Epistemic interpretations of quantum theory have a measurement problem

Joshua B. Ruebeck,[1,*] Piers Lillystone,[1] and Joseph Emerson[2,3]

[1] *Institute for Quantum Computing and Department of Physics and Astronomy University of Waterloo, Waterloo, Ontario N2L 3G1, Canada*

[2] *Institute for Quantum Computing and Department of Applied Math University of Waterloo, Waterloo, Ontario N2L 3G1, Canada*

[3] *Canadian Institute for Advanced Research, Toronto, Ontario M5G 1Z8, Canada*

(Dated: March 1, 2019)

Epistemic interpretations of quantum theory maintain that quantum states only represent incomplete information about the physical states of the world. A major motivation for this view is the promise to provide a reasonable account of state update under measurement by asserting that it is simply a natural feature of updating incomplete statistical information. Here we demonstrate that all known epistemic ontological models of quantum theory in dimension $d \geq 3$, including those designed to evade the conclusion of the PBR theorem, cannot represent state update correctly. Conversely, interpretations for which the wavefunction is real evade such restrictions despite remaining subject to long-standing criticism regarding physical discontinuity, indeterminism and the ambiguity of the Heisenberg cut. This revives the possibility of a no-go theorem with no additional assumptions, and demonstrates that what is usually thought of as a strength of epistemic interpretations may in fact be a weakness.

I. INTRODUCTION

There are many interpretations of quantum theory[1]. Among the many differences between these interpretations, one that often takes center stage is the stance that they take towards the wavefunction or quantum state. Three broad categories have been identified which capture a number of interpretations. Two of these categories are more commonly juxtaposed: ontic interpretations [1–14] posit that the quantum state is a part of the *real* (physical) state of a system, whereas epistemic interpretations [15–23] argue that the quantum state is merely a state of *knowledge* about the real state of the system. A very thorough review of these two stances can be found in [24]. A third recently articulated category of doxastic interpretations [25–31] argue that the quantum state is a state of *belief*, and are distinguished from epistemic interpretations by the fact that they deny that a system has some 'real state.' While not all interpretations conform to these three descriptors, they are useful categories insofar as they allow us to qualitatively discuss certain features separately from the particular interpretation in which they are embedded.

The epistemic class has garnered attention as a view which provides very appealing explanations of otherwise paradoxical features of quantum theory like the state update rule [17], the classical limit under quantum chaos [20], no cloning [32], and entanglement [33]. For example, through this viewpoint state update is not a physical 'collapse' process and therefore not subject to paradoxes, indeterminism, and discontinuity; rather it is understood as analogous to the non-pardoxical 'collapse' of a subjective probability distribution via Bayes' rule upon consideration of new information. While several epistemic models have been proposed [16, 22, 34–38], they generally have undesirable features or are restricted to a subtheory of full quantum theory. None achieve all of the features that an optimistic epistemicist would expect.

This suggests the possibility that a fully satisfactory epistemic interpretation cannot actually explain all of quantum theory despite the qualitatively compelling features of such a view[2]. This suspicion has led to a number of no-go theorems in recent years which establish that, given at least one additional assumption, any consistent interpretation of quantum theory cannot be epistemic [36, 40–43]. These no-go theorems are generally proven within the ontological models formalism, which describes a large class of existing interpretations of quantum theory [21, 44, 45].

Within the ontological models formalism, epistemic models can be given a precise mathematical definition called the ψ-epistemic criterion [22]. This precise criterion allows the possibility of conclusively ruling out this type of model. Outside of this framework, it is unlikely that ψ-epistemic models can be precluded with any kind of certainty; doxastic interpretations, for example, do not fit neatly into the ontological models framework and thus are not necessarily ruled out by these no-go theorems. This is despite the fact that they share many of the features which make ψ-epistemic interpretations appealing. In the present paper we restrict our attention to the ontological models framework.

* jbruebeck@uwaterloo.ca

[1] This is an understatement.

[2] As is well-known, the work of Bell [39] has shown that no epistemic interpretation can evade the non-locality that manifests trivially in ontic interpretations; this trivial manifestation of nonlocality in ontic interpretations is an oft-forgotten insight from Einstein [15, 18, 23].

The fact that an extra assumption is required to rule out ψ-epistemic theories has purportedly been demonstrated by the existence of ψ-epistemic models which, while being individually unsatisfactory for various reasons, do satisfy at least the bare minimum requirements of a ψ-epistemic theory [35–37]. All of these models were specified within a prepare-measure framework, so they have been proven to reproduce quantum statistics for all experiments that involve preparing a state and then measuring it once. In this paper we show that, if we allow sequential measurements in the operational description, these models cannot reproduce operational statistics.

Our main contribution in this work is thus to demonstrate that the state update rule imposes severe constraints on ψ-epistemic models. This is in contrast to the prevailing view that, as articulated by Liefer, "a straightforward resolution of the collapse of the wavefunction, the measurement problem, Schrödinger's cat and friends is one of the main advantages of ψ-epistemic interpretations" [24]. As a consequence, we revive the possibility of a general no-go theorem for ψ-epistemic models that doesn't rely on an additional assumption such as the locality assumption required in [40] which conflicts with the non-locality that is implied by Bell's theorem [43, 46].

Although state update under measurement has been described in a few specific models [22, 34, 47, 48] and discussed with regards to contextuality [49], it has yet to be treated generally or in relation to the ψ-epistemic/ψ-ontic distinction[3]. Here we take some preliminary steps in both of these directions, and argue that ψ-epistemic models are the natural arena in which to investigate interesting behavior of state update under measurement.

In Section II, we describe the ontological models formalism, adding a description of state update under measurement and motivating its importance. Despite this motivation, one might still argue that many distinctly quantum phenomena (e.g. Bell inequality violations) can be described without reference to state update; thus, from an operationalist point of view, we shouldn't need to consider state update in order to investigate these phenomena. However, we show in Section III that the consideration of state update actually places nontrivial restrictions on how one can represent even a prepare-and-measure-once experiment. Thus our results are directly applicable to models which have only specified behavior for a single measurement. Section IV reviews a number of examples of ontological models from the literature; in each case we either specify its state update rule (in dimension $d = 2$ and for ψ-ontic models) or prove its impossibility (for all known ψ-epistemic models in dimension $d \geq 3$). Finally, we discuss the implications of our results and describe some open questions in Section V.

[3] Although the Leggett-Garg inequalities [50, 51] might be construed as a general treatment of state update, it is more accurate to say that they are about the *absence* of state update.

II. DEFINING MEASUREMENT UPDATE IN ONTOLOGICAL MODELS

A. The ontological models formalism

In the standard treatment, an *operational theory* [44] is described by a set of preparations $\mathcal{P}$, a set of transformations $\mathcal{T}$, and a set of measurements $\mathcal{M}$ along with a probability distribution

$$\Pr(k|M, T, P). \tag{1}$$

This quantity describes the probability of some measurement outcome $k \in \mathbb{Z}$ given an experimenter's choice of $P \in \mathcal{P}$, $T \in \mathcal{T}$, and $M \in \mathcal{M}$. When considering transformations this is called the *prepare-transform-measure* operational framework, and when we omit transformations it is the *prepare-measure* framework. Often we take $\mathcal{P}$ to be the set of pure quantum state preparations, $\mathcal{T}$ to be the full set unitary maps on a Hilbert space, and $\mathcal{M}$ to be all projective measurements on this Hilbert space. In this case, we say we are describing the full quantum theory[4]; in contrast, a *subtheory* is described by taking subsets of $\mathcal{P}, \mathcal{T}, \mathcal{M}$ for the full quantum theory. For example, in quantum information settings we often consider only measurements in the standard basis.

Note that this standard definition involves a single measurement and a single measurement outcome despite the fact many important quantum experiments (e.g. Stern-Gerlach, double slit [52]) and quantum algorithms (e.g. measurement-based error correction [53]) involve multiple measurements. Thus we will refer to the usual definition of prepare-measure as prepare-and-measure-once experiments. In this paper, we are concerned with multiple measurements, so we will also have to describe probabilities like

$$\Pr(k_2, k_1|M_2, M_1, P). \tag{2}$$

Additionally, we note that positive operator valued measures (POVMs) do not fully specify how a measurement updates a state. Although one can obtain a POVM $\{E_k\}$ from a set of generalized measurement operators $\{M_k\}$ by the relation $E_k = M_k^\dagger M_k$, the decomposition of $\{E_k\}$ into $\{M_k\}$ is not unique. Thus although $\mathcal{M}$ is often described by POVMs, consideration of state update requires that we specify generalized measurement operators instead. As an example of when this is important, consider a coarse-graining of the measurement $\{M_k\} = \{[0], [1], [2]\}$, where we denote the projector onto a state

$$[\psi] = |\psi\rangle\langle\psi| \tag{3}$$

[4] Larger sets can be considered (e.g. including mixed states, CPTP maps, or non-projective measurements), but all of the models studied in this paper fit the given definition.

as in [24]. We can either coarse-grain coherently, i.e. measure $\{M'_k\} = \{[0]+[1],[2]\}$ or we can coarse-grain decoherently by measuring $\{M_k\}$ and then combining outcomes 0 and 1 into a single measurement result and 'forgetting' which one actually occurred. While these two processes are represented by the same POVM $\{E_k\} = \{[0]+[1],[2]\}$, their state update behavior is different: if the state $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ is measured, it will stay the same in the coherent case or update to the mixed state $\frac{1}{2}([0]+[1])$ in the decoherent case.

