes represent the transformations. In what follows, including in the quantum case, I'm going to stick to this dual representation, and draw boxes to represent the transformations on the systems and use wires to denote inputs to the transformations. Doing so unfortunately breaks the visual analogy with the DAGs, but allows for greater conceptual clarity.)

Now, let's turn to the quantum version of such an analysis. Here, the analogue of the deterministic functions of variables is unitary transformations on a Hilbert space. The corresponding analogue of probabilistic causal models is obtained by tracing over the subsystems of the Hilbert space which we are discarding. A map obtained by taking a unitary map and discarding its action on auxiliary subsystems is called a *quantum channel*, which, mathematically is represented by so-called *completely positive* maps. These maps take density matrices—which are obtained by tracing over subsystems of a quantum wavefunction—to other density matrices.

Quantum channels can also be thought of as the quantum generalization of the classical channels of Shannon's information theory. Classical channels are represented by conditional probability functions: a channel $X \to Y$ is associated to the distribution $P(Y|X)$, which represents the probability of receiving a message $Y = y$ given that the message $X = x$ was sent. Thus, while in the case of classical causal diagrams, the causal links between variables are associated with conditional probabilities, in the quantum case, the causal links between subsystems are associated with quantum channels $\mathcal{E}_{X\to Y}$.[25]

Thus, in analogy with the classical causal modeling framework, the inputs are subsystems of the global Hilbert space and the transformations are quantum channels, which are to be thought of as being obtained by tracing over auxiliary quantum subsystems of the Hilbert space. This is depicted in Fig. 8, where a quantum channel is derived from a unitary transformation by tracing over auxiliary degrees of freedom. Allen et al. (2017) show how we can obtain a picture of DAGs from the starting points in the brief sketch presented here. They also present a generalization of the causal Markov condition. I refer the reader to Allen et al. (2017) for more details.

Figure 7: A deterministic structural equation model (SEM) and a classical probabilistic causal model obtained from it. (a) A very simple deterministic structural equation: $Y = f(\lambda, X)$. (b) A very simple classical channel obtained from this deterministic SEM by averaging over λ.

How should we model interventions in quantum causal models? In the context of classical causal models, the idea of an intervention is modeled, mathematically, using the *do*-calculus. We intervene on a particular variable X by setting the variable equal to a particular value: written $do(X = x)$. This breaks the arrows between X and its parents. We can then look at how intervening on X changes its descendants. In the quantum mechanical case, how do we represent the *do* operation? By a *quantum instrument*: a collection of completely positive maps

25. Note that Allen et al. (2017) employ a particular representation of quantum channels, called the Choi-Jamiolkowski representation. On this representation, we can associate a density matrix $\rho_{A|B}$ with the channel itself, which makes the analogy with conditional probability functions sharper. I do not present it this way here to avoid complicating the presentation.

