arXiv:1708.00137v2 [quant-ph] 14 May 2018

Causation does not explain contextuality

Sally Shrapnel and Fabio Costa

Centre for Engineered Quantum Systems, School of Mathematics and Physics, The University of Queensland, St Lucia, QLD 4072, Australia

15th May 2018

Realist interpretations of quantum mechanics presuppose the existence of elements of reality that are independent of the actions used to reveal them. Such a view is challenged by several no-go theorems that show quantum correlations cannot be explained by non-contextual ontological models, where physical properties are assumed to exist prior to and independently of the act of measurement. However, all such contextuality proofs assume a traditional notion of causal structure, where causal influence flows from past to future according to ordinary dynamical laws. This leaves open the question of whether the apparent contextuality of quantum mechanics is simply the signature of some exotic causal structure, where the future might affect the past or distant systems might get correlated due to non-local constraints. Here we show that quantum predictions require a deeper form of contextuality: even allowing for arbitrary causal structure, no model can explain quantum correlations from non-contextual ontological properties of the world, be they initial states, dynamical laws, or global constraints.

Introduction

The appeal of an operational physical theory is that it makes as few unwarranted assumptions about nature as possible. One simply assigns probabilities to experimental outcomes, conditioned on the list of experimental procedures required to realise these outcomes. Ideally, such operational theories are *minimal*: procedures that cannot be statistically discriminated are given the same representation in the theory. Quantum mechanics is an example of such a minimal operational theory: all the statistically significant information about the preparation procedure is contained in the quantum state, and the probability of an event (labelled by a Positive Operator Valued Measure (POVM) element) does not depend on any other information regarding the manner in which the measurement was achieved (such as the full POVM). However, one of the most debated questions in the foundations of the theory is whether one can go beyond this statistical level and also provide an *ontological* description of some actual state of affairs that occurs during each run of an experiment. That is, a statement about the world that tells us what is responsible for the observed experimental outcomes.

The task of providing such an ontological model for quantum theory has proven to be exceedingly difficult. A plethora of no-go theorems exists that describe the various natural assumptions one must forgo in order to produce an ontological model that accords

Sally Shrapnel: s.shrapnel@uq.edu.au
Fabio Costa: f.costa@uq.edu.au

with experiment. One such caveat is *non-contextuality*. Ultimately an *apriori* assumption, non-contextual theories posit the existence of physical properties that do not depend on the way they are measured. There is a large literature discussing the various ways one may wish to cash out this notion more precisely. Broadly speaking, non-contextuality no-go theorems fall into two distinct categories. Kochen-Specker style proofs show that quantum measurements cannot be regarded as deterministically uncovering pre-existing, or ontic, properties of systems [1–3]. Spekkens style proofs, on the other hand, show that one cannot explain quantum statistics via ontological properties that mirror the context-independence seen at the operational level [4–8]. While both approaches are well justified and have led to interesting and relevant results, our own definition of non-contextuality is more closely related to the latter. This particular view of non-contextuality can more broadly be seen as an analogue of the no fine-tuning argument from causal modelling [9], an analogue of Leibniz's principle of the Identity of Indiscernibles [4,6], and a methodological assumption akin to Occam's razor.

Non-contextuality no-go theorems are not merely of foundational interest but can also serve as security proofs for a range of simple cryptographic scenarios [10, 11], can herald a quantum advantage for computation [12], and also for state discrimination [8]. Such results, however, require the assumption of a fixed background causal structure; at the very minimum, a single causal arrow from preparation to measurement. This leaves open the question of whether one can produce a non-contextual ontological model by allowing for a suitably exotic causal structure. Some authors attempt to explain quantum correlations by positing backwards-in-time causal influences [13–19], while others claim it is the existence of non-local constraints that does the explanatory work [20,21]. The rationale in both cases is that non-contextuality could emerge naturally in such models: physical properties might well be "real" and "counterfactually definite", but depend on future or distant measurements because of some physically motivated—although radically novel—causal influence. Such proposals do not fit neatly within the classical causal modelling framework, and so are not ruled out by recent work in this direction [9,22], nor by any of the existing no-go theorems.

In this paper, we characterise a new ontological models framework to prove that even if one allows for *arbitrary* causal structure, ontological models of quantum experiments are necessarily contextual. Crucially, what is contextual is not just the traditional notion of "state", but any supposedly objective feature of the theory, such as a dynamical law or boundary condition. Our finding suggests that *any* model that posits unusual causal relations in the hope of saving "reality" will necessarily be contextual. Finally, this work also represents a possible approach to how we ought to think of the generalised quantum processes of recent work [23–37]. It is clear that any ontological reading of such processes will have to contend with the spectre of contextuality.

The paper is organised as follows. In section 1 we present the traditional ontological models framework and clarify the rationale behind retrocausal explanations of quantum statistics. In section 2 we introduce and justify the four primitive elements required to define our operational model: local regions, local controllables, outcomes and an environment. In section 3 we define the three classes of operationally indistinguishable elements: events, instruments and processes. In Section 4 we characterise instrument and process non-contextuality according to these equivalence classes, and provide a generalised framework for a non-contextual ontological model. As this is the conceptual heart of our result, in Section 5 we clarify the scope and applicability of this framework via three examples. Using standard quantum theory and results from previous work [29, 37], in section 6 we characterise an operational model that accords with the experimental predictions of quantum theory. Section 7 puts these elements together to prove that one cannot produce an on-

tological model that is both process and instrument non-contextual and accords with the predictions of quantum theory. In Section 8 we consider the constraints imposed on ontological models when one only assumes instrument non-contextuality. We finish with a discussion.

1 An introduction to ontological models and retrocausal approaches.

The ontological models framework assumes that systems possess well defined properties at all times [18, 38, 39]. The starting point is the very general claim that all experiments can be modelled operationally as sets of preparations, followed by transformations, followed by measurements, all performed upon some physical system. The set of all possible preparations, transformations and measurements is regarded as capturing the entire possibility space of any experiment and can be associated with the operational predictions of a particular theory. For example, an experiment can involve choices of possible preparation settings (labelled by the random variable P) and choices of possible measurement settings (M) with associated outcomes (a).[1] An operational model then predicts probabilities for outcomes for all possible combinations of preparations and measurements:

$$\forall a, M, P: \; p(a|M, P). \tag{1}$$

Such probabilistic predictions should coincide with the operational predictions of the theory in question. For example, in the case of quantum theory each preparation choice is modelled as a density operator (ρ_P) on a Hilbert space associated to a quantum system ($\mathcal{H}_A$). Similarly, each measurement choice M is associated with a positive operator valued measure $\{E_{a|M}\}$, whose elements correspond to particular outcomes a. The probabilities predicted by the theory are:

$$p(a|P, M) = \mathrm{Tr}(\rho_P E_{a|M}). \tag{2}$$

An *ontological extension* of such an operational model further assumes that the system possesses well defined ontological properties between the time of preparation and measurement. Such properties are collectively known as the "ontic state" and typically denoted by λ. In the ontological models framework each preparation procedure P is presumed to select a particular ontic state λ according to a fixed probability distribution: $\mu_P(\lambda)$, and each measurement choice is presumed to output a particular outcome according to a fixed response function: $\xi_{a|M}(\lambda)$. That is, (i) every preparation P can be associated to a normalised probability distribution over the ontic state space $\mu_P(\lambda)$, such that $\int \mu_P(\lambda) d\lambda = 1$, and (ii) every measurement M, with outcomes a, can be associated to a set of response functions $\{\xi_{a|M}(\lambda)\}$ over the ontic states, satisfying $\sum_a \xi_{a|M}(\lambda) = 1$ for all λ.