An *ontological model* [21, 44, 45] supplements this operational point of view by asserting that a system has a state λ, called an *ontic state*. To specify an ontological model, we first choose an *ontic state space* Λ. Then, preparations are described by a *preparation distribution* $\mu(\lambda|P)$, which is the probability of preparing some state $\lambda \in \Lambda$ given the preparation P. Transformations are described by a *transition matrix* $\Gamma(\lambda'|\lambda, T)$, which is the probability of preparing a new state λ' given the previous state λ and the choice T of transformation. Finally, measurements are represented by a *response function* $\xi(k|\lambda, M)$ which describes the probability of an outcome k given the ontic state λ and the choice of measurement M. We say that an ontological model successfully reproduces quantum theory if

$$\int_\Lambda \mathrm{d}\lambda' \int_\Lambda \mathrm{d}\lambda\, \xi(k|\lambda', M)\Gamma(\lambda'|\lambda, T)\mu(\lambda|P) = \mathrm{Pr}_Q(k|M,T,P) \quad \forall P \in \mathcal{P}, T \in \mathcal{T}, M \in \mathcal{M}, \tag{4}$$

or, in a prepare-and-measure-once experiment,

$$\int_\Lambda \mathrm{d}\lambda\, \xi(k|\lambda, M)\mu(\lambda|P) = \mathrm{Pr}_Q(k|M,P) \quad \forall P \in \mathcal{P}, M \in \mathcal{M}. \tag{5}$$

In both of these cases, Pr_Q indicates the outcome probability calculated by operational quantum theory for the particular experiment under consideration.

Again, we must supplement this definition in order to model repeated measurement. In textbook quantum theory, the state updates during a measurement in a way that depends on the previous state, the measurement procedure, and the measurement outcome. We allow for dependence on all of these things by choosing to represent this via a *state update rule* $\eta(\lambda'|k, \lambda, M)$. Although this object looks very similar to the transition matrix for transformations, it is distinguished by two important features which we emphasize by choosing a new symbol to represent it.

The first distinction is simple, in that η depends on a measurement outcome k, while Γ does not; this is analogous to the fact that generally in quantum theory we can only implement measurement update maps probabilistically (i.e. by post-selecting on a not-necessarily-deterministic measurement outcome).

The second distinction is the fact that $\eta(\lambda'|k, \lambda, M)$ is not defined for all $\lambda \in \Lambda$. Roughly speaking, it doesn't make sense to ask "What is the new state λ' after an outcome k?" if the outcome k could not have occurred given the previous state λ. To express this formally, we define the *support* of a distribution

$$\mathrm{Supp}(\xi(k|\cdot)) = \{\lambda \in \Lambda : \xi(k|\lambda) > 0\} \tag{6}$$

as the set of ontic states on which it is nonzero. Using this, we can say that $\eta(\lambda'|k, \lambda, M)$ is well-defined only for $\lambda \in \mathrm{Supp}(\xi(k|\cdot))$. This fact is central to the main result of this paper; by showing that η is non-normalizable on some domain, we are able to conclude that that domain cannot be part of the support of ξ. This property of η is analogous to the fact that while post-selected state update maps are completely positive (CP) in quantum theory, they have a nontrivial kernel and so are not trace-preserving (TP).

The consistency condition with quantum theory is given by equations similar to Eqs. 4 and 5; see Appendix A for more formal treatments of the properties of η and other extensions of the ontological models formalism discussed so far. The rigorous treatment in the appendix requires drawing on the stochastic processes and hidden Markov models literature, which additionally provides insight into the basic assumptions of the ontological models formalism itself.

This paper is not explicitly concerned with contextuality [16, 44], but we are careful to ensure that we do not assume noncontextuality. Under the definition of generalized contextuality given in [44], an ontological model is *noncontextual* if two operationally equivalent preparation, transformation, or measurement procedures are always represented by equivalent preparation distributions, transition matrices, or response functions, respectively[5]. It is *contextual* otherwise. While this generalized definition may be too broad, accounting for contextuality under this definition also accounts for the traditional definition [16], and is therefore more inclusive. Although, as stated above, we do not assume noncontextuality of any kind, most of our results hold for a projector independent of its full measurement context. We are explicit about this when it is the case, and write $\xi(\Pi|\lambda)$ rather than $\xi(k=0|\lambda, M=\{\Pi, \ldots\})$ for notational convenience; $\eta(\lambda'|\lambda, \Pi)$ is defined similarly.

B. ψ-epistemic models

Here we focus on a set of precise criteria for epistemic interpretations within the ontological models formalism. Consider first the ψ-epistemic criterion, proposed in [23] as a test for whether an interpretation admits at least some quantum states that are not uniquely determined by the underlying state of reality. Following [49], we

[5] This definition is actually complicated slightly by the inclusion of repeated measurements, but we do not discuss this here.

account for potential preparation noncontextuality by defining

$$\Delta_\psi = \bigcup_{P_\psi \in \mathcal{P}_\psi} \mathrm{Supp}(\mu(\cdot|P_\psi)) \tag{7}$$

where $\mathcal{P}_\psi \subseteq \mathcal{P}$ is the set consisting of every possible preparation of $|\psi\rangle$. We refer to Δ_ψ as the support of a state $|\psi\rangle$, to distinguish it from the support of a particular preparation P_ψ. Then a pair of states $|\psi\rangle, |\phi\rangle$ is *ontologically distinct* in a particular model if $\Delta_\phi \cap \Delta_\psi = \emptyset$, and *ontologically indistinct* otherwise[6]. This leads us to the standard definition of a ψ-epistemic ontological model [23, 24]:

Definition 1 (ψ-epistemic)**.** *An ontological model is ψ-epistemic if there exists a pair of states $|\psi\rangle, |\phi\rangle$ that are ontologically indistinct; i.e. $\exists |\psi\rangle, |\phi\rangle : \Delta_\phi \cap \Delta_\psi \neq \emptyset$.*

As noted in [24], this definition is highly permissive in the sense that, if an ontological model were to contain only a single pair of quantum states that are ontologically indistinct, then it would not achieve the full explanatory power expected of the ψ-epistemic viewpoint; this is exactly the case with the ABCL_0 model discussed in Section IV B 4.

There are, however, proposals to strengthen the notion of ψ-epistemicity, two of which are relevant to our discussion [24, 54].

Definition 2 (Pairwise ψ-epistemic)**.** *An ontological model is* pairwise ψ-epistemic *if, for all pairs $|\psi\rangle, |\phi\rangle$ of nonorthogonal quantum states, $|\psi\rangle$ and $|\phi\rangle$ are ontologically indistinct.*

Definition 3 (Never ψ-ontic)**.** *An ontological model is* never ψ-ontic *if every ontic state $\lambda \in \Lambda$ is in the support of at least two quantum states:*

$$\forall \lambda \in \Lambda : \exists \psi, \phi : \lambda \in \Delta_\psi \cap \Delta_\phi. \tag{8}$$

Note that both of these definitions imply the weaker notion of ψ-epistemicity, but are independent from one another.

C. Some easy cases of state update rules

There are two cases in which, given a prepare-and-measure-once ontological model for quantum theory, we can always augment it with a state update rule for measurement. First, if we only include rank-1 projective measurements in the subtheory we're modeling, we can simply re-prepare in the measured (unique, pure) state:

$$\eta(\lambda'|k, \lambda, M_{\{\Pi_i\}}) = \mu(\lambda'|P_{\Pi_k}) \text{ for } \mathrm{tr}(\Pi_k) = 1. \tag{9}$$

This is normalized for all λ since μ is normalized, and faithfully reproduces quantum statistics since μ does. It is independent of the previous state λ, which, besides being unsatisfying, is also not possible in general (this follows from Section III).

Second, it is quick to prove, again by construction, that ψ-ontic models can always be given a state update rule. Since there is a unique quantum state $|\psi_\lambda\rangle$ associated with every ontic state λ, we can define

$$\eta(\lambda'|k, \lambda, M) = \mu\left(\lambda' \middle| P = \frac{M_k[\psi_\lambda]M_k^\dagger}{\mathrm{tr}\left(M_k[\psi_\lambda]M_k^\dagger\right)}\right). \tag{10}$$

Again, normalization and faithfulness follow because μ has these properties. Note that this construction works for any kind of measurement, not just projective measurements.

These observations together suggest that in order to find anything interesting involving state update, we ought to examine higher-rank measurements in ψ-epistemic theories. This suspicion will be confirmed by the main result of this paper, which applies to exactly these types of measurements.

III. THE CONSEQUENCES OF STATE UPDATE

We now prove our central claim that consideration of a rule for state update under measurement has consequences for the response function of an ontological model, so that consistent state update puts restrictions on how one may represent even a prepare-and-measure-once experiment. We begin with a lemma that articulates a general property of the update rule η, and then examine its consequences for response functions ξ.

Lemma 1. *Suppose we have an ontological model with ontic space Λ, preparation distributions $\mu(\lambda|P)$, indicator functions $\xi(k|\lambda, M)$, and state update maps $\eta(\lambda'|k, \lambda, M)$. For a particular ontic state λ and measurement projector Π, we define the set $S_{\lambda,\Pi}$ of quantum states that one could obtain after measurement of any quantum state consistent with λ:*

$$S_{\lambda,\Pi} = \left\{ \frac{\Pi|\phi\rangle}{\sqrt{\langle\phi|\Pi|\phi\rangle}} \,\middle|\, \forall |\phi\rangle : \lambda \in \Delta_\phi \right\}. \tag{11}$$

It is then true that, independently of the measurement context of Π,

$$\mathrm{Supp}(\eta(\cdot|\lambda, \Pi)) \subseteq \bigcap_{|\psi\rangle \in S_{\lambda,\Pi}} \Delta_\psi. \tag{12}$$

[6] For this purposes of this paper we assume there exists a measure that is absolutely continuous to all other measures in the ontological model. Therefore, we can work with probability densities, rather than the full measure-theoretic treatment. While this assumption is not strictly true in all of our models, it does not affect our results.