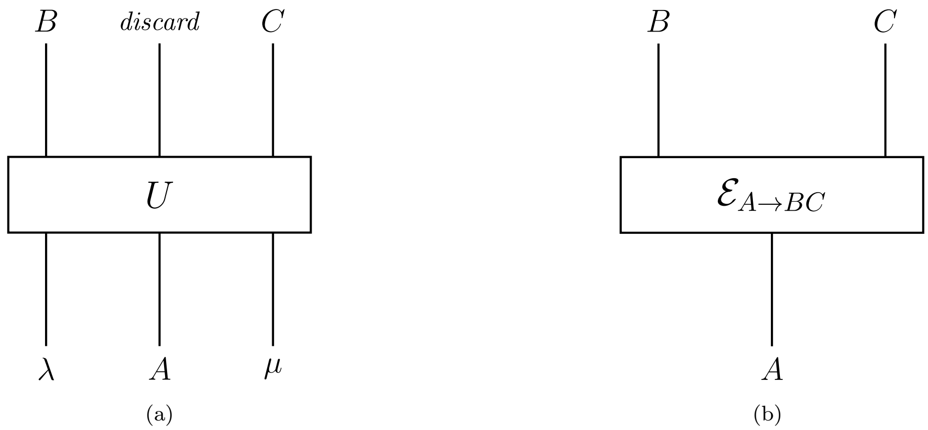

Figure 8: A deterministic quantum map and a quantum causal model obtained from it. (a) A very simple deterministic quantum transformation: $|\psi_{\text{out}}\rangle = U(|\psi_{\lambda A\mu}\rangle$. (b) A very simple quantum channel obtained from this deterministic quantum map by tracing over an auxiliary system [depicted in (a) by *discard*]: $\rho_{BC} \equiv \text{Tr}_{\text{aux}} |\psi_{\text{out}}\rangle \langle\psi_{\text{out}}| \equiv \mathcal{E}_{A\to BC}(\rho_A)$, where $\rho_A = \text{Tr}_{\lambda\mu}(|\psi_{\lambda A\mu}\rangle \langle\psi_{\lambda A\mu}|)$.

$\mathcal{E}_k$, where k is a classically readable result of applying the quantum instrument on the system. Thus, if the system was in state ρ before the interaction with the quantum instrument and the quantum instrument reads out k, then the quantum state of the system is now $\mathcal{E}_k(\rho)/\text{Tr}(\mathcal{E}_k(\rho))$. The denominator represents the probability that the instrument will record k.

Why quantum instruments? Interventions fix the system in a particular definite state. In classical physics, what one can do is, in effect, *erase* the previous state of the system and *rewrite* it into a new state. This is not possible in general for quantum systems, where transformations have to be *linear*. (If a linear map takes all elements of a vector space to a single vector, then that map has to be the zero map.) The closest we can come to deliberately putting the system in a desired state is operate on the system with a kind of measuring device, which, with some probability, puts the system into one of many different possible states. This operation is what is mathematically modeled using the framework of quantum instruments I just sketched.

With all this set up, representing a Bell-violation experiment is really quite straightforward; see Fig. 9. The two particles of the entangled pair correspond to two subsystems, carted off to two far away regions while they maintain their entanglement (this is where control of entanglement is essential). Then, they are brought into interaction with the measurement devices at the two ends by a unitary process, producing post-measurement systems. The unitary interaction with the measurement devices is consistent with the Everettian framework: there is no collapse, and hence no globally unique outcome. These post-measurement systems are then brought together where they interact unitarily with a device that measures the correlations between the two post-measurement systems, producing a quantum system, which represents as a diagonal density matrix the correlations. According to the Everett interpretation, this diagonal density matrix will correspond to different observers on different branches of the wavefunction seeing different sets of outcomes with different probabilities and Bell-violating statistics will have the overwhelming weight.

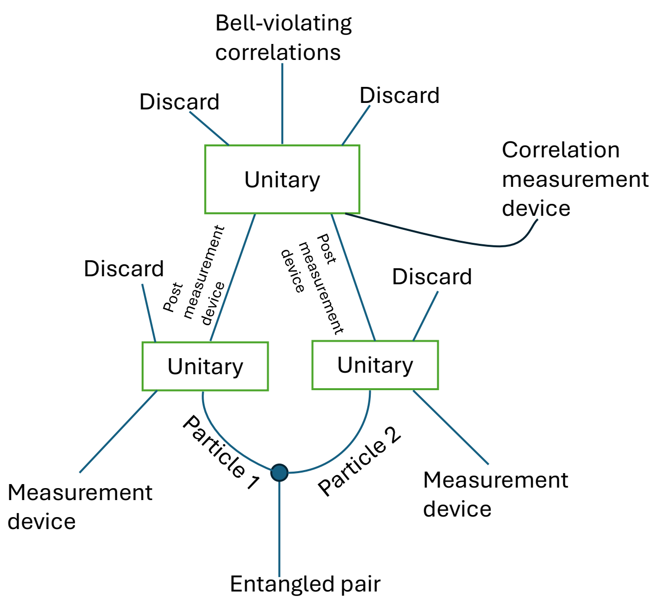

Figure 9: A quantum causal diagram representing the Everettian story of what is happening in the experiments showing violations of Bell inequalities.

One point that becomes clear from this quantum causal model is that intervening on one side of the experiment has no causal effect on the other side: we can only intervene using a quantum instrument, which doesn't allow you to select a state to "collapse" to. This is made particularly clear within the Everettian framework, which has no collapses. This is *not* transparent either in the "textbook" formulation of quantum mechanics or in attempts to represent quantum processes using classical causal models. In the "textbook" formulation, a measurement on one end instantaneously *collapses* the entangled state, causing an instantaneous definite state to obtain on the other end.