As the ontic states are not directly observed, the operational statistics are obtained via marginalisation and we have:

$$\forall a, M, P: \; p(a|M, P) = \int \xi_{a|M}(\lambda) \mu_P(\lambda) d\lambda, \tag{3}$$

where for quantum theory:

$$\forall a, M, P: \; \mathrm{Tr}(\rho_P E_{a|M}) = \int \xi_{a|M}(\lambda) \mu_P(\lambda) d\lambda. \tag{4}$$

[1] For this example we assume that any transformation between preparation and measurement is trivial.

The ontological models framework has been used in numerous works to clarify the manner in which quantum theory should be considered contextual [4–8, 40]. The key assumption is that one can infer ontological equivalence from operational equivalence: for example, if two preparation procedures produce the same distributions over outcomes for all possible measurements, then any differences between them do not play a role in determining the ontic states of the system in question. Thus, the justification for why one can't distinguish between the two equivalent preparations at the *operational* level is because there is no difference between the role the preparations play at the *ontological* level. The view is that each use of a preparation device selects one from a set of possible ontic states according to exactly the same probability distribution in each run. Formally, if $\forall M$ and outcome a

$$p(a|M, P_1) = p(a|M, P_2), \tag{5}$$

then both preparations specify the same distribution over ontic states:

$$\mu_{P_1}(\lambda) = \mu_{P_2}(\lambda). \tag{6}$$

Similarly, if two measurements result in the same outcome statistics for all possible preparations then both measurements are represented by the same fixed response function. Formally, if $\forall P$ and outcomes a

$$p(a|M_1, P) = p(a|M_2, P), \tag{7}$$

then both measurements specify the same distribution over ontic states:

$$\xi_{a|M_1}(\lambda) = \xi_{a|M_2}(\lambda). \tag{8}$$

Thus, in a non-contextual ontological model one can account for operational statistics according to Eq. 3. Implicit in this model is the belief that the ontic state screens off the preparations from the measurements (a property also known as λ-mediation [18]). In [40] it was shown that one can use the assumptions above to derive a contextuality proof: no model of the form of Eq. 3 can explain the statistics of quantum theory.

In this approach to contextuality [4–8, 40], one assumes that ontic states determine correlations according to some fixed causal order. Formally, this is captured by Eq. 3: the preparation is assumed to cause the selection of a particular ontic state λ according to a fixed distribution $\mu(\lambda)$, and the measurement choice does not alter this value but merely determines the outcome probability, also according to a fixed probability distribution. This leaves open the question of whether one can explain the contextuality of quantum theory by postulating an alternative, retrocausal ontology: If the future can affect the past, then the state λ could depend on the measurement setting M, and Eq. (3) would not be justified.

Generally speaking, retrocausal approaches posit the existence of backwards-in-time causal influences to explain quantum correlations. The stated appeal of such approaches is that the consequent explanations retain some element of our classical notion of reality: local causality, determinate ontology, and counterfactual definiteness. For example, Price and Wharton explain Bell correlations by including a "zig-zag" of causal influence, passing via hidden variables that travel backwards in time from one measurement event to the source and then forwards in time to the distant measurement event [14]. Although not explicitly stated, there is also one further assumption underlying these approaches: such causal influences follow some kind of law-like behaviour. That is, one would not expect the *rules* by which such retrocausal influences propagate, or backward-in-time states evolve, to be completely ad hoc.

As stated in the introduction, we follow the Spekkens-style approach and also define non-contextuality in terms of operational equivalences. Where we depart however, is in our particular choice of operational primitives. The usual primitives of preparations, transformations and measurements do not permit one to consider causal scenarios that move beyond the most simple causally ordered situations; in these models the notion of reality is defined in terms of properties that exist *before* a measurement takes place. The underlying ontology is therefore assumed to follow some ordinary causal structure, akin to the directed acyclic graphs of causal models [41]. In our model we wish to be able to consider more general situations, for example where we include *any* possible global dynamics, causal structure, space-time geometry or global constraints. In order to provide this alternative perspective we consider the primitive operational elements to be sets of labelled local regions, locally controllable properties and an environment.

2 Operational primitives

We define an operational model of any experiment to consist of *local labelled regions* $(A, B, C, \dots)$ where one can perform *controlled operations* that can be associated with *outcomes*. The regions align with concepts such as local laboratories, communicating parties (e.g. Alice and Bob) and local space-time regions (similar, e.g., to the operational framework of [42]). There is no *apriori* assumption that these regions be "fixed" or preassigned in some manner; they are simply labels for the locus of a set of controlled operations. Controlled operations generalise the notion of preparations, measurements, transformations, and can include the addition or subtraction of ancillary systems. Examples include the orientation of a wave-plate, the instigation of a microwave pulse, and the use of a photodetector. We call such local operations the *local controllables*. Each local controllable is represented as $\tilde{\mathfrak{I}}^X$, where the superscript $X = A, B, \dots$ labels the associated region. We consider outcomes as labels associated to the result of choosing a particular local controllable; the outcomes for region A are labelled $a = 0, 1, 2, \dots$. Examples include the number of detected photons, the result of a spin measurement or the time of arrival of a photon. We allow the outcomes to have infinite possible values as this enables us to use the same variable for local controllables that have different numbers of possible outcomes. In general however, we expect that only a finite number of such outcomes is associated with non-zero probability.

Finally, we consider all the possible properties that could account for correlations between outcomes in the local regions. These include any global properties, initial states, connecting mechanisms, causal influence, or global dynamics. We call this the *environment*, $\tilde{W}$. Note that in our operational model environments and local controllables are by construction always uncorrelated. That is, if we see a property change in relation to a choice of local controllable we label this as an *outcome* and do not classify it as part of the environment.

We can thus describe an experiment by a set of regions, outcomes, local controllables and an environment. If we consider a particular run of an experiment there will in general be a collection of outcomes that occur, one for each local region. One can associate a joint probability to this set of outcomes and empirically verify probability assignments for each possible set of outcomes. An operational model for such an experiment allows one to calculate expected probabilities:

$$p(a, b, c, \dots | \tilde{\mathfrak{I}}^A, \tilde{\mathfrak{I}}^B, \tilde{\mathfrak{I}}^C, \dots, \tilde{W}). \tag{9}$$

The operational model thus specifies a distribution over outcomes for local controllables

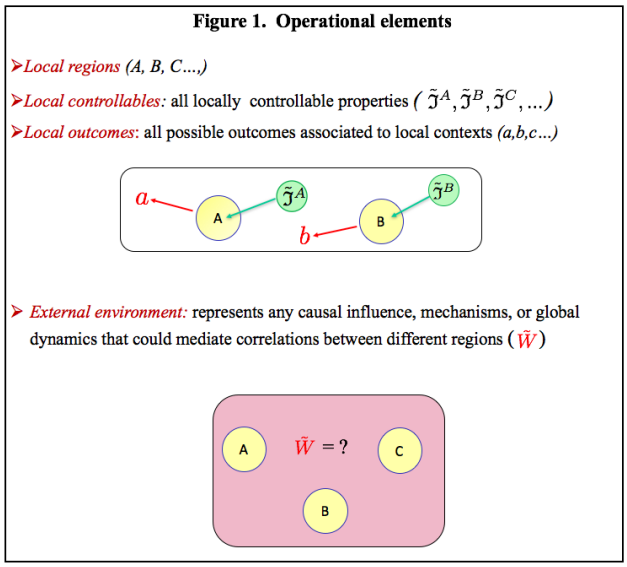

Figure 1: **Operational primitives.**

$\tilde{\mathfrak{I}}^A, \tilde{\mathfrak{I}}^B, \ldots,$ and a shared environment $\tilde{W}$, Fig. 1. Note that it should be possible to have ignorance over part of the environment and characterise this accordingly using the operational model. More explicitly, if $\tilde{\xi}$ represents the part of the environment about which we are ignorant, then the operational probabilities given the known part of the environment are obtained by marginalising over $\tilde{\xi}$:

$$\begin{aligned} p(a,b,c,\ldots|\tilde{\mathfrak{I}}^A,\tilde{\mathfrak{I}}^B,\tilde{\mathfrak{I}}^C,\ldots,\tilde{W}) &= \int d\tilde{\xi} p(a,b,c,\ldots,\tilde{\xi}|\tilde{\mathfrak{I}}^A,\tilde{\mathfrak{I}}^B,\tilde{\mathfrak{I}}^C,\ldots,\tilde{W}) \\ &= \int d\tilde{\xi} p(a,b,c,\ldots|\tilde{\mathfrak{I}}^A,\tilde{\mathfrak{I}}^B,\tilde{\mathfrak{I}}^C,\ldots,\tilde{W},\tilde{\xi}) p(\tilde{\xi}|\tilde{W}), \end{aligned} \tag{10}$$

where the second equality comes from the assumption that the local controllables are uncorrelated with the environment.