Proof. Suppose that measuring a state $|\phi\rangle$ with a measurement M results in the updated state $|\psi\rangle$ when we get outcome k, where Π is the kth projector in M. Then let $P_{M,k,P_\phi} \in \mathcal{P}_\psi$ be the preparation procedure associated with post-selection of this measurement outcome after a particular preparation P_ϕ. It must be normalized on Δ_ψ:

$$\begin{aligned} 1 &= \int_{\Delta_\psi} \mathrm{d}\lambda' \, \mu(\lambda'|P_{M,k,P_\phi}) \\ &= \int_{\Delta_\psi} \mathrm{d}\lambda' \int_{\Delta_\phi} \mathrm{d}\lambda \, \eta(\lambda'|k,\lambda,M)\mu(\lambda|P_\phi) \\ &= \int_{\Delta_\phi} \mathrm{d}\lambda \, \mu(\lambda|P_\phi) \int_{\Delta_\psi} \mathrm{d}\lambda' \, \eta(\lambda'|k,\lambda,M) \end{aligned}$$

Since η is always positive, normalization of $\mu(\lambda|P_\phi)$ then implies that

$$\int_{\Delta_\psi} \mathrm{d}\lambda' \, \eta(\lambda'|k,\lambda,M) = 1 \qquad \forall\lambda \in \Delta_\phi.$$

If η is normalized on a region, its support must be contained in that region. Thus for all λ that are consistent with some preparation $|\phi\rangle$ that could result in the post-measurement state $|\psi\rangle$,

$$\mathrm{Supp}(\eta(\cdot|k,\lambda,M)) \subseteq \Delta_\psi.$$

The fact that this is true for all $|\psi\rangle$ that could result from the measurement leads to Eq. 12. □

We note that Lemma 1 is easy to account for in ψ-ontic theories and for rank-1 measurements, since in both cases $S_{\lambda,\Pi}$ has a single element. This is why we were able to write down update rules for these situations in Section II C. Outside of these trivial cases, Eq. (12) is a very restrictive condition; depending on the structure of $S_{\lambda,\Pi}$, the intersection may be a very small set. In particular, if any two of the post-selected quantum states are orthogonal, then $S_{\lambda,\Pi}$ is empty.

Theorem 1 (Main theorem)**.** *Suppose that a projector* Π *maps two states* $|\alpha\rangle, |\beta\rangle$ *to ontologically distinct states* $\Pi|\alpha\rangle, \Pi|\beta\rangle$*. Then the response function for* Π *cannot have support on the overlap of* $|\alpha\rangle, |\beta\rangle$ *for any measurement context of* Π*; i.e.*

$$\Delta_{\Pi|\alpha\rangle} \cap \Delta_{\Pi|\beta\rangle} = \emptyset \implies \xi(\Pi|\lambda) = 0 \ \ \forall\lambda \in \Delta_\alpha \cap \Delta_\beta. \quad (13)$$

Proof. Pick some $\lambda \in \Delta_\alpha \cap \Delta_\beta$. By Lemma 1, $\mathrm{Supp}(\eta(\cdot|\lambda,\Pi)) = \emptyset$ so $\eta(\lambda'|\lambda,\Pi)$ is not normalizable. As discussed in Section II A, this is only allowable if $\xi(\Pi|\lambda) = 0$. □

Both the lemma and the theorem hold for non-projective measurements as well. We emphasize that this result does not say anything *directly* about the overlap of the supports of quantum states, just their overlap within the support of a particular response function. In the following section, we deploy this theorem by showing that, in every known ψ-epistemic model for $d \geq 3$, measurements of this type exist *and* have support on the relevant overlaps, leading to contradiction and demonstrating that these models cannot reproduce state update under measurement.

IV. EXAMPLES OF STATE UPDATE UNDER MEASUREMENT (OR ITS IMPOSSIBILITY)

We provide a number of examples of ontological models from the literature, illustrating some of the properties described in the previous sections. For each model, we either specify its state update rule or prove that it cannot reproduce state update. Although many of these models can be easily defined for arbitrary types of measurements, we only consider projective measurements for simplicity and notational consistency.

A. ψ-epistemic models of a qubit

1. Kochen-Specker model

The Kochen-Specker model of a qubit [16, 24] is an exemplar of what we look for in a ψ-epistemic theory, with the unfortunate feature that it only works in $d = 2$ dimensions. It is both pairwise ψ-epistemic and never ψ-ontic, and provides a very intuitively pleasing interpretation of the statistical nature of quantum theory. We take the ontic space to be the unit sphere S^2, and denote by $\vec{\psi}$ the Bloch vector corresponding to $|\psi\rangle$ under the usual mapping. Preparations and measurement outcomes are represented by distributions over hemispheres, with response functions uniform and preparation distributions peaked towards the center (Fig. 1a). Unitary transformations are represented by rotations of the sphere. Since the only nontrivial measurements on a qubit are rank-1 measurements, this is a case where we can use the state update rule described in Eq. 9 and just re-prepare the measured state for our update rule:

$$\begin{aligned} \Lambda &= S^2 \\ \mu(\vec{\lambda}|P_\psi) &= \frac{1}{\pi}\Theta(\vec{\psi}\cdot\vec{\lambda})\vec{\psi}\cdot\vec{\lambda} \\ \Gamma(\vec{\lambda}'|\vec{\lambda},T_U) &= \delta(\vec{\lambda}' - R_U\vec{\lambda}) \\ \xi(k|\vec{\lambda},M_\phi) &= \Theta(k\vec{\phi}\cdot\vec{\lambda}) \\ \eta(\vec{\lambda}'|k,\vec{\lambda},M_\phi) &= \frac{1}{\pi}\Theta(k\vec{\phi}\cdot\vec{\lambda}')k\vec{\phi}\cdot\vec{\lambda}' \end{aligned} \quad (14)$$

Here Θ is the Heaviside step function, R_U is the rotation of the Bloch sphere corresponding to a unitary U, and $k \in \{+1, -1\}$. This particular state update rule is not particularly satisfactory in an explanatory sense, since it is independent of the previous ontic state.

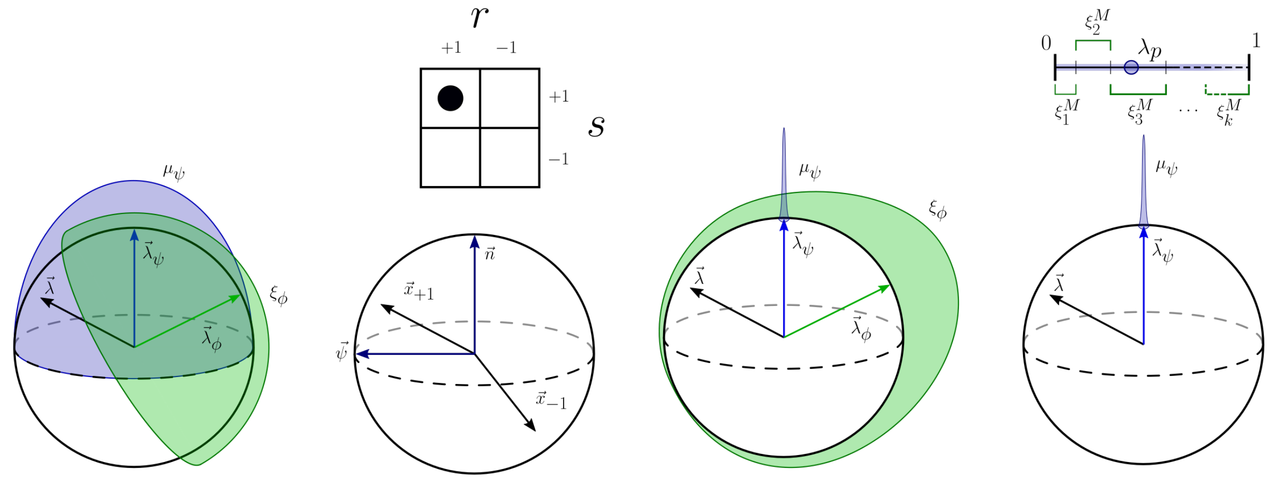

(a) Kochen-Specker model

(b) Montina model

(c) Beltrametti-Bugajski model

(d) Bell Model

Figure 1: Visualizations of the state space, preparation distributions, and response functions for (a) the Kochen-Specker model, (b) Montina's model, (c) the Beltrametti-Bugajski model for $d = 2$, and (d) Bell's model for $d = 2$. Blue represents the support of preparations, and green the support of the response functions, where possible. Black objects are generic elements of the state space.

2. *Montina model*

In [34], Montina introduces an ontological model based on the Kochen-Specker model. The model was constructed to show that state update in a qubit can be successfully modeled by only updating a finite amount of information in the ontic state. To do so, Montina extends the ontic space of the Kochen-Specker model by taking two vectors $\vec{x}_{+1}, \vec{x}_{-1}$ on the Bloch sphere and adding two bits, labeled r and s, such that the vectors on the Bloch sphere are dynamic under transformations but remain static under measurement update (Fig. 1b). The bit $r \in \{-1, +1\}$ acts as an index which decides which of the two Bloch vectors is 'active;' $s \in \{-1, +1\}$ stores the result of a hypothetical standard basis measurement on the state. We take the standard basis to be defined by a special vector $\vec{n}$ pointing along the z-axis.

As in the Kochen-Specker model, unitary transformations act by rotating the Bloch vectors; additionally, if the vector $\vec{x}_r$ (i.e. the 'active' Bloch vector) crosses the horizontal equator of the sphere during this transformation, then the bit s flips to $-s$. r does not change during a transformation.

A measurement in the standard basis simply reveals the value of s, and then updates r so that the active vector is the one which was more closely aligned with $\vec{n}$ at the time of measurement. For any other basis, we apply the unitary that maps our desired measurement basis to the standard basis, measure, and then rotate back—this whole process has been wrapped into our definitions of η and ξ below. In either case, the vectors $\vec{x}_{+1}, \vec{x}_{-1}$ do not change during measurement.