Another point becomes clear as well. The quantum causal framework doesn't require any fine-tuning to explain Bell correlations. This is made precise by formulating the faithfulness condition within this quantum causal framework. Barrett, Lorenz, and Oreshkov (2021) define faithfulness for quantum causal models roughly as follows: an assignment of an initial state and a collection of quantum channels (which results in an assignment of density matrices to subsystems) is faithful to a causal structure (represented by a DAG), insofar as that assignment allows for signaling between any two systems connected by a channel. What does "allow for signaling" mean? Essentially, it means that one can intervene (using a quantum instrument) on one of the systems and effect a change in the probability distribution of potential measurements at the other end. Thus, an assignment of channels/density matrices is unfaithful if, despite a causal pathway between two systems, we cannot signal between them.

From this we can see that there is no worry about faithfulness violations in the Everettian/quantum causal way of thinking about Bell correlations. Wherever there are causal connections between systems, those connections can be used to signal, but none of these are problematic. Signaling between the causally connected systems in Fig. 9 doesn't violate the no-signaling theorem or require invoking physics in tension with relativity. In particular, that diagram has no connections between the left wing and right wing until the worldlines of the two experimenters come into contact.

9 Clarifications and Qualifications

The message of this paper has been that classical causal explanations struggle to deal with Bell correlations, while the combination of Everettian quantum mechanics and a recently developed framework of quantum causal models can provide realistic, local, and non-fine-tuned causal explanations for Bell experiments.

One might think a bait and switch has happened in this paper. One was puzzled by the mysterious correlations in quantum experiments. One was hoping that this mysteriousness would be dispelled by a causal story. But does the Everettian quantum causal story given above *really* dispel that mystery?

It's true that we didn't provide the kind of causal story one typically expects when there are puzzling correlations. But that's because *you just can't have* that kind of causal explanation of Bell correlations: that's the upshot of EPR and Bell and Wood & Spekkens. One simply can't have a fully satisfying classical causal explanation of quantum correlations. You can't always get what you want. What we have instead is a story that allows us to talk the causal influences that *do* obtain in these experiments.

Another question that might arise is whether non-Everettian realists about quantum mechanics are somehow unable to use quantum causal models. They can, but for the leading non-Everettian realist views about quantum mechanics, there won't be much motivation to reach for quantum causal models. For instance, a Bohmian is committed to being able to explain quantum correlations using broadly the kind of explanation that one uses in classical contexts; viewing quantum systems as particles with definite locations is central to their appeal. Thus, it makes more sense within the Bohmian worldview to bite the bullet of faithfulness violations and avoid being pressed towards quantum causal models.

This doesn't mean that non-Everettians can't ever find use in quantum causal models. For instance, the Bohmian can focus on the state vector evolving unitarily and then they can perfectly well employ the quantum causal modeling framework above. However, the Bohmian will treat the quantum causal modeling framework as largely instrumental; the real causal relations will involve the Bohmian particles. What's distinctive about the Everettians' use of the quantum causal modeling framework is that they can straightforwardly interpret the framework realistically. This is a special case of a broader point: since the mathematical framework of the Everettian is just unitary quantum mechanics, they can adopt any well-developed quantum-theoretic framework 'off-the-shelf', without having to retool it to fit their worldview.

One final question: Are quantum causal explanations really causal explanations? We took the classical causal modeling framework seriously because we thought it embedded the core assumptions of what good causal explanations consist in. If these core assumptions no longer hold when it comes to quantum causal models, then to what extent can I say that the quantum causal explanations are actually causal explanations?

I don't think the concern applies to my work because what I'm challenging when it comes to the causal modeling framework are not the assumptions that embody the principles we expect causal explanations to satisfy—such as the Markov condition or faithfulness—but rather the background *physical* presuppositions that that framework employs. Indeed, the new quantum causal framework is trying to keep the older principles of causal explanations around—i.e., it is trying to retain notions of intervention and trying to formulate appropriate versions of the Markov condition and faithfulness for the quantum context. Instead, what my arguments have done is to show that there was always a *physical* assumption (about the representative aptness of random variables) built into the causal modeling framework and that that *physical* assumption is *defeated* in the context of quantum mechanics. This shows us the need to develop a version of our causal modeling that accords with our knowledge of the physical world, while retaining well-tested principles of causal explanations.