As a concrete example, $\tilde{W}$ can describe the axis along which a spin-$\frac{1}{2}$ particle is prepared, while $\tilde{\xi}$ represents whether the spin is prepared aligned or anti-aligned with that axis.[2] The marginal (10) then describes a scenario where there is some probabilistic uncertainty of the spin's direction i.e. which value of ξ occurs in any given run. Note that, for the particular case $p(\tilde{\xi}|\tilde{W}) = \frac{1}{2}$, we obtain the maximally mixed state irrespective of the axis, making the variable $\tilde{W}$ redundant. Such redundancies can be taken into account via operational equivalences.

3 Operational equivalences

We next characterise the appropriate operational equivalences in order to define our ontological model. Notationally, we omit the 'tilde' for each equivalence class.

[2]Here (and again in Section 6), we take for simplicity a scenario with a single region where a measurement is performed, so the specification of a process is equivalent to the specification of a state. More generally, the variables W and ξ could describe quantum channels, quantum networks, or more general quantum processes.

3.1 Events

We say that a pair composed of an outcome and the respective local controllable $(a, \tilde{\mathfrak{I}}^A)$ is operationally equivalent to the pair $(a', \tilde{\mathfrak{I}}'^A)$ if the joint probabilities for $a, b, c, \ldots$ and $a', b, c, \ldots$ are the same for all possible outcomes and local controllables in the other regions $B, C, \ldots$, and for all environments $\tilde{W}$.

$$p(a, b, c, \ldots | \tilde{\mathfrak{I}}^A, \tilde{\mathfrak{I}}^B, \ldots \tilde{W}) = p(a', b, c, \ldots | \tilde{\mathfrak{I}}'^A, \tilde{\mathfrak{I}}^B, \ldots \tilde{W}), \tag{11}$$
$$\forall (b, c, \ldots, \tilde{\mathfrak{I}}^B, \tilde{\mathfrak{I}}^C, \ldots, \tilde{W}).$$

We denote an equivalence class of such pairs of outcomes and local controllables as an *event*:

$$M^A = [(a, \tilde{\mathfrak{I}}^A)]. \tag{12}$$

3.2 Instruments

We define an instrument as the list of *possible events* for a local controllable $\tilde{\mathfrak{I}}^A$, where an event $M^A = [(a, \tilde{\mathfrak{I}}^A)]$ is possible for $\tilde{\mathfrak{I}^A}$ if

$$p(a, b, c, \ldots | \tilde{\mathfrak{I}}^A, \tilde{\mathfrak{I}}^B, \ldots \tilde{W}) \neq 0, \tag{13}$$

for some

$$(b, c, \ldots, \tilde{\mathfrak{I}}^B, \tilde{\mathfrak{I}}^C, \ldots, \tilde{W}).$$

We say that $\tilde{\mathfrak{I}}^A$ is equivalent to $\tilde{\mathfrak{I}}'^A$ if they define the same list of possible events and we denote the equivalence class $\mathfrak{I}^A := [\tilde{\mathfrak{I}}^A] \equiv \{M_1^A, \ldots, M_n^A\}$. Note that our definition allows distinct instruments to share one or more events. Note also, our definition implies that the probability for an event doesn't depend on the particular instrument $\mathfrak{I}$, once we assume the event is possible given the instrument. This property we call *operational instrument equivalence.*[3]

3.3 Process

The process captures those physical features responsible for generating the joint statistics for a set of events, independently of the choice of local instruments. A process is defined as an equivalence class of environments, $W := [\tilde{W}]$, where $\tilde{W}$ is equivalent to $\tilde{W}'$, if

$$p(a, b, c, \ldots | \tilde{\mathfrak{I}}^A, \tilde{\mathfrak{I}}^B, \ldots, \tilde{W}) = p(a, b, c, \ldots | \tilde{\mathfrak{I}}^A, \tilde{\mathfrak{I}}^B, \ldots \tilde{W}'), \tag{14}$$
$$\forall (a, b, c, \ldots, \tilde{\mathfrak{I}}^B, \tilde{\mathfrak{I}}^C, \ldots,).$$

A simple example is the spatio-temporal ordering of regions. It is clear that the operational statistics of events in regions A and B can be different for the following two causal orderings: (i) A is before B, (ii) B is before A; thus the respective environments, $\tilde{W}_{(i)}$ and $\tilde{W}_{(ii)}$, will not be equivalent. On the other hand, for certain experiments we would not expect any difference in statistics for a simple rotation of the whole experiment by 45 degrees; these two environments will be represented by the same process W.

The above equivalences allow us to define a joint probability distribution over the space of *events* (rather than outcomes) conditioned on *instruments* (rather than local controllables) and the *process* (rather than the environment). As discussed above, this distribution

[3] In other work, where we are not concerned with the possibility of non-contextual hidden variable theories, we refer to this property as instrument non-contextuality [37]. Here we reserve the term non-contextuality to refer to an ontological model.

satisfies operational instrument equivalence, which means that the joint probability for a set of events is either zero or independent of the respective instruments. Therefore, it can be expressed in terms of a *frame function* f_W that maps events to probabilities and is normalised for each instrument:

$$p(M^A, M^B, \dots | \mathfrak{I}^A, \mathfrak{I}^B, \dots, W) = f_W(M^A, M^B, \dots) \prod_{X=A,B,\dots} \chi_{\mathfrak{I}^X}(M^X), \tag{15}$$

where, for a set S, χ_S is the indicator function, $\chi_S(s) = 1$ for $s \in S$ and $\chi_S(s) = 0$ for $s \notin S$. Note that the indicator functions are necessary to make the whole expression a valid probability distribution, normalised over the *entire* space of events. Furthermore, and in contrast to similar expressions involving POVMs, the dependency on the instruments is crucial to allow for causal influence across the regions: Integrating over the events of, say, region A, can result in a marginal distribution that still depends on A's instrument and displays signalling from A to other regions. However, the fact that the dependency on the instruments is solely through the indicator functions tells us that the causal relations can be attributed to the particular events realised in each experimental run, rather than to the whole instruments (which include the specification of events that did not happen). In other words, the event "screens off" the instrument: once the event in a local region is known, further knowledge of the instrument does not allow for any better prediction about events in other regions.

4 Ontological model

The purpose of an ontological model is to introduce possible elements of reality. Typically, one assumes that the ontology is encoded in a "state", representing the physical properties of a system at a given time. Here we shift the focus from states to more general properties of the environment that are responsible for mediating correlations between regions. We represent the collection of all such properties by a single variable ω, named the *ontic process*. We wish to clarify at this point that our ontic process captures the physical properties of the world that remain invariant under our local operations. That is, although we allow local properties to *change* under specific operations, we wish our *ontic process* to capture those aspects of reality that are independent of this probing. The interpretation of ontic processes and the relation with the usual notion of ontic states can be seen via the examples of the following section.