Finally, we prepare a state $\vec{\psi}$ by measuring in the basis $\{\vec{\psi}, -\vec{\psi}\}$ and applying a rotation that maps $-\vec{\psi} \to \vec{\psi}$ if we measured $-\vec{\psi}$. Summarizing these constructions, we can write

$$
\begin{aligned}
\Lambda &= S^2 \times S^2 \times \{-1, +1\} \times \{-1, +1\} \\
\lambda &= (\vec{x}_{+1}, \vec{x}_{-1}, r, s) \\
\mu(\lambda|P_\psi) &= \frac{1}{(4\pi)^2} \Theta\left[s\left(\vec{x}_r \cdot \vec{\psi}\right)(\vec{x}_r \cdot \vec{n})\right] \\
&\quad \cdot \Theta\left[r\left[\left(\vec{x}_{+1} \cdot \vec{\psi}\right)^2 - \left(\vec{x}_{-1} \cdot \vec{\psi}\right)^2\right]\right] \\
\Gamma(\lambda'|\lambda, T_U) &= \delta\left(\vec{x}'_{+1} - R_U \vec{x}_{+1}\right) \delta\left(\vec{x}'_{-1} - R_U \vec{x}_{-1}\right) \\
&\quad \cdot \Theta\left[ss'\left(\vec{x}'_r \cdot \vec{n}\right)(\vec{x}_r \cdot \vec{n})\right] \Theta[rr'] \\
\xi(k|\lambda, M_\phi) &= \Theta\left[ks\,(\vec{x}_r \cdot \vec{n})\left(\vec{x}_r \cdot \vec{\phi}\right)\right] \\
\eta(\lambda'|k, \lambda, M_\phi) &= \delta\left(\vec{x}'_{+1} - \vec{x}_{+1}\right) \delta\left(\vec{x}'_{-1} - \vec{x}_{-1}\right) \\
&\quad \cdot \Theta\left[ss'(\vec{x}_r \cdot \vec{n})\left(\vec{x}_r \cdot \vec{\phi}\right)(\vec{x}_{r'} \cdot \vec{n})\left(\vec{x}_{r'} \cdot \vec{\phi}\right)\right] \\
&\quad \cdot \Theta\left[r'\left[\left(\vec{x}_{+1} \cdot \vec{\phi}\right)^2 - \left(\vec{x}_{-1} \cdot \vec{\phi}\right)^2\right]\right]
\end{aligned}
\tag{15}
$$

The original presentation is not stated in terms of the ontological models formalism, and only explicitly models measurements in the standard basis. This lead to the claim that state update under measurement is accounted for by updating a single bit, but it is clear from the form of η above that by including all measurements in our subtheory we have caused both bits to be updated during measurement.

We include this model here for two reasons. First, it is one of the few ontological models in the literature that has explicitly considered state update under measurement. Second, it demonstrates that the generic rank-1 update (Eq. 9) that we used for the Kochen-Specker model is not the only possibility; even though all measurements in this model are rank-1, η has nontrivial dependence on the previous ontic state λ. Thus just because we *can* construct a trivial update rule in some cases does not mean that there is then nothing interesting to investigate. It also includes these features while remaining pairwise ψ-epistemic and never ψ-ontic.

B. Models of full quantum theory for arbitrary dimension

1. Beltrametti-Bugajski model

The Beltrametti-Bugajski model [9, 24] is perhaps the simplest ontological model that describes a system of arbitrary dimension. Although it is ψ-ontic, it is the starting point for the construction of the next three models in this section. For a d-dimensional quantum system, we take the ontic space to be the quantum state space, which we denote $\mathcal{PH}^{d-1}$ (the projective Hilbert space of dimension $d-1$). Preparations, transformations, measurements, and state update rules then follow directly from the usual quantum rules:

$$\begin{aligned}\Lambda &= \mathcal{PH}^{d-1}\\ \mu(\lambda|P_\psi) &= \delta(|\lambda\rangle - |\psi\rangle)\\ \Gamma(\lambda'|\lambda, T_U) &= \delta(|\lambda'\rangle - U|\lambda\rangle)\\ \xi(k|\lambda, M_{\{\Pi_i\}}) &= \langle\lambda|\Pi_k|\lambda\rangle\\ \eta(\lambda'|k,\lambda, M_{\{\Pi_i\}}) &= \delta\left(|\lambda'\rangle - \frac{\Pi_k|\lambda\rangle}{\sqrt{\langle\lambda|\Pi_k|\lambda\rangle}}\right)\end{aligned} \tag{16}$$

This provides an example of the generic update-rule for ψ-ontic models (Eq. 10), and is depicted in Fig. 1c.

2. Bell's model

Lewis et al. [35] extended a model of a qubit orignally proposed by Bell [39] to arbitrary dimension, which can be seen as a modification of the Beltrametti-Bugajski model [24]. The ontic space is the Cartesian product of the projective Hilbert space with the unit interval $[0,1]$. Now we write λ as an ordered pair $\lambda = (|\lambda\rangle, p_\lambda)$ where $|\lambda\rangle \in \mathcal{PH}^{d-1}$, as in the Beltrametti-Bugajski model, and $p_\lambda \in [0,1]$. Preparations remain essentially the same, becoming a product distribution of a delta function on the quantum state space with a uniform distribution over the unit interval. The response functions divide up the unit interval into lengths corresponding to probabilities of measuring each outcome, and respond with outcome k when p_λ is in the corresponding interval (Fig. 1d). This has the effect of making the model outcome deterministic.

$$\begin{aligned}\Lambda &= \mathcal{PH}^{d-1}\times[0,1]\\ \mu(\lambda|P_\psi) &= \delta(|\lambda\rangle - |\psi\rangle)\\ \Gamma(\lambda'|\lambda, T_U) &= \delta(|\lambda'\rangle - U|\lambda\rangle)\\ \xi(k|\lambda, M_{\{\Pi_i\}}) &= \Theta\left[p_\lambda - \sum_{j=0}^{k-1}\mathrm{tr}(\Pi_j[\lambda])\right]\\ &\quad\cdot\Theta\left[-p_\lambda + \sum_{j=0}^{k}\mathrm{tr}(\Pi_j[\lambda])\right]\\ \eta(\lambda'|k,\lambda, M_{\{\Pi_i\}}) &= \delta\left(|\lambda'\rangle - \frac{\Pi_k|\lambda\rangle}{\sqrt{\langle\lambda|\Pi_k|\lambda\rangle}}\right)\end{aligned} \tag{17}$$

Since this model is still ψ-ontic, we once again use the generic state update rule for ψ-ontic models. In this case, we can also see that this works because of the structure of the preparations as product distributions. Since every state has a uniform distribution over p_λ, and this is uncorrelated with $|\lambda\rangle$, we don't need to update any information about p_λ and so can just re-use the Beltrametti-Bugajski model update rule.

3. LJBR model

In [35], Lewis et al. define a ψ-epistemic model based on their generalization of Bell's model. Referred to here as the LJBR model, it is motivated by the observation that the order of segments in the response function of Bell's model does not matter: a re-ordering of these segments allows arbitrary modification of preparation distributions within a subset of the ontic space, so they can be made to overlap. We present here a brief description of the 'most epistemic' version of this model, and refer the reader to [35] for a more thorough construction and motivation.

The LJBR model has the same ontic space as the Bell model, so we again write ontic states as $\lambda = (|\lambda\rangle, p_\lambda)$. It is constructed in a preferred basis $\{|j\rangle\}$, which we use in defining two helper functions. First,

$$z_j(|\lambda\rangle) = \inf_{|\phi\rangle:\mathrm{tr}([j][\phi])\geq 1/d}\mathrm{tr}([\lambda][\phi]). \tag{18}$$

Note that $z_j(|\lambda\rangle) > 0$ if and only if $\mathrm{tr}([j][\lambda]) > \frac{d-1}{d}$, so $z_j(|\lambda\rangle)$ is nonzero for at most a single element of the preferred basis; we denote this unique vector as $|j_\lambda\rangle$. Second, we define a permutation $\pi_{M,\lambda}$ for each measurement M and ontic state λ:

$$\begin{aligned}\mathrm{tr}\big(M_{\pi_{M,\lambda}(0)}[j_\lambda]\big) \geq \mathrm{tr}\big(M_{\pi_{M,\lambda}(1)}[j_\lambda]\big) \geq \cdots\\ \cdots \geq \mathrm{tr}\big(M_{\pi_{M,\lambda}(|M|-1)}[j_\lambda]\big).\end{aligned} \tag{19}$$

If there is no $|j_\lambda\rangle$, i.e. $z_j(|\lambda\rangle) = 0$ for all j, then we take $\pi_{M,\lambda}$ to be the identity permutation. The final element we need before defining the model itself is a set

$$\mathcal{E}_j = \{\lambda : z_j(|\lambda\rangle) > 0\} \tag{20}$$

defined for each basis vector. Without further ado, the full specification of the model:

$$\begin{aligned}
\Lambda &= \mathcal{PH}^{d-1} \times [0,1] \\
\lambda &= (|\lambda\rangle, p_\lambda) \\
\mu(\lambda|P_\psi) &= \delta(|\lambda\rangle - |\psi\rangle) \prod_j \Theta[p_\lambda - z_j(|\psi\rangle)] \\
&\quad + \sum_j z_j(|\psi\rangle)\mu_{\mathcal{E}_j}(\lambda) \\
\xi(k|\lambda, M_{\{\Pi_i\}}) &= \Theta\left[p_\lambda - \sum_{l=0}^{k-1} \mathrm{tr}\big(\Pi_{\pi_{M,\lambda}(l)}[\lambda]\big)\right] \\
&\quad \cdot \Theta\left[-p_\lambda + \sum_{l=0}^{k} \mathrm{tr}\big(\Pi_{\pi_{M,\lambda}(l)}[\lambda]\big)\right]
\end{aligned} \tag{21}$$

where $\mu_{\mathcal{E}_j}(\lambda)$ is the uniform distribution over $\mathcal{E}_j$.

Roughly, all quantum states $|\psi\rangle$ with $\mathrm{tr}([j][\psi]) > \frac{d-1}{d}$ will have support on $\mathcal{E}_j$, and so will all overlap with each other. The permutation included in the definition of the measurements is constructed so that this shared support does not affect the prepare-and-measure-once statistics: the measurement ordered first by the permutation has a support which entirely contains $\mathcal{E}_j$. This model is not pairwise ψ-epistemic, nor is it never ψ-ontic. Note additionally that it was originally only defined for rank-1 projective measurements, but it works just as well for higher-rank projective measurements without modification.

This model is the first to fall to Theorem 1:

Theorem 2. *The LJBR model cannot represent state update under measurement in dimension $d \geq 3$.*

Proof. The general idea of the proof is to find a measurement which maps any two nonidentical states $|\alpha\rangle, |\beta\rangle$ to two again nonidentical states $\Pi|\alpha\rangle, \Pi|\beta\rangle$ which both have no support on any of the $\mathcal{E}_j$, and so are ontologically distinct.