Acknowledgments. Special thanks to David Wallace and Jim Woodward for valuable discussions and comments on multiple drafts. Thanks to Gal Ben-Porath, Bixin Guo, Nick Huggett, and especially, Alyssa Ney for useful comments and discussions. Thanks to audiences at the Everett and Locality workshop in New Orleans and at the PSA 2024, and specifically, Emily Adlam, Jer Steeger, Alex Meehan, and Doreen Fraser for their questions. Thanks also to those who came to my poster on this topic at the PSA 2022 (Pittsburgh) and Foundations of Physics 2023 (Bristol); in particular to Neil Dewar, John Dougherty, Alex Franklin, and Katie Robertson.

References

Albert, David Z, and Rivka Galchen. 2009. "A quantum threat to special relativity." *Scientific American* 300 (3): 32–39.

Allen, John-Mark A., Jonathan Barrett, Dominic C. Horsman, Ciarán M. Lee, and Robert W. Spekkens. 2017. "Quantum Common Causes and Quantum Causal Models." *Phys. Rev. X* 7 (3): 031021.

Andersen, Holly. 2013. "When to Expect Violations of Causal Faithfulness and Why It Matters." *Philosophy of Science,* no. 5, 672–683.

Barrett, Jonathan, Robin Lorenz, and Ognyan Oreshkov. 2021. "Cyclic quantum causal models." *Nature Communications* 12 (1): 1–15.

Bell, J. S. 1964. "On the Einstein Podolsky Rosen paradox." *Physics Physique Fizika* 1 (3): 195–200.

Bell, John S. 1975. "The theory of local beables." *Epistemological Letters* 9:11–24. Reprinted in: Bell, John S., Abner Shimony, Michael A. Horne, and John F. Clauser. "An exchange on local beables." *Dialectica* (1985): 85-110. https://www.jstor.org/stable/42970534.

Brunner, Nicolas, Daniel Cavalcanti, Stefano Pironio, Valerio Scarani, and Stephanie Wehner. 2014. "Bell nonlocality." *Rev. Mod. Phys.* 86 (2): 419–478.

Egg, Matthias, and Michael Esfeld. 2014. "Non-local common cause explanations for EPR." *European Journal for Philosophy of Science* 4 (2): 181–196.

Einstein, A., B. Podolsky, and N. Rosen. 1935. "Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?" *Phys. Rev.* 47 (10): 777–780.

Evans, Peter W. 2021. "A Sideways Look at Faithfulness for Quantum Correlations." *Journal of Philosophy* 118 (1): 28–42.

Giustina, Marissa, Alexandra Mech, Sven Ramelow, Bernhard Wittmann, Johannes Kofler, Jörn Beyer, Adriana Lita, Brice Calkins, Thomas Gerrits, Sae Woo Nam, et al. 2013. "Bell violation using entangled photons without the fair-sampling assumption." *Nature* 497 (7448): 227–230.

Goldstein, Sheldon, and Ward Struyve. 2007. "On the uniqueness of quantum equilibrium in Bohmian mechanics." *Journal of Statistical Physics* 128:1197–1209.

Harrigan, Nicholas, and Robert W Spekkens. 2010. "Einstein, incompleteness, and the epistemic view of quantum states." *Foundations of Physics* 40 (2): 125–157.

Hausman, Daniel M. 1999. "Lessons From Quantum Mechanics." *Synthese* 121 (1-2): 79–92.

Hausman, Daniel M., and James Woodward. 1999. "Independence, Invariance and the Causal Markov Condition." *The British Journal for the Philosophy of Science* 50 (4): 521–583.

———. 2004. "Modularity and the Causal Markov Condition: A Restatement." *The British Journal for the Philosophy of Science* 55 (1): 147–161.

Healey, Richard. 2023. "Quantum-Bayesian and Pragmatist Views of Quantum Theory." In *The Stanford Encyclopedia of Philosophy,* Winter 2023, edited by Edward N. Zalta and Uri Nodelman. Metaphysics Research Lab, Stanford University.