Our ontological model specifies a joint probability for a set of outcomes, one at each local region, given the ontic process, the environment, and the set of local controllables. This joint probability reduces to the operational joint probability when the value of the ontic process is unknown:

$$p(a, b, c, \dots, | \tilde{\mathfrak{I}}^A, \tilde{\mathfrak{I}}^B, \dots \tilde{W}) = \int d\omega p(a, b, c, \dots, \omega | \tilde{\mathfrak{I}}^A, \tilde{\mathfrak{I}}^B, \dots \tilde{W}). \tag{16}$$

There are three natural assumptions one might require of an ontological model defined according to these operational equivalences:

Assumption 1. ***ω-mediation*** *The ontic process mediates all the correlations between regions, thus ω screens off outcomes from the environment, and we have:*

$$p(a, b, c, \dots | \tilde{\mathfrak{I}}^A, \tilde{\mathfrak{I}}^B, \dots, \tilde{W}) = \int d\omega p(a, b, c, \dots | \omega, \tilde{\mathfrak{I}}^A, \tilde{\mathfrak{I}}^B, \dots) p(\omega | \tilde{W}). \tag{17}$$

Assumption 2. ***Instrument non-contextuality.*** *Operationally indistinguishable pairs of outcomes and local controllables should remain indistinguishable at the ontological level. That is, for operationally equivalent pairs* $(a, \tilde{\mathfrak{I}}^A), (a', \tilde{\mathfrak{I}}'^A)$,

$$p(a, b, c, \dots, |\omega, \tilde{\mathfrak{I}}^A, \tilde{\mathfrak{I}}^B, \dots) = p(a', b, c, \dots, |\omega, \tilde{\mathfrak{I}}'^A, \tilde{\mathfrak{I}}^B, \dots), \tag{18}$$

$$\forall (b, c, \dots, \tilde{\mathfrak{I}}^B, \dots), \forall \omega.$$

which means that we can define a probability distribution on the space of events, conditioned on instruments and on the ontic process, in terms of a frame function f_ω*, such that:*

$$p(M^A, M^B, \dots |\omega, \mathfrak{I}^A, \mathfrak{I}^B, \dots,) = \prod_X \chi_{\mathfrak{I}^X}(M^X) f_\omega(M^A, M^B \dots), \tag{19}$$

where χ *is the indicator function,* $\chi_X(x) = 1$ *for* $x \in X$ *and* $\chi_X(x) = 0$ *for* $x \notin X$*, and* f_ω *maps events to probabilities:*

$$f_\omega(M^A, M^B, \dots) \in [0, 1], \tag{20}$$

and is normalised for each set of events that corresponds to a particular instrument:

$$\sum_{\substack{M^A \in \mathfrak{I}^A \\ M^B \in \mathfrak{I}^B \\ M^C \in \mathfrak{I}^C \\ \dots}} f_\omega(M^A, M^B, M^C, \dots) = 1. \tag{21}$$

Assumption 3. ***Process non-contextuality.***
For operationally equivalent processes $\tilde{W}, \tilde{W}'$ *the assumption of process non-contextuality implies:*

$$p(\omega|\tilde{W}) = p(\omega|\tilde{W}'), \tag{22}$$

and we can define a function $g_W(\omega)$ *that maps ontic processes to probabilities, given each process* W*:*

$$g_W(\omega) = p(\omega|\tilde{W}), \quad W = \left[\tilde{W}\right], \tag{23}$$

that is normalised for all ω*:*

$$\int d\omega \; g_W(\omega) = 1. \tag{24}$$

For an ontological model that satisfies the above three assumptions, the operational probability can now be expressed in terms of events, instruments and processes as:

$$p(M^A, M^B, \dots |\mathfrak{I}^A, \mathfrak{I}^B, \dots, W) = \prod_{X=A,B,\dots} \chi_{\mathfrak{I}^X}(M^X) \int d\omega \; f_\omega(M^A, M^B, \dots) g_W(\omega). \tag{25}$$

Although ontic states, as they are usually understood, are not represented explicitly in our framework, they are not excluded. In the following section we present three examples to illustrate how such ontic states, with or without retrocausality, can be represented in our model.

5 Examples

5.1 Deterministic, classical models

5.1.1 Causally-ordered models

As a first example, let us consider a classical, deterministic scenario (without retrocausality) with two regions, A in the past of B, each delimited by a past and future space-like boundary, see Fig. 2a. For a classical system, we can assign input states λ_I^A and λ_I^B to the past boundaries of A and B, respectively, and output states λ_O^A and λ_O^B to the respective future boundaries. As measurements can be performed without disturbance on a classical system, we associate the input state in each region with the respective measurement outcome: $a \equiv \lambda_I^A$ and $b \equiv \lambda_I^B$. As local controllables we take deterministic local operations, defined as functions f^X that map the input state of each region to the corresponding output:

$$\lambda_I^X \mapsto \lambda_O^X = f^X\left(\lambda_I^X\right), \tag{26}$$

where X denotes the respective local region, A or B. Assuming ordinary dynamical laws, the input state at B can depend on the output at A through some function:

$$\lambda_O^A \mapsto \lambda_I^B = w^B\left(\lambda_O^A\right). \tag{27}$$

The input state at A, on the other hand, does not depend on B, and thus has to be specified as an independent environment variable. The ontic process for this model is thus identified with the pair

$$\omega = \left(\lambda_I^A, w^B\right). \tag{28}$$

Indeed, knowing ω and the choice of local operations is sufficient to fully determine the measured outcomes:

$$a = \lambda_I^A, b = w^B\left(f^A\left(\lambda_I^A\right)\right). \tag{29}$$

As the model is fully deterministic, and we have not introduced any redundant variables, there are no non-trivial equivalence classes. Explicitly, an event in region A (and similarly for B) is given by the pair $\left(a, f^A\right)$, or equivalently by the input-output pair

$$\lambda^A := \left(\lambda_I^A, \lambda_O^A = f^A\left(\lambda_I^A\right)\right), \tag{30}$$

while the instrument is given by the collection of events given a choice of operation,

$$\mathfrak{I}^A = \left\{\left(\lambda_I^A,\, f^A\left(\lambda_I^A\right)\right)\right\}_{\lambda_I^A}, \tag{31}$$

which is just to say the instrument can be identified with the function f^A. We see in this example that the ontology, as traditionally understood, lies in the event variables λ. These variables are *not* independent of the local controllables, because the event at B can depend on the operation performed at A. However, there is still an aspect of the ontology that does not depend on the operations: the initial state λ_I^A and the functional relation w^B. It is this invariant aspect of the ontology that we call a process.

5.1.2 Time-travelling classical systems

General Relativity allows for space-time geometries with closed time-like curves, where a system can travel back in time and interact with its past self [43], thus providing physically-motivated examples of scenarios that defy ordinary forward causality. Notably, qualitative

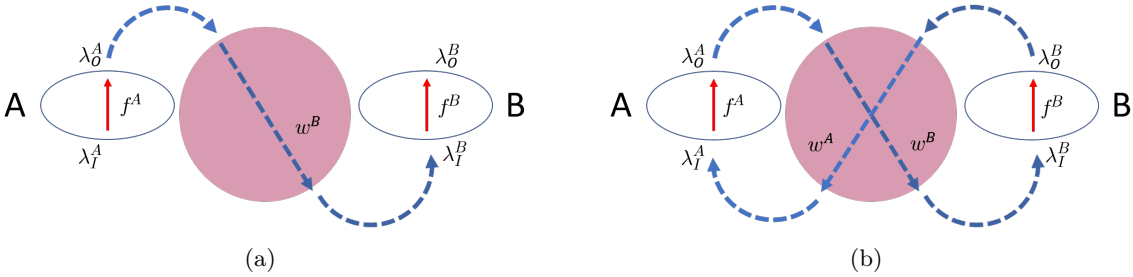

Figure 2: **Classical process with ontological interpretation.** (a) We assign input states λ_I^A and λ_I^B to the past boundaries of A and B, respectively, and output states λ_O^A and λ_O^B to the respective future boundaries. A deterministic local operation is a function f^X that maps the input state of each region to the corresponding output. (b) An example of ontic process is one describing classical closed time-like curves, defined by a pair of functions $\omega = \left(w^A, w^B\right)$, where $\lambda_O^B \mapsto w^A\left(\lambda_O^B\right) = \lambda_I^A$ and similarly for w^B.

analogies between quantum phenomena and classical time-travelling systems have been suggested [44], making the latter an interesting test-bed for generalised ontological models.