Consider the preferred basis $|j\rangle$ of the LJBR model and the generalized x-basis defined by

$$|X_k\rangle = \frac{1}{\sqrt{d}} \sum_{j=0}^{d-1} \omega^{jk} |j\rangle, \qquad \omega = e^{2\pi i/d}. \tag{22}$$

These x-basis states have the property $\mathrm{tr}([X_j][k]) = \frac{1}{d}$ for all j, k. There must exist two elements $|X_{k_1}\rangle, |X_{k_2}\rangle$ of the x-basis such that $|\alpha\rangle, |\beta\rangle$ differ on that two-dimensional subspace or else $|\alpha\rangle, |\beta\rangle$ would be identical. Pick two such elements, and consider the projector $\Pi = [X_{k_1}] + [X_{k_2}]$. The quantum overlap of the post-measurement state $\Pi|\alpha\rangle$ with any basis vector $|j\rangle$ is, using the submultiplicativity of the trace,

$$\frac{\mathrm{tr}([j]\Pi[\alpha]\Pi)}{\mathrm{tr}(\Pi[\alpha])} \leq \mathrm{tr}([j]\Pi) = \frac{2}{d} \leq \frac{d-1}{d} \tag{23}$$

and the same is true for $\Pi|\beta\rangle$. As described above, only states with $\mathrm{tr}([j][\psi]) > \frac{d-1}{d}$ have overlap with any other states in the LJBR model, so the post-measurement states are ontologically distinct; by Theorem 1, $\xi(\Pi|\lambda) = 0$ for all $\lambda \in \mathcal{E}_j$ for all j since $|\alpha\rangle$ and $|\beta\rangle$ were arbitrary.

However, when measured in the context of the rest of the rank-1 x-basis projectors, Π will be ordered first by $\pi_{M,\lambda}$ for all λ since

$$\mathrm{tr}(\Pi[j]) = \frac{2}{d} > \frac{1}{d} \tag{24}$$

for all j. Thus, by the construction of the LJBR model, $\xi(\Pi|\lambda) = 1$ for all $\lambda \in \mathcal{E}_j$, resulting in a contradiction. □

4. ABCL models

In [36], Aaronson et. al. construct two ψ-epistemic models. The first, which we will call ABCL_0 , is very closely related to the LJBR model but is not identical; rather than continuous regions of quantum states which overlap, this model has exactly one pair of quantum states which are ontologically indistinct. However, it gains the feature that *any* two nonorthogonal quantum states can be chosen as the single pair that overlaps. With malice aforethought, we will call this defining pair $|\alpha\rangle, |\beta\rangle$.

The second, ABCL_1 , is a convex mixture (to be defined) of the ABCL_0 model constructed for all $|\alpha\rangle, |\beta\rangle$ and is intended to demonstrate the possibility of a pairwise ψ-epistemic model. This is the only known example of a pairwise ψ-epistemic model in $d \geq 3$, but it still is not never ψ-ontic [24, 54]. These models have come under criticism for their "unnaturalness," but we show here that their problems go deeper due to an inability to represent state-update.

We begin with ABCL_0 , defining a couple of helper functions like in the LJBR model. Rather than ordering measurements with respect to traces with a preferred basis, we use the defining states $|\alpha\rangle, |\beta\rangle$ and a function

$$g_{\alpha\beta}(\Pi) = \min\{\mathrm{tr}(\Pi[\alpha]), \mathrm{tr}(\Pi[\beta])\}. \tag{25}$$

We now define a new permutation σ_M [7] for each measurement M [24]:

$$g_{\alpha\beta}(M_{\sigma_M(0)}) \geq g_{\alpha\beta}(M_{\sigma_M(1)}) \geq \cdots \geq g_{\alpha\beta}(M_{\sigma_M(|M|-1)}). \tag{26}$$

---

[7] To be precise, we should label this with α, β as well to emphasize that it belongs to the model defined by that particular pair of states.

With this, we can specify the ABCL_0 model.

$$
\begin{aligned}
&\Lambda = \mathcal{PH}^{d-1} \times [0,1] \\
&\lambda = (|\lambda\rangle, p_\lambda) \\
&\mu(\lambda|P_\psi) = \begin{cases} \Big[\Theta(p_\lambda - \varepsilon)\delta(|\lambda\rangle - |\psi\rangle) \\ \quad + \dfrac{1}{2}\Theta(\varepsilon - p_\lambda)[\delta(|\lambda\rangle - |\alpha\rangle) + \delta(|\lambda\rangle - |\beta\rangle)]\Big] \\ \hfill \text{if } |\psi\rangle = |\alpha\rangle, |\beta\rangle \\ \delta(|\lambda\rangle - |\psi\rangle) \hfill \text{otherwise} \end{cases} \\
&\xi(k|\lambda, M_{\{\Pi_i\}}) = \Theta\left[p_\lambda - \sum_{j=0}^{k-1} \mathrm{tr}\big(\Pi_{\sigma_M(j)}[\lambda]\big)\right] \\
&\qquad\qquad \cdot \Theta\left[-p_\lambda + \sum_{j=0}^{k} \mathrm{tr}\big(\Pi_{\sigma_M(j)}[\lambda]\big)\right] \qquad (27)
\end{aligned}
$$

for $\varepsilon \le \frac{|\langle\alpha|\beta\rangle|}{d}$. Now the preparation distributions for $|\alpha\rangle$ and $|\beta\rangle$ overlap on $\{|\alpha\rangle, |\beta\rangle\} \times [0, \varepsilon]$; as in the LJBR model, the permutation in ξ ensures that the preparation change doesn't affect the prepare-and-measure-once statistics by making sure measurements whose support must contain this overlap region are ordered first. Once again, this model fails to meet the conditions required in order to faithfully represent state update:

Theorem 3. *The $ABCL_0$ model cannot represent state update under measurement in dimension $d \ge 3$.*

Proof. Call the two states defining the model $|\alpha\rangle, |\beta\rangle$. Let $\Pi = [\alpha] + [\gamma]$, where $|\gamma\rangle$ is some state such that

$$\langle\alpha|\gamma\rangle \ne 0 \quad \text{and} \quad 0 < |\langle\gamma|\beta\rangle|^2 < 1 - |\langle\alpha|\beta\rangle|^2. \qquad (28)$$

Under this measurement, $|\alpha\rangle$ maps to $|\alpha\rangle$ and $|\beta\rangle$ does not get mapped to either $|\alpha\rangle$ or $|\beta\rangle$. Thus the post-measurement states are ontologically distinct, so by Theorem 1, $\xi(\Pi|\lambda) = 0$ for all $\lambda \in \Delta_\alpha \cap \Delta_\beta$.

For the other half of the contradiction, note that $\mathrm{tr}(\Pi[\alpha]) = 1$ means $g_{\alpha\beta}(\Pi) = \mathrm{tr}(\Pi[\beta]) > 0$ (since $|\alpha\rangle, |\beta\rangle$ are nonorthogonal) and $g_{\alpha\beta}(\mathbb{I} - \Pi) = 1 - \mathrm{tr}(\Pi[\alpha]) = 0$, so Π is ordered first by σ_M when measured in the context $M = \{\Pi, \mathbb{I} - \Pi\}$. Thus $\xi(\Pi|\lambda) = 1$ for all $\lambda \in \Delta_\alpha \cap \Delta_\beta$, resulting in a contradiction. □

We outline the ABCL_1 model schematically and refer the reader to [24, 36] for details. Given two ontological models specified by Λ_1, μ_1, ξ_1 and Λ_2, μ_2, ξ_2 respectively, the authors define a convex combination of these models as Λ_3, μ_3, ξ_3 such that

$$
\begin{aligned}
\Lambda_3 &= \Lambda_1 \oplus \Lambda_2 \\
\mu_3 &= p\mu_1 + (1-p)\mu_2 \\
\xi_3 &= \xi_1 + \xi_2 \qquad (29)
\end{aligned}
$$

Here $p \in (0, 1)$ is some mixing parameter. If there's overlap between two states in either of models 1 or 2, then model 3 has overlap on these states. The ABCL_1 model is then defined essentially as a convex mixture of the ABCL_0 models for *all* pairs $|\alpha\rangle, |\beta\rangle$, taking care with respect to the uncountable size of this set.

In order to include state update in a convex combination of ontological models, the most obvious (and perhaps only) option is to specify

$$\eta_3 = \eta_1 + \eta_2. \qquad (30)$$

The failure of the ABCL_1 model to reproduce state update follows directly from the failure of the ABCL_0 model.

Theorem 4. *The $ABCL_1$ model cannot represent state update under measurement in dimension $d \ge 3$*

Proof. When we take a convex combination of models, we see that

$$
\begin{aligned}
\mathrm{Supp}(\xi_3(k|\cdot, M)) &= \mathrm{Supp}(\xi_1(k|\cdot, M)) \cup \mathrm{Supp}(\xi_2(k|\cdot, M)) \\
\mathrm{Supp}(\eta_3(\cdot|k, \lambda, M)) &= \mathrm{Supp}(\eta_1(\cdot|k, \lambda, M)) \\
&\qquad \cup \mathrm{Supp}(\eta_2(\cdot|k, \lambda, M))
\end{aligned}
$$

Thus if either of models 1 or 2 violates Theorem 1, model 3 must violate it as well. Since all of the ABCL_0 models being mixed violate Theorem 1, so must ABCL_1 . □

5. A note on transformations

As Leifer notes, transformations also play a role in restricting the structure of ψ-epistemic ontological models [24, Section 8.1]. If an ontological model successfully represents all unitary transformations, then $|\langle\psi|\phi\rangle| = |\langle\psi'|\phi'\rangle|$ implies that $|\psi\rangle, |\phi\rangle$ are ontologically distinct if and only if $|\psi'\rangle, |\phi'\rangle$ are ontologically distinct. If the model also includes all CPTP maps, then $|\langle\psi|\phi\rangle| \ge |\langle\psi'|\phi'\rangle|$ implies that $|\psi'\rangle, |\phi'\rangle$ are ontologically distinct if $|\psi\rangle, |\phi\rangle$ are.

It immediately follows that the LJBR and ABCL_0 models cannot faithfully represent unitary transformations. In each model there exist quantum states which are ontologically distinct from every other state; pick one of these states, and it is easy to find examples of pairs of ontologically distinct states with any inner product.