Hensen, Bas, Hannes Bernien, Anaïs E Dréau, Andreas Reiserer, Norbert Kalb, Machiel S Blok, Just Ruitenberg, Raymond FL Vermeulen, Raymond N Schouten, Carlos Abellán, et al. 2015. "Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres." *Nature* 526 (7575): 682–686.

Hitchcock, Christopher. 2020. "Causal Models." In *The Stanford Encyclopedia of Philosophy,* Summer 2020, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University.

Howard, Don. 1985. "Einstein on locality and separability." *Studies in History and Philosophy of Science Part A* 16 (3): 171–201.

Malinsky, Daniel, and David Danks. 2018. "Causal Discovery Algorithms: A Practical Guide." *Philosophy Compass* 13 (1): e12470.

Maudlin, Tim. 2014. "What Bell did." *Journal of Physics A: Mathematical and Theoretical* 47 (42): 424010.

———. 2019. *Philosophy of Physics: Quantum Theory.* Princeton University Press.

Myrvold, Wayne, Marco Genovese, and Abner Shimony. 2021. "Bell's Theorem." In *The Stanford Encyclopedia of Philosophy,* Fall 2021, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University.

Näger, Paul. 2016. "The Causal Problem of Entanglement." *Synthese* 193 (4): 1127–1155.

Pearl, Judea. 2000. *Causality.* Cambridge University Press.

Rauch, Dominik, Johannes Handsteiner, Armin Hochrainer, Jason Gallicchio, Andrew S. Friedman, Calvin Leung, Bo Liu, et al. 2018. "Cosmic Bell Test Using Random Measurement Settings from High-Redshift Quasars." *Phys. Rev. Lett.* 121 (8): 080403.

Reichenbach, Hans. 1956. *The Direction of Time.* Edited by Maria Reichenbach. Mineola, N.Y.: Dover Publications.

Saunders, Simon, Jonathan Barrett, Adrian Kent, and David Wallace. 2010. *Many worlds?: Everett, quantum theory, & reality.* Oxford University Press.

Shalm, Lynden K., Evan Meyer-Scott, Bradley G. Christensen, Peter Bierhorst, Michael A. Wayne, Martin J. Stevens, Thomas Gerrits, et al. 2015. "Strong Loophole-Free Test of Local Realism." *Phys. Rev. Lett.* 115 (25): 250402.

Shimony, A., M.A. Horne, and J.F. Clauser. 1976. "Comment on "Theory of Local Beables"." *Epistemological Letters* 13:1–8. Reprinted in: Bell, John S., Abner Shimony, Michael A. Horne, and John F. Clauser. "An exchange on local beables." *Dialectica* (1985): 85-110. https://www.jstor.org/stable/42970534.

Shrapnel, Sally. 2014. "Quantum Causal Explanation: Or, Why Birds Fly South." *European Journal for Philosophy of Science* 4 (3): 409–423.

Spirtes, Peter, Clark Glymour, and Richard Scheines. 2000. *Causation, Prediction, and Search.* MIT Press: Cambridge.

Suárez, Mauricio, and Iñaki San Pedro. 2010. "Causal Markov, Robustness and the Quantum Correlations." In *Probabilities, Causes and Propensities in Physics,* edited by Mauricio Suárez, 173–193. Springer.

Van Fraassen, Bas C. 1982. "The Charybdis of Realism: Epistemological Implications of Bell's Inequality." *Synthese* 52 (1): 25–38.

Wallace, David. 2012. *The emergent multiverse: Quantum theory according to the Everett interpretation.* Oxford University Press.

Weinberger, Naftali. 2018. "Faithfulness, Coordination and Causal Coincidences." *Erkenntnis* 83 (2): 113–133.

Weinberger, Naftali, Porter Williams, and James Woodward. 2024. "The Worldly Infrastructure of Causation." *The British Journal for the Philosophy of Science.*

Wood, Christopher J, and Robert W Spekkens. 2015. "The lesson of causal discovery algorithms for quantum correlations: causal explanations of Bell-inequality violations require fine-tuning." *New Journal of Physics* 17, no. 3 (March): 033002.

Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation.* Oxford University Press.