The example in the previous subsection can be readily generalised to a deterministic model of classical system near closed time-like curves by allowing the input state at A to depend on B through some function $\lambda_O^B \mapsto \lambda_I^A = w^A\left(\lambda_O^B\right)$. The process is now given by two functions, $\omega \equiv \left(w^A, w^B\right)$, Fig. 2b, with the causally-ordered case recovered when one of the two is a constant. Compatibility with arbitrary local operations imposes constraints on the function w^A, w^B and, in the two-region case, it turns out that one of them has in fact to be constant [45, 46]. However, for three or more regions, it is possible to find deterministic processes, with no constant component, that are still consistent with arbitrary local operations[4].

Also in this case, the observed outcomes are fully determined once the process and the local operations are specified, as the unique fixed points $a \equiv \lambda_I^A, b \equiv \lambda_I^B, \ldots,$ of the function obtained by composing the process $\omega = \left(w^A, w^B, \ldots\right)$ with the operations $\left(f^A, f^B, \ldots\right)$ (see Ref. [46] for more details). Crucially, in this case the events in each region can depend on the choice of operation in all regions, $\lambda^A = \lambda^A(f^A, f^B, \ldots)$. Thus, from the perspective of ordinary ontological models, time-travelling systems appear contextual, since it is impossible to assign a "state" to any region independently of the operations. Nonetheless, the relation between events, captured by the process, does not depend on the operations. Thus, following the terminology introduced here, models such as the above are both instrument and process non-contextual. (As in the previous causally-ordered example, there are no non-trivial equivalence classes, so non-contextuality is straightforward.)

More general models of classical closed time-like curves might impose restrictions on the accessible local operations[5]. Even more generally, one can consider models where

[4]The incompatibility of such processes with an underlying causal order can be demonstrated rigorously by showing they can be used to violate causal inequalities [47], device-independent constraints on probabilities imposed by a definite causal order [29, 48].

[5] A constraint on the accessible operations is often invoked to solve "paradoxes", such as a time-traveller killing their past self. Although classical studies of closed time-like curves do not support the need for this type of restriction [49–53], it might be necessary in a general theory. Such a restriction on an agent's actions is sometimes interpreted as a violation of "free will". This worry is however misplaced, since an agent can still be (or fail to be) free to perform all the physically available operations. A different set of

instruments are not associated with local input-to-output functions but with more general sets of input-output pairs, $\mathfrak{I}^A = \Lambda^A \subsetneq \Lambda_I^A \times \Lambda_O^A$, where $\Lambda_{I(O)}^A$ is the state space associated with the past (future) boundary of the local region. In such models, a choice of instrument selects which pairs of input-output states are *possible*, while a deterministic process would determine, given all choices of instruments, which pairs are actually realised. Thus, in such models both the state in the past and in the future of a local region depend on the choice of instrument, thus again they are necessarily contextual from the point of view of traditional ontological models. Yet, they remain instrument and process non-contextual as long as deterministic processes are considered.

In the above deterministic examples ω-mediation is satisfied trivially, because ontic and operational processes coincide. This can be generalised to situations where we have only partial knowledge about the environment. For example, we might not have full knowledge of the initial state, but only know the temperature T of a thermal bath from which the state is extracted; or the system might get coupled to some external environment during the evolution from one region to another. In all cases, we end up with partial knowledge of the ontic process, expressed by some probability $p(\omega|W)$ where W represents all relevant accessible information about the environment (the temperature of the bath or other noise parameters). The resulting probabilistic operational model naturally satisfies the property of ω-mediation, because knowing the temperature or noise parameters does not provide more information than already encoded in the ontic process, namely in the underlying microstates and functional relations.

Note also that our construction of an ontological model respects the mobility of the boundary between local instruments and processes that one sees in ordinary applications of quantum theory. As a simple example, consider a preparation P of a quantum system, followed by a measurement M. This can be modelled in three different ways: (i) with P as part of the environment $\tilde{W}$, and M as an instrument associated to a single local region, (ii) with P and M as instruments in two distinct local regions, and $\tilde{W}$ capturing both a channel between preparation and measurement, plus any additional information about the environment, or (iii) with both P and M characterising the instruments in a single local region and all other information about the environment modelled as $\tilde{W}$. For classical processes characterised as causal models, such a shift in perspective is formalised by the notion of "latent variables" [41]. An analogue notion of "latent laboratories" exists for quantum processes characterised as quantum causal models, and this formal structure likewise characterises the mobility of the boundary to which we refer [34].

5.1.3 All-at-once stochastic models

In the above examples of time-travelling systems, the ontic process (or at least certain aspects of it) can be understood as describing the dynamical evolution of systems between regions. Some retrocausal approaches attempt to provide an ontology for quantum mechanics that does not rely on any dynamical process; rather, one should consider all relevant events in space-time "at once". The appearance of quantum probabilities is then justified by the fact that the information available at a given time is not sufficient to fully determine the state of the system at all times (with the missing information possibly contained in some unknown boundary condition in the future). Our framework naturally captures all such models, because an ontic process need not be interpreted as a transformation: it simply represents the rule generating all relevant events given the local operations.

operations would simply represent a deviation from classical physics in the local region where the agent acts.

An instructive example is a toy model by Wharton [16], which represents a space-time scenario as a system in thermal equilibrium, with events at different space-time locations represented as states at different points in space. While having a clear ontological interpretation, this model offers qualitative analogies with quantum interference and, when analysed from an ordinary time-evolution perspective, displays an apparent contextuality. We show in detail in appendix A how (a generalisation of) Wharton's model fits within our framework and satisfies the requirements of ω-mediation and instrument and process non-contextuality.

The above three examples illustrate that it is indeed easy to represent many possible physical scenarios via ontological models that are both instrument and process non-contextual. Given the exotic nature of the latter two examples, it seems plausible that one could also produce such a model to explain quantum correlations. In the following sections we prove that this is not the case.

6 Quantum models

If one assumes that the results of experiments in local regions accord with quantum mechanics, then events can be associated with *completely positive trace-non-increasing* (CP) maps $\mathcal{M}^A : A_I \to A_O$, where input and output spaces are the spaces of linear operators over input and output Hilbert spaces of the local region, $A_I \equiv \mathcal{L}(\mathcal{H}^{A_I})$, $A_O \equiv \mathcal{L}(\mathcal{H}^{A_O})$ respectively [54]. Each set $\mathfrak{I}^A$ of CP maps that sums to a completely positive trace preserving (CPTP) map is a *quantum instrument* [55]:

$$\mathrm{Tr}\left[\sum_{\mathcal{M}^A \in \mathfrak{I}^A} \mathcal{M}^A(\rho)\right] = \mathrm{Tr}(\rho). \tag{32}$$

An instrument thus represents the collection of all possible events that can be observed given a specific choice of local controllable.