However, transformations cannot necessarily rule out the ABCL_1 model: since it is pairwise ψ-epistemic, the unitary condition could *in principle* be satisfied. That said, the transformation rule would be complicated because it would have to map *between* models that are mixed together, so it is certainly an open question whether this is actually possible.

C. Models of subtheories

Although we have dealt so far with models that include the full quantum set of preparations, transformations, and measurements, there is the possibility that we can retain ψ-epistemic models of subtheories. It turns out that although the stabilizer subtheory can be represented by a ψ-epistemic model, the more general Kitchen Sink model which models any finite subtheory cannot in general represent state update under measurement.

1. *Kitchen Sink model*

The Kitchen Sink model is a ψ-epistemic ontological model for any finite subtheory of quantum theory [37, Section IIIC]. Given a finite set of projective measurements $\mathcal{M} = \{M^{(i)}\}$, we choose our ontic states to be a list of measurement outcomes $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_{|\mathcal{M}|})$. That is, $\lambda_i = k$ means that if $M^{(i)} = \{\Pi_j^{(i)}\}$ is measured on the ontic state λ, the outcome Π_k will occur with certainty. For a system of dimension d, the maximum number of projectors in any given measurement is d, so we pad all of our measurements with 0s until they have d elements. The Kitchen Sink model is then defined by

$$\begin{aligned} \Lambda &= \mathbb{Z}_d^{|\mathcal{M}|} \\ \mu(\lambda|\psi) &= \prod_{i=1}^{|\mathcal{M}|} \operatorname{tr}\left(\Pi_{\lambda_i}^{(i)}[\psi]\right) \\ \xi(k|\lambda, M^{(i)}) &= \delta(k, \lambda_i) \end{aligned} \tag{31}$$

The Kitchen Sink is pairwise ψ-epistemic for any subtheory, and also never ψ-ontic if we include all pure states in our subtheory. Transformations can additionally be modeled under the assumption of a closed subtheory.

We can only rule out the Kitchen Sink model for certain subtheories, as it is easy to construct subtheories with trivial update rules (e.g. by only including rank-1 measurements). That said, our requirements are few and are satisfied by the multi-qupit stabilizer subtheory, arguably the most important subtheory of quantum theory. Specifically, we only need to include two states $|\alpha\rangle, |\beta\rangle$ and two measurements $M^{(1)} = \{\Pi, \mathbb{I}-\Pi\}, M^{(2)}$ satisfying

$$\langle\alpha|\beta\rangle \neq 0 \tag{32}$$

$$\langle\alpha|\Pi|\alpha\rangle \neq 0 \tag{33}$$

$$\langle\beta|\Pi|\beta\rangle \neq 0 \tag{34}$$

$$M^{(2)} \text{ distinguishes } \Pi|\alpha\rangle \text{ and } \Pi|\beta\rangle \tag{35}$$

The first is required because we don't expect orthogonal states to be ontologically indistinguishable. The next two stipulate that there is a nonzero chance of obtaining an outcome Π when measuring $|\alpha\rangle$ and $|\beta\rangle$, so that its support overlaps with their support. The last condition implies that the post-measurement states are ontologically distinct [24].

Theorem 5. *For any finite subtheory containing states and measurements satisfying the conditions given in Eqs. 32–35, the Kitchen sink model cannot model state update under measurement.*

Proof. Theorem 1 implies that, if Π maps $|\alpha\rangle, |\beta\rangle$ to ontologically distinct states, then

$$\int_\Lambda \mathrm{d}\lambda\, \xi(\Pi|\lambda)\mu(\lambda|\alpha)\mu(\lambda|\beta) = 0. \tag{36}$$

We evaluate this quantity for the Kitchen Sink's response functions and preparation distributions, using the states and measurement $M^{(1)}$ satisfying Eqs. 32–35:

$$\begin{aligned} &\int_\Lambda \mathrm{d}\lambda\, \xi(k=0|\lambda, M^{(1)})\mu(\lambda|\alpha)\mu(\lambda|\beta) \\ &\quad= \sum_{\lambda\in\mathbb{Z}_r^{|\mathcal{M}|}} \delta(0,\lambda_1) \prod_{j=1}^{|\mathcal{M}|} \operatorname{tr}\left(M_{\lambda_j}^{(j)}[\alpha]\right) \operatorname{tr}\left(M_{\lambda_j}^{(j)}[\beta]\right) \\ &\quad= \sum_{\lambda_1\in\mathbb{Z}_r} \delta(0,\lambda_1) \operatorname{tr}\left(M_{\lambda_1}^{(1)}[\alpha]\right) \operatorname{tr}\left(M_{\lambda_1}^{(1)}[\beta]\right) \\ &\qquad\quad \cdot \prod_{j=2}^{|\mathcal{M}|} \sum_{l\in\mathbb{Z}_r} \operatorname{tr}\left(M_l^{(j)}[\alpha]\right) \operatorname{tr}\left(M_l^{(j)}[\beta]\right) \\ &\quad= \operatorname{tr}\left(M_0^{(1)}[\alpha]\right) \operatorname{tr}\left(M_0^{(1)}[\beta]\right) \\ &\qquad\quad \cdot \prod_{j=2}^{|\mathcal{M}|} \sum_{l\in\mathbb{Z}_r} \operatorname{tr}\left(M_l^{(j)}[\alpha]\right) \operatorname{tr}\left(M_l^{(j)}[\beta]\right) \end{aligned} \tag{37}$$

This final expression will be zero if and only if at least one of its factors is 0. The first two factors are nonzero by Eqs. 33 and 34. The rest of the factors are nonzero due to Eq. 32 and the completeness condition on the measurements. Thus Eq. 37 is nonzero and we have a contradiction, so state update cannot be represented faithfully. □

This demonstrates that Theorem 1 can create trouble even in subtheories. In particular, the stabilizer subtheory satisfies the requirements in Eqs. 32–35, so it cannot be modeled by the Kitchen Sink. That said, we can show that the stabilizer subtheory still supports a ψ-epistemic interpretation using other models.

2. *Qupit stabilizer subtheory*

We begin with the straightforward case of n p-dimensional systems, for p an odd prime. In this case, the stabilizer subtheory has an ontological model given by the discrete Wigner function [33, 55] (see Appendix B for definitions of the stabilizer subtheory and the phase-point operaters A_λ):

$$\begin{aligned} \Lambda &= \mathbb{Z}_p^n \times \mathbb{Z}_p^n \\ \mu(\lambda|P_\psi) &= \frac{1}{p^n} \operatorname{tr}(A_\lambda[\psi]) \\ \xi(k|\lambda, M_{\{\Pi_i\}}) &= \operatorname{tr}(\Pi_k A_\lambda) \end{aligned} \tag{38}$$

As the Wigner function is a quasi-probability distribution [56], it necessarily takes on negative values if we try to model the full quantum theory. If, however, we restrict to modeling preparations, transformations, and measurements in the qu$p$it stabilizer subtheory, then the representation is positive and it forms a well-defined ontological model [33, 57]. It is both pairwise ψ-epistemic and never ψ-ontic.

The stabilizer subtheory presents a challenge in that measuring a single qu$p$it is described by a rank-p^{n-1} measurement, which may run into trouble due to Theorem 1. In particular, there are many examples in the stabilizer subtheory of the type of measurements that broke the Kitchen Sink model. Nonetheless, we can specify the update rule

$$\eta(\lambda'|k, \lambda, M_{\{\Pi_i\}}) = \frac{1}{p^n} \frac{\mathrm{tr}(A_\lambda \Pi_k A_{\lambda'} \Pi_k)}{\mathrm{tr}(\Pi_k A_\lambda)}. \tag{39}$$

Note the normalization factor in η which makes clear that η is only defined in the support of ξ. The fact that this successfully reproduces quantum statistics follows from the fact that post-selected measurement is a completely positive map, and this is how completely positive maps are represented in quasi-probability representations [56]. η is always positive for the stabilizer subtheory, though we do not include the proof here.

For the case $p = 2$, Lillystone and Emerson construct a ψ-epistemic model of the n-qubit stabilizer formalism that successfully represents state update under measurement [48]. This model starts from the Kitchen Sink model and augments Λ so that the problematic overlaps of the Kitchen sink are removed. This model is not pairwise ψ-epistemic, but a modified version (see appendix of [48]) is never-ψ-ontic. We don't present the construction here because it is significantly more convoluted than the model above for $p \geq 3$. This reflects the often-observed ill-behaved nature of the qubit stabilizer subtheory.

Tangentially, if we extend the Wigner function to the full quantum theory, we get negatively represented state update, as expected. One consequence of this is that some state updates can't be normalized, so the Wigner function state update must include a renormalization step not allowed in ontological models or quasiprobability representations. Although further discussion of state update under measurement in quasi-probability representations is beyond the scope of this paper, we note that Theorem 1 does not hold for quasi-probability representations so this could be one potential direction for related future work.

V. DISCUSSION

We have demonstrated that state update under measurement poses a serious challenge to ψ-epistemic interpretations of quantum theory in the ontological models framework: all currently known ψ-epistemic models for full quantum theory in $d \geq 3$ cannot faithfully represent state update. This runs in direct contrast to the prevailing view that ψ-epistemic models provide a compelling explanation of state update.

There are a number of remaining open questions. Most pressingly, we have re-opened the possibility of proving a general ψ-onticity result without additional assumptions—will the methods of this paper be useful in doing so?

On the one hand, the proofs above do not rule out the possibility of extending the ontology of the broken models in order to represent state update under measurement while still retaining the epistemicity of the model. This is exactly the route taken in [48] for the n-qubit stabilizer subtheory. Granted, the n-qubit stabilizer subtheory was brought within an inch of ψ-onticity by this process, so it seems unlikely that a similar technique will work for the full quantum theory.

In the other direction, we've shown that consideration of state update puts powerful constraints on the structure of ψ-epistemic models. These restrictions would ideally lead to a categorical statement like "ψ-epistemic models cannot represent state-update," but there are challenges to achieving this conclusion. In particular, we note that all of the ψ-epistemic models that we considered share the property of outcome determinism, which means that Theorem 1 may be less trouble in non-outcome-deterministic models. At the very least, any no-go theorem will have to include measurements, states, and/or transformations from outside the stabilizer subtheory since we have shown that ψ-epistemic models for this subtheory exist.