Given these definitions of events and instruments, one can predict the joint probability over possible events using a generalised form of the Born rule:

$$p(M^A, M^B, \dots | \mathfrak{I}^A, \mathfrak{I}^B, \dots, W) = f_W(M^A, M^B, \dots) \prod_X \chi_{\mathfrak{I}^X}(M^X), \tag{33}$$

$$f_W(M^A, M^B \dots) = \mathrm{Tr}\left[(M^A \otimes M^B \otimes \dots)W\right], \tag{34}$$

where $M^A, M^B \dots$ are the Choi-Jamiołkowski representations of the local CP maps associated to particular events, and W is a positive, semi-definite operator associated to the relevant process [23, 26, 29]. We call W the *process matrix*, using the terminology of Ref. [29].

It is possible to derive this trace rule for probabilities by assuming linearity [29], or alternatively one can *derive* linearity (and the trace rule) from the assumption of operational instrument equivalence alone [37]. The significance of this latter derivation is that the condition of operational instrument equivalence is formally identical to that of instrument non-contextuality, with the only difference that the latter includes the ontic process. Therefore, for each ontic process ω, the corresponding frame function can be expressed as:

$$f_\omega(M^A, M^B, \dots) = \mathrm{Tr}\left[\sigma(\omega)\mathbf{M}\right], \tag{35}$$

where we introduced the short-hand notation $\mathbf{M} \equiv M^A \otimes M^B \otimes \dots$ and $\sigma(\omega)$ is a process matrix [37]. We now wish to show that the function $g_W(\omega)$ that features in our ontological

model, under the assumption of process non-contextuality, can be represented as

$$g_W(\omega) = \mathrm{Tr}\left[\eta(\omega)W\right], \tag{36}$$

where $\{\eta(\omega)\}_{\omega\in\Omega}$, Ω being the set of ontic processes, is a quantum instrument.

It is common in non-contextuality no-go theorems (as well as in the process matrix formalism) to assume preservation of probabilistic mixtures as an assumption that is independent of the assumption of non-contextuality. Here we rather derive it from our assumption of process non-contextuality. Consider two classical variables ξ, W used to describe the process, where we already take operational equivalences into account. Following the earlier example, we can think of W as describing a cartesian axis, while ξ—the aspect of the process about which we are ignorant—describes whether a spin-$\frac{1}{2}$ particle is prepared aligned or anti-aligned to this axis. The operational probabilities given W, and the corresponding decomposition for ontological probabilities, are obtained by marginalisation:

$$\begin{aligned} &p(M^A, M^B, \dots|\mathfrak{I}^A, \mathfrak{I}^B, \dots, W) \qquad (37)\\ &= \int d\xi\, d\omega\, p(M^A, M^B, \dots|\omega, \mathfrak{I}^A, \mathfrak{I}^B, \dots, W, \xi)p(\omega|\mathfrak{I}^A, \mathfrak{I}^B, \dots, W, \xi)p(\xi|\mathfrak{I}^A, \mathfrak{I}^B, \dots, W)\\ &= \int d\xi\, d\omega\, p(M^A, M^B, \dots|\omega, \mathfrak{I}^A, \mathfrak{I}^B, \dots)p(\omega|W, \xi)p(\xi|W), \end{aligned}$$

where, in the last identity, we use the fact that $p(\omega|W, \xi)$ does not depend on the local controllables (and thus on the instruments) due to the assumption of ω-mediation; and $p(\xi|W)$ is due to our assumption that the environment and local controllables (and thus process and instruments) are uncorrelated. Additionally, due to ω-mediation, we no longer need to condition the $M^A, M^B, \dots$ directly on W and ξ.

Now let us write W_ξ for the process corresponding to the pair W, ξ. We have

$$g_{\int d\xi\ W_\xi p(\xi|W)}(\omega) = g_W(\omega) = p(\omega|W) = \int d\xi\ g_{W_\xi}(\omega)p(\xi|W), \tag{38}$$

thus $g_W(\omega)$ is convex-linear in W. The first identity in Eq. (38) comes from the fact that probabilistic mixtures of quantum processes are represented as convex combinations, thus $W = \int d\xi\ W_\xi p(\xi|W)$. This in turn is a consequence of the trace formula for operational quantum probabilities (which is itself a consequence of operational instrument equivalence):

$$p(M^A, M^B, \dots|\mathfrak{I}^A, \mathfrak{I}^B, \dots, W) = \mathrm{Tr}\left[(M^A \otimes M^B \otimes \dots)W\right]\prod_X \chi_{\mathfrak{I}^X}(M^X) \tag{39}$$

$$\begin{aligned} &= \int d\xi\ p(M^A, M^B, \dots|\mathfrak{I}^A, \mathfrak{I}^B, \dots, W, \xi)p(\xi|W)\prod_X \chi_{\mathfrak{I}^X}(M^X)\\ &= \mathrm{Tr}\left[(M^A \otimes M^B \otimes \dots)\int d\xi\ W_\xi p(\xi|W)\right]\prod_X \chi_{\mathfrak{I}^X}(M^X) \end{aligned} \tag{40}$$

for all CP maps $M^A, M^B, \dots$

Using standard linear-algebra arguments, $g_W(\omega)$ can be extended to a linear function over W, leading to the representation (36), $g_W(\omega) = \mathrm{Tr}\left[\eta(\omega)W\right]$. Positivity and normalisation of probabilities then imply

$$g_W(\omega) \geq 0 \Rightarrow \eta(\omega) \geq 0 \quad \forall\omega, \tag{41}$$

$$\int d\omega g_W(\omega) = 1 \Rightarrow \mathrm{Tr}\left[\int d\omega\eta(\omega)W\right] = 1 \quad \forall W. \tag{42}$$

Operators $\eta(\omega)$ as defined above can be understood as the Choi representation of CP maps that sum up to a trace preserving map, namely $\{\eta(\omega)\}_{\omega\in\Omega}$ defines an instrument. In general, the CP maps $\eta(\omega)$ do not have to factorise over the separate regions, therefore it might not be possible to interpret them as local operations. This is not an obstacle, as such an interpretation is not required for the rest of the argument.

7 A quantum contradiction

To summarise the results so far, we have an operational rule for the predictions of the joint probabilities of outcomes according to quantum theory:

$$p(M^A, M^B, \dots | \mathfrak{I}^A, \mathfrak{I}^B, \dots, W) = \prod_X \chi_{\mathfrak{I}^X}(M^X)\, \mathrm{Tr}\,[\mathbf{M}\, W]\,. \tag{43}$$

We also have an ontological model for predicting the joint probabilities under the assumptions of ω-mediation, instrument non-contextuality and process non-contextuality:

$$p(M^A, M^B, \dots | \mathfrak{I}^A, \mathfrak{I}^B, \dots, W) = \prod_X \chi_{\mathfrak{I}^X}(M^X) \int d\omega f_\omega(M^A, M^B, \dots) g_W(\omega), \tag{44}$$

which given the results of the last section, becomes:

$$p(M^A, M^B, \dots | \mathfrak{I}^A, \mathfrak{I}^B, \dots, W) = \prod_X \chi_{\mathfrak{I}^X}(M^X) \int d\omega\, [\mathrm{Tr}\,\sigma(\omega)\mathbf{M}]\, [\mathrm{Tr}\,\eta(\omega) W]\,. \tag{45}$$

If this accords with quantum predictions then we should have:

$$\mathrm{Tr}\,[\mathbf{M}\, W] = \int d\omega\ \ [\mathrm{Tr}\,\sigma(\omega)\mathbf{M}]\, [\mathrm{Tr}\,\eta(\omega) W]\ \ \forall M, W. \tag{46}$$

It has been noted [40] that a decomposition of the form (44) is akin to the expression of expectation values in terms of quasi-probability distributions [56, 57]. However, the non-contextuality assumptions force both f_ω and g_W to be ordinary, positive probability distributions. It is well known that quantum expectation values cannot be expressed in such a way. It is however instructive to consider an explicit contradiction within the present process framework.

From (46),

$$\mathrm{Tr}\,[\mathbf{M}\, W] = \mathrm{Tr}\left[\mathbf{M} \int d\omega\ \sigma(\omega)\ g_W(\omega)\right] \quad \forall \mathbf{M} \tag{47}$$

$$\rightarrow W = \int d\omega\ \sigma(\omega)\ g_W(\omega), \tag{48}$$

which follows from the fact that $\mathbf{M}$ span a complete set of the joint linear space $A^I \otimes A^O \otimes B^I \otimes B^O, \dots$

Eq. (48) tells us that W is a convex mixture of the operators $\sigma(\omega)$. If W is extremal, namely if it cannot be decomposed into a non-trivial convex combination of other processes, then $W \propto \sigma(\omega)$ for $g_W(\omega) \neq 0$. Denoting the support of g_W by Ω_W, i.e., $\omega \in \Omega_W \Leftrightarrow g_W(\omega) \neq 0$, we have $W \propto \sigma(\omega)\ \forall \omega \in \Omega_W$ for an extremal W.

Consider now a process W that can be decomposed into two distinct mixtures of two sets of extremal processes W_j and W'_k (we take discrete sets for simplicity):

$$W = \sum_j q_j W_j = \sum_k p_k W'_k. \tag{49}$$

Since g_W is convex-linear in W, we have $g_W = \sum_j q_j g_{W_j}$. This means that, for every $\omega \in \Omega_W$, there must be a j such that $g_{W_j}(\omega) \neq 0$. In other words, $\Omega_W = \bigcup_j \Omega_{W_j}$. By a similar argument, we have that $\Omega_W = \bigcup_k \Omega_{W'_k}$. We thus see that each convex decomposition of W into distinct extremal processes corresponds to a partition of W's support into the extremal processes' supports. This in turns implies that each ω belongs to both Ω_{W_j} and $\Omega_{W'_k}$, for some j and k. As we have seen, this would imply

$$\sigma(\omega) \propto W_j \propto W'_k. \tag{50}$$

However, one can find many examples where no process in one decomposition is proportional to any process in the other. This implies a contradiction and shows that a decomposition such as (46) cannot exist for all CP maps and quantum processes. As a particular example to show the above contradiction, consider a process W corresponding to a quantum channel from a region with a two-level output, A_O to a region with a two-level input, B_I:

$$W = \sum_j q_j W_j = \sum_k p_k W'_k, \tag{51}$$

formed from the following two combinations of extremal processes:

$$W_1 = [[\mathbb{1}]] = \mathbb{1} + X \otimes X - Y \otimes Y + Z \otimes Z, \tag{52}$$
$$W_2 = [[X]] = \mathbb{1} + X \otimes X + Y \otimes Y - Z \otimes Z, \tag{53}$$
$$W_3 = [[Y]] = \mathbb{1} - X \otimes X - Y \otimes Y - Z \otimes Z, \tag{54}$$
$$W_4 = [[Z]] = \mathbb{1} - X \otimes X + Y \otimes Y + Z \otimes Z. \tag{55}$$

$$W'_1 = [[U\mathbb{1}]] = \mathbb{1} + X \otimes UXU^\dagger - Y \otimes UYU^\dagger + Z \otimes UZU^\dagger, \tag{56}$$
$$W'_2 = [[UX]] = \mathbb{1} + X \otimes UXU^\dagger + Y \otimes UYU^\dagger - Z \otimes UZU^\dagger, \tag{57}$$
$$W'_3 = [[UY]] = \mathbb{1} - X \otimes UXU^\dagger - Y \otimes UYU^\dagger - Z \otimes UZU^\dagger, \tag{58}$$
$$W'_4 = [[UZ]] = \mathbb{1} - X \otimes UXU^\dagger + Y \otimes UYU^\dagger + Z \otimes UZU^\dagger. \tag{59}$$

where X, Y and Z are the Pauli matrices, U is a unitary, and we used the notation $[[V]] := \sum_{rs} |r\rangle \langle s| \otimes V |r\rangle \langle s| V^\dagger$ for the Choi representation of a unitary V.

It is clear that no W_j is proportional to any W'_k for an appropriate choice of U, and we have a contradiction with (50).

8 Process-contextual extensions of quantum theory

Contextuality proofs do not always require both preparation and measurement non-contextuality. Indeed, many no-go theorems focus on the requirement of measurement non-contextuality alone. Interestingly, even without preparation non-contextuality, measurement non-contextuality imposes strong constraints on the ontology. Essentially, any non-contextual ontology must reduce to the Beltrametti-Bugajski (BB) model [58], which identifies elements of reality with the quantum wave function. An important consequence of this result is that no measurement non-contextual extension of quantum theory exists that can provide more accurate predictions of experimental outcomes [5].

It is thus interesting to consider dropping the requirement of process non-contextuality in our framework, leaving instrument non-contextuality as the sole requirement. It is easy

to see that instrument non-contextual, process-*contextual* models are possible. An example is a model where the ontic process is directly identified with the quantum process:

$$g_W(\omega) = \delta(W - \omega). \tag{60}$$

Operational probabilities are then recovered simply by using the "quantum process rule", Eq. (34), for the ontic frame function:

$$f_\omega(M^A, M^B \dots) = \mathrm{Tr}[\mathbf{M}\,\omega]. \tag{61}$$

This "crude" ontological model is similar to the BB model. A difference is that the BB model only identifies *pure* quantum states with elements of reality, while in Eq. (60) *any* process counts as ontic, including those corresponding to mixed states or noisy channels. One could refine the above model by only allowing an appropriately defined "pure process" to be ontic. (See however Ref. [59] for possible ambiguities regarding such a definition.)

A similar non-extendability result to that of [5] also holds in our case. As already discussed above, the only instrument non-contextual frame function must be given by Eq. (35), namely to every ontic process ω is associated a process matrix $\sigma(\omega)$. The implication is that an instrument non-contextual hidden variable cannot provide more information than that contained in a process matrix. We thus conclude that quantum mechanics admits no non-trivial, instrument non-contextual extension. Indeed, this result holds independently of any assumptions one may make about the causal structure of a possible underlying ontology. Therefore, even instrument non-contextuality alone poses strong restrictions on hidden variable models that attempt to leverage exotic causal structures to recover a non-contextual notion of reality.

9 Discussion

We have shown that it is not possible to construct an ontological model that is both instrument and process non-contextual and also accords with the predictions of quantum mechanics. We take both forms of non-contextuality to be very reasonable assumptions if one wishes some aspect of "reality" to be describable in a manner that is independent of the act of experimentation. Thus our work shows that models that posit unusual causal, global or dynamical relations will not solve a key quantum mystery, that of contextuality.

Standard no-go theorems show that quantum theory is not consistent with ontological models where the properties of a system exist prior to and independently of the way they are measured. A possible interpretation is that properties *do* exist, but they are in fact dependent on future actions. Here we have shown that hidden variable models that attempt to leverage such influence from the future have to violate some broader form of non-contextuality. This new notion of non-contextuality refers to the rules that dictate how local actions influence observed events, rather than to states and measurements.

We have introduced three assumptions in order to analyse non-contextuality in such scenarios where influence from the future is possible. The core idea is captured by the assumption of ω-mediation. This states that an agent's actions should effect the world according to rules or laws that do not themselves depend on such actions. Indeed, if the rules changed every time we changed how we intervened on the world, we would not call them "rules" to begin with. In the context of ontological models, this assumption allows one to assume that experiments uncover an aspect of nature that is unchanging.