What import does our result have for the general interpretational project of quantum theory? First of all, we have demonstrated that ψ-epistemicists have yet another challenge to overcome: a successful explanation of state-update. This is in contrast with the usual claim that this arena is one where epistemic interpretations have an advantage over ontic interpretations. As we emphasized in the introduction, our results only strictly apply to interpretations that can be described by the ontological models formalism, but there may be a qualitative message for epistemic and doxastic interpretations that are outside this formalism as well.

VI. ACKNOWLEDGEMENTS

This research was supported by the Canada First Research Excellence Fund and the NSERC Discovery program. J.R. acknowledges funding from an Ontario Graduate Scholarship. The authors would like to thank Matt Leifer and David Schmid for productive discussions.

[1] D. Bohm, Physical Review **85**, 166 (1952).
[2] H. Everett, Reviews of Modern Physics **29**, 454 (1957).

[3] P. Dirac, *The Principles of Quantum Mechanics* (Oxford University Press, Oxford, 1958).
[4] G. C. Ghirardi, A. Rimini, and T. Weber, Physical Review D **34**, 470 (1986).
[5] P. Pearle, Physical Review A **39**, 2277 (1989).
[6] A. Valentini, *On the Pilot-Wave Theory of Classical, Quantum, and Subquantum Physics*, Ph.D. Thesis (1991).
[7] D. Bohm and B. J. Hiley, *The Undivided Universe* (Routledge, New York, 1993).
[8] P. R. Holland, *The Quantum Theory of Motion* (Cambridge University Press, New York, 1993).
[9] E. G. Beltrametti and S. Bugajski, Journal of Physics A: Mathematical and General **28**, 3329 (1995).
[10] D. Dieks and P. E. Vermaas, *The Modal Interpretation of Quantum Mechanics* (Springer Netherlands, Dordrecht, 1998) oCLC: 851391542.
[11] P. Pearle, in *Open Systems and Measurement in Relativistic Quantum Theory*, Vol. 526, edited by H.-P. Breuer and F. Petruccione (Springer Berlin Heidelberg, 1999) pp. 195–234.
[12] J. A. Barrett, *The Quantum Mechanics of Minds and Worlds* (Oxford University Press, Oxford, 1999).
[13] O. Lombardi and D. Dieks, The Stanford Encyclopedia of Philosophy.
[14] A. Bassi, K. Lochan, S. Satin, T. P. Singh, and H. Ulbricht, Reviews of Modern Physics **85**, 471 (2013).
[15] A. Einstein, B. Podolsky, and N. Rosen, Physical Review **47**, 777 (1935).
[16] S. Kochen and E. Specker, Indiana University Mathematics Journal **17**, 59 (1967).
[17] L. E. Ballentine, Reviews of Modern Physics **42**, 358 (1970).
[18] D. Howard, Studies in History and Philosophy of Science Part A **16**, 171 (1985).
[19] L. E. Ballentine, Y. Yang, and J. P. Zibin, Physical Review A **50**, 2854 (1994).
[20] J. Emerson, arXiv:quant-ph/0211035 (2002), arXiv:quant-ph/0211035.
[21] T. Rudolph, Arxiv (2006), arXiv:quant-ph/0608120.
[22] R. W. Spekkens, Physical Review A **75** (2007), 10.1103/PhysRevA.75.032110, arXiv:quant-ph/0401052.
[23] N. Harrigan and R. W. Spekkens, Foundations of Physics **40**, 125 (2010), arXiv:0706.2661.
[24] M. S. Leifer, Quanta **3**, 67 (2014), arXiv:1409.1570.
[25] A. Cabello, arXiv:1509.04711 [quant-ph] , 138 (2017), arXiv:1509.04711 [quant-ph].
[26] R. Schack, "Participatory realism," (2016).
[27] Werner Heisenberg, *Physics & Philosophy* (1958).
[28] C. A. Fuchs, N. D. Mermin, and R. Schack, American Journal of Physics **82**, 749 (2014).
[29] C. Rovelli, International Journal of Theoretical Physics **35**, 1637 (1996).
[30] R. Healey, arXiv:1008.3896 [quant-ph] (2010), arXiv:1008.3896 [quant-ph].
[31] J. Bub, arXiv:1711.01604 [physics, physics:quant-ph] (2017), arXiv:1711.01604 [physics, physics:quant-ph].
[32] C. A. Fuchs, arXiv:quant-ph/0205039 (2002), arXiv:quant-ph/0205039.
[33] D. Gross, Journal of Mathematical Physics **47**, 122107 (2006), arXiv:quant-ph/0602001.
[34] A. Montina, Physical Review Letters **108** (2012), 10.1103/PhysRevLett.108.160501, arXiv:1108.5138.
[35] P. G. Lewis, D. Jennings, J. Barrett, and T. Rudolph, Physical Review Letters **109** (2012), 10.1103/PhysRevLett.109.150404.
[36] S. Aaronson, A. Bouland, L. Chua, and G. Lowther, Physical Review A **88** (2013), 10.1103/PhysRevA.88.032111.
[37] N. Harrigan, T. Rudolph, and S. Aaronson, arXiv:0709.1149 [quant-ph] (2007), arXiv:0709.1149 [quant-ph].
[38] J. J. Wallman and S. D. Bartlett, Physical Review A **85** (2012), 10.1103/PhysRevA.85.062121, arXiv:1203.2652.
[39] J. S. Bell, Reviews of Modern Physics **38**, 447 (1966).
[40] M. F. Pusey, J. Barrett, and T. Rudolph, Nature Physics **8**, 476 (2012), arXiv:1111.3328.
[41] R. Colbeck and R. Renner, New Journal of Physics **19**, 013016 (2017).
[42] L. Hardy, International Journal of Modern Physics B **27**, 1345012 (2013).
[43] S. Mansfield, Physical Review A **94** (2016), 10.1103/PhysRevA.94.042124.
[44] R. W. Spekkens, Physical Review A **71** (2005), 10.1103/PhysRevA.71.052108, arXiv:quant-ph/0406166.
[45] N. Harrigan and T. Rudolph, Arxiv (2007), arXiv:0709.4266.
[46] J. Emerson, D. Serbin, C. Sutherland, and V. Veitch, arXiv:1312.1345 [quant-ph] (2013), arXiv:1312.1345 [quant-ph].
[47] L. Catani and D. E. Browne, New Journal of Physics **19**, 073035 (2017), arXiv:1701.07801.
[48] P. Lillystone and J. Emerson, forthcoming.
[49] A. Karanjai, J. J. Wallman, and S. D. Bartlett, arXiv:1802.07744 [quant-ph] (2018), arXiv:1802.07744 [quant-ph].
[50] A. J. Leggett and A. Garg, Physical Review Letters **54**, 857 (1985).
[51] C. Emary, N. Lambert, and F. Nori, Reports on Progress in Physics **77**, 016001 (2014), arXiv:1304.5133.
[52] J. J. Sakurai and J. Napolitano, *Modern Quantum Mechanics*, 2nd ed. (Addison-Wesley, Boston, 2011).
[53] D. Gottesman, arXiv:0904.2557 [quant-ph] (2009), arXiv:0904.2557 [quant-ph].
[54] A. Montina, "Comment on math overflow question: 'psi-epistemic theories' in 3 or more dimensions," (2012).
[55] V. Veitch, *Negative Quasi-Probability in the Context of Quantum Computation*, Ph.D. thesis, University of Waterloo.
[56] C. Ferrie, Reports on Progress in Physics **74**, 116001 (2011), arXiv:1010.2701.
[57] D. Gross, Applied Physics B **86**, 367 (2007), arXiv:quant-ph/0702004.
[58] N. Barnett and J. P. Crutchfield, Journal of Statistical Physics **161**, 404 (2015), arXiv:1412.2690.
[59] R. A. Howard and J. E. Matheson, Decision Analysis **2**, 127 (2005).
[60] M. S. Leifer and M. F. Pusey, Proc. R. Soc. A **473**, 20160607 (2017).

Appendix A: Ontological models as hidden Markov models of stochastic channels

In the context of state update under measurement, it is illuminating to motivate the definition of an ontological model from the point of view of the hidden Markov models (HMM) literature. We do this in order to (a) provide a rigorous treatment of multiple-measurement scenarios presented informally in Section II A and (b) clarify the assumptions that define the ontological models framework.

We picture a quantum circuit as a memoryful stochastic channel (Fig. 2). The channel that we often discuss with regards to a quantum circuit is the (quantum) channel that takes the input quantum state and maps it to the output quantum state. For present purposes, we will instead think of it as a channel from the experimenter to individual measurement outcomes used repeatedly at each time step. Pictorially, one might think of this as 'rotating the channel ninety degrees' in a circuit diagram.

The input string $\overleftrightarrow{a_0} = \ldots a_{-2}a_{-1}a_0a_1a_2\ldots$ of the channel is the experimenter's choice of action, which we take to be an operational preparation, transformation, or measurement procedure. The output string $\overleftrightarrow{k_0}$ of the channel reports either the results of measurements or a trivial output for preparations and transformations. The subscript 0 indicates in both cases the time that we take as an origin/reference point. The channel is then described by the conditional probability distribution $\Pr(\overleftrightarrow{k_0}|\overleftrightarrow{a_0})$. Following [58], we denote substrings with $a_{t:t+L} = a_t a_{t+1}\ldots a_{t+L-1}$, and also define the past $\overleftarrow{a_t} = a_{-\infty:t}$ and future $\overrightarrow{a_t} = a_{t:\infty}$. We can now define two properties of stochastic channels:

Definition 4 (Stationary)**.** *A* stationary *channel is one that has time-translation symmetry, so statistics are not affected by our choice of time-origin:*

$$\Pr(k_{t:t+L}|\overleftrightarrow{a_t}) = \Pr(k_{0:L}|\overleftrightarrow{a_o}) \qquad \textit{and}$$
$$\Pr(\overleftrightarrow{k_t}|\overleftrightarrow{a_t}) = \Pr(\overleftrightarrow{k_0}|\overleftrightarrow{a_0}) \qquad \forall t, L, \overleftrightarrow{a}. \tag{A1}$$

Definition 5 (Causal)**.** *A* causal *channel is one for which a finite output substring depends only on input symbols in its past:*

$$\Pr(k_{t:t+L}|\overleftrightarrow{a}) = \Pr(k_{t:t+L}|\overleftarrow{a}_{t+L}). \qquad \forall t, L, \overleftrightarrow{a} \tag{A2}$$

It is shown in [58] that a channel satisfying these two properties can be specified entirely by the single-symbol recurrence relation

$$\Pr(k_0|a_0, \overleftarrow{a_0}, \overleftarrow{k_0}). \tag{A3}$$

Note that this does not imply a Markov process, since it depends in general on the entire histories $\overleftarrow{a_0}$ and $\overleftarrow{k_0}$. All we mean by single-symbol is that we are not specifying the probabilities over the whole future, just a single output symbol. We now construct an HMM as follows:

Definition 6 (Hidden Markov Model)**.** *A* hidden Markov model *(HMM) of a stationary, causal channel is specified by an additional random variable λ taking values in a state space Λ. It is given a joint probability distribution over $\overleftrightarrow{\lambda_0}, \overleftrightarrow{a_0}, \overleftrightarrow{k_0}$ so that the recurrence relation above (Eq. A3) becomes*

$$\Pr(k_0, \lambda_1|a_0, \lambda_0, \overleftarrow{a}_0, \overleftarrow{k}_0, \overleftarrow{\lambda}_0) = \Pr(k_0, \lambda_1|a_0, \lambda_0). \tag{A4}$$

In other words, the state λ renders the future conditionally independent of the past and induces a Markov process over the state space Λ that mediates the channel statistics. An influence diagram [59] of a stationary, causal channel is shown in Figure 3 before and after the specification of an HMM.

To see that specification of an HMM as in Eq. A4 is equivalent to the definition of an ontological model given in Section II A, we first note that generally we don't think of preparations and transformations having output; to account for this, we stipulate that they give a trivial, deterministic output $k_0 = 0$. We then factor the probability distribution from Eq. A4 and look separately at the cases where a_0 is a preparation, transformation, or measurement:

$$\begin{aligned}
&\Pr(k_0, \lambda_1|\lambda_0, a_0) \\
&\quad = \Pr(\lambda_1|k_0, \lambda_0, a_0)\Pr(k_0|\lambda_0, a_0) \\
&\quad = \begin{cases} \mu(\lambda_1|P)\delta_{k_0,0} & a_0 = P \in \mathcal{P} \\ \Gamma(\lambda_1|\lambda_0, T)\delta_{k_0,0} & a_0 = T \in \mathcal{T} \\ \eta(\lambda_1|k_0, \lambda_0, M)\xi(k_0|\lambda_0, M) & a_0 = M \in \mathcal{M} \end{cases}
\end{aligned} \tag{A5}$$

where δ is the Kronecker delta. The state update map η emerges naturally from this perspective of a quantum experiment as a stochastic process, and here we see another reason why it is only defined in the support of ξ. If

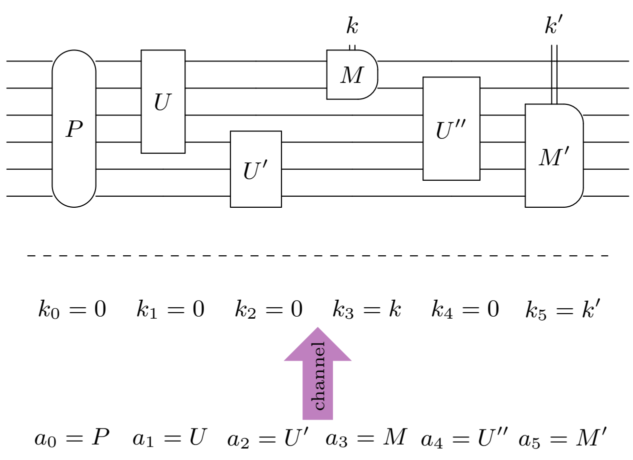

Figure 2: A quantum circuit can be pictured as a stochastic channel, as described in the text. The inputs to the channel a_t are the choice of operation, and the outputs k_t report the results of measurements.

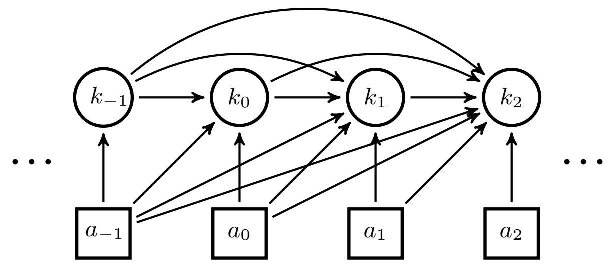

(a) A causal and stationary channel

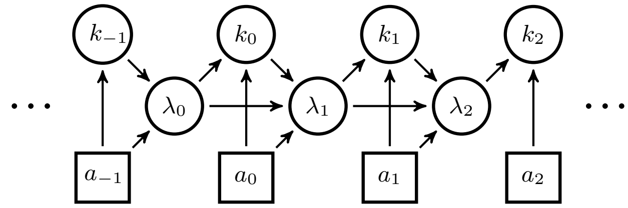

(b) A hidden Markov model of the above channel

Figure 3: Influence diagrams [59] for stochastic channels. The boxes represent choices made by the experimenter, circles represent random variables, and arrows represent a possible causal influence. Note that no arrows point backwards in time, and that in the hidden Markov model the state λ_t mediates all causal influences through time.

we take the joint distribution $\Pr(k_0, \lambda_1|\lambda_0, a_0)$ to be the more fundamental object, then it is clear we can obtain ξ directly by marginalization

$$\xi(k_0|\lambda_0, M) = \int_\Lambda \mathrm{d}\lambda_1 \Pr(k_0, \lambda_1|\lambda_0, M) \tag{A6}$$

which is always well defined, and then find η by rearranging Eq. A5:

$$\eta(\lambda_1|k_0, \lambda_0, M) = \frac{\Pr(k_0, \lambda_1|\lambda_0, M)}{\xi(k_0|\lambda_0, M)} \tag{A7}$$

Thus clearly $\eta(\lambda_1|k_0, \lambda_0, M)$ is only well-defined when $\xi(k_0|\lambda_0, M) \neq 0$.

Definitions 4–6 constitute an equivalent formulation of the ontological models formalism. The assumptions of this construction can be broken down as follows: (a) quantum theory is described by a stochastic channel, (b) this channel is stationary, (c) it is causal, and (d) we assign the system a state which acts as an HMM of the channel. The authors of [60] identify (c) and (d), calling them non-retrocausality and λ-mediation, respectively.

Appendix B: A brief introduction to the stabilizer subtheory

We focus here on the stabilizer subtheory for n qu$p$its, where p is a prime. For a more detailed exposition, we refer the reader to [33, 55].

a. Mathematical objects The generalizations of the X and Z operators to a single qu$p$it are defined by their action on the standard basis $\{|j\rangle\}$ for $j = 0, 1, \ldots, p-1$:

$$X|j\rangle = |j+1\rangle \tag{B1}$$

$$Z|j\rangle = \omega^j |j\rangle \tag{B2}$$

$$\omega = e^{2\pi i/p} \tag{B3}$$

All integer arithmetic is done $\bmod p$. Then the full set of generalized Pauli operators on n qu$p$its is given by

$$T_{(\mathbf{x},\mathbf{z})} = \begin{cases} \bigotimes_{j=0}^{n-1} X^{x_j} Z^{z_j} & p = 2 \\ \bigotimes_{j=0}^{n-1} \omega^{x_j z_j/2} X^{x_j} Z^{z_j} & p > 2 \end{cases} \tag{B4}$$

$$\text{for } \mathbf{x} = (x_0, x_1, \ldots, x_{n-1}) \in \mathbb{Z}_p^n$$
$$\text{and } \mathbf{z} = (z_0, z_1, \ldots, z_{n-1}) \in \mathbb{Z}_p^n$$

We also define the symplectic inner product as

$$[(\mathbf{x}, \mathbf{z}), (\mathbf{x}', \mathbf{z}')] = \mathbf{z} \cdot \mathbf{x}' - \mathbf{x} \cdot \mathbf{z}'. \tag{B5}$$

Finally, the phase-point operators A_λ, for $\lambda = (\mathbf{x}, \mathbf{z}) \in \mathbb{Z}_p^n \times \mathbb{Z}_p^n$, are a symplectic Fourier transform of the Pauli operators:

$$A_\lambda = \frac{1}{p^n} \sum_{\lambda' \in \mathbb{Z}_p^n \times \mathbb{Z}_p^n} \omega^{[\lambda, \lambda']} T_{\lambda'} \tag{B6}$$

b. The stabilizer subtheory A stabilizer group S is a set of p^n mutually commuting Pauli operators, which can be specified by a set of n generators. There is a unique state (up to global phase) which is an eigenvector of all of these operators with eigenvalue $+1$; we say that S *stabilizes* this state. For example, for two qubits, the Bell state

$$|\Psi\rangle = |00\rangle + |11\rangle \tag{B7}$$

is stabilized by

$$S_\Psi = \langle Z_1 Z_2, X_1 X_2 \rangle \tag{B8}$$

$$= \{\mathbb{I}, Z_1 Z_2, X_1 X_2, -Y_1 Y_2\}. \tag{B9}$$

Here a subscript indicates on which qubit the operator is acting, e.g. $Z_1 = Z \otimes \mathbb{I}$ describes Z acting on the first qubit. A stabilizer state, then, is a state which is stabilized by a group of p^n Pauli operators.

Stabilizer measurements are simply measurements of the Pauli operators. Note that since each Pauli operator has p eigenvalues, these amount to a measurement of p projectors, each with rank p^{n-1}. Lower-rank projectors can be constructed by performing commuting Pauli measurements sequentially.

Finally, the transformations of the stabilizer subtheory are called Clifford transformations. These are the transformations that map the set of Pauli operators to itself, up to a global phase. In other words, it is the normalizer of the Pauli group.

The stabilizer subtheory thus consists of preparations corresponding the set of stabilizer states, measurements of Pauli observables, and Clifford transformations, along with convex combinations thereof.