The second assumption, instrument non-contextuality, states that operationally equivalent interventions should not produce distinct effects at the ontological level. We have

shown that this assumption is compatible with scenarios that *would* be interpreted as contextual when viewed from an ordinary, time-oriented perspective. For example, we have illustrated that time-travelling models where states *can* depend on future interventions satisfy the requirement of instrument non-contextuality. Despite this generality, instrument non-contextuality is nonetheless sufficient to rule out all non-trivial hidden-variable extensions of quantum theory: Any additional variable that could provide better predictions for quantum statistics than ordinary quantum mechanics must be instrument *contextual*.

Our third assumption, process non-contextuality, states that the probabilistic assignment of the ontic description of an experiment should reflect the operationally equivalent arrangements of the same experiment. Here by "experiment" we mean the specification of the set of conditions under which agents can operate. That is, we include in this description all aspects of a physical scenario other than the choices of settings and the observed outcomes. Such aspects include what kind of systems are involved, the laws describing such systems, boundary conditions, etc. We have shown that no ontic model can satisfy this requirement of process non-contextuality, including those that directly identify quantum objects as ontic.

The distinction between background environment variables and locally controllable settings that one makes when describing experiments using our approach is of course mobile. What counts as a freely chosen parameter in one situation can count as a fixed parameter in another. Our result is robust under such a shift in perspective: no matter how we decide to describe a quantum experiment, it will not be possible to find an ontic representation for it that is both instrument and process non-contextual.

Finally, we draw attention to the fact that our results rely on complete matching to the operational predictions of quantum theory. This is a recognised feature of all ontological models that rely on operational equivalence classes and leaves open the possibility that particular ontological models might allow for some experimentally testable, different predictions. Thus, for proponents of particular retrocausal models, the door remains open to develop their ontology such that they can predict some possible deviation from quantum statistics. In the face of such statistical deviation, the possibility of a non-contextual ontological model remains open.

Acknowledgments

We thank Časlav Brukner, Eric Cavalcanti, Ravi Kunjwal, Matthew Leifer, Gerard Milburn, Alberto Montina, David Schmid, Robert Spekkens, and Ken Wharton for helpful discussions. This work was supported by an Australian Research Council Centre of Excellence for Quantum Engineered Systems grant (CE 110001013), and by the Templeton World Charity Foundation (TWCF 0064/AB38). F.C. acknowledges support through an Australian Research Council Discovery Early Career Researcher Award (DE170100712). This publication was made possible through the support of a grant from the John Templeton Foundation. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation. We acknowledge the traditional owners of the land on which the University of Queensland is situated, the Turrbal and Jagera people.

References

[1] S. Kochen and E. Specker, "The problem of hidden variables in quantum mechanics," *J. Math. Mech.* **17**, 59–87 (1967).

[2] J. S. Bell, "On the problem of hidden variables in quantum mechanics," *Rev. Mod. Phys.* **38**, 447–452 (1966).

[3] A. Cabello, "Experimentally testable state-independent quantum contextuality," *Phys. Rev. Lett.* **101**, 210401 (2008).

[4] R. W. Spekkens, "Contextuality for preparations, transformations, and unsharp measurements," *Phys. Rev. A* **71**, 052108 (2005).

[5] Z. Chen and A. Montina, "Measurement contextuality is implied by macroscopic realism," *Phys. Rev. A* **83**, 042110 (2011).

[6] R. Kunjwal, "Contextuality beyond the Kochen-Specker theorem," `arXiv:1612.07250 [quant-ph]`.

[7] M. D. Mazurek, M. F. Pusey, R. Kunjwal, K. J. Resch, and R. W. Spekkens, "An experimental test of noncontextuality without unphysical idealizations," *Nat. commun.* **7**, 11780 (2016).

[8] D. Schmid and R. W. Spekkens, "Contextual Advantage for State Discrimination," *Phys. Rev. X* **8**, 011015 (2018).

[9] E. G. Cavalcanti, "Classical Causal Models for Bell and Kochen-Specker Inequality Violations Require Fine-Tuning," *Phys. Rev. X* **8**, 021018 (2018).

[10] A. Chailloux, I. Kerenidis, S. Kundu, and J. Sikora, "Optimal bounds for parity-oblivious random access codes," *New J. Phys.* **18**, 045003 (2016).

[11] R. W. Spekkens, D. H. Buzacott, A. J. Keehn, B. Toner, and G. J. Pryde, "Preparation contextuality powers parity-oblivious multiplexing," *Phys. Rev. Lett.* **102**, 010401 (2009).

[12] M. Howard, J. Wallman, V. Veitch, and J. Emerson, "Contextuality supplies the 'magic' for quantum computation," *Nature* **510**, 351–355 (2014).

[13] H. Price, "Does time-symmetry imply retrocausality? How the quantum world says "Maybe"?," *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* **43**, 75–83 (2012).

[14] H. Price and K. Wharton, "Disentangling the Quantum World," *Entropy* **17**, 7752–7767 (2015).

[15] P. W. Evans, H. Price, and K. B. Wharton, "New Slant on the EPR-Bell Experiment," *Brit. J. Philos. Sci.* **64**, 297–324 (2013).

[16] K. Wharton, "Quantum States as Ordinary Information," *Information* **5**, 190–208 (2014).

[17] Y. Aharonov, E. Cohen, and T. Shushi, "Accommodating Retrocausality with Free Will," *Quanta* **5**, 53–60 (2016).

[18] M. S. Leifer and M. F. Pusey, "Is a time symmetric interpretation of quantum theory possible without retrocausality?," *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **473**, (2017).

[19] R. I. Sutherland, "How retrocausality helps," *AIP Conference Proceedings* **1841**, 020001 (2017).

[20] A. Carati and L. Galgani, "Nonlocality of classical electrodynamics of point particles, and violation of Bell's inequalities," *Nuovo Cimento B* **114**, 489–500 (1999).

[21] S. Weinstein, "Nonlocality Without Nonlocality," *Found. Phys.* **39**, 921–936 (2009).

[22] C. J. Wood and R. W. Spekkens, "The lesson of causal discovery algorithms for quantum correlations: Causal explanations of Bell-inequality violations require fine-tuning," *New J. Phys.* **17**, 033002 (2015).

[23] G. Gutoski and J. Watrous, "Toward a general theory of quantum games," in *Proceedings of 39th ACM STOC*, pp. 565–574. 2006. `arXiv:quant-ph/0611234`.

[24] G. Chiribella, G. M. D'Ariano, and P. Perinotti, "Quantum Circuit Architecture," *Phys. Rev. Lett.* **101**, 060401 (2008).

[25] G. Chiribella, G. M. D'Ariano, and P. Perinotti, "Memory Effects in Quantum Channel Discrimination," *Phys. Rev. Lett.* **101**, 180501 (2008).

[26] G. Chiribella, G. M. D'Ariano, and P. Perinotti, "Theoretical framework for quantum networks," *Phys. Rev. A* **80**, 022339 (2009).

[27] A. Bisio, G. Chiribella, G. D'Ariano, and P. Perinotti, "Quantum networks: General theory and applications," *Acta Physica Slovaca. Reviews and Tutorials* **61**, 273–390 (2011). `arXiv:1601.04864 [quant-ph]`.

[28] A. Bisio, G. M. D'Ariano, P. Perinotti, and M. Sedlák, "Optimal processing of reversible quantum channels," *Physics Letters A* **378**, 1797 – 1808 (2014).

[29] O. Oreshkov, F. Costa, and Č. Brukner, "Quantum correlations with no causal order," *Nat. Commun.* **3**, 1092 (2012).

[30] K. Modi, "Operational approach to open dynamics and quantifying initial correlations," *Sci. Rep.* **2**, 581 (2012).

[31] M. S. Leifer and R. W. Spekkens, "Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference," *Phys. Rev. A* **88**, 052130 (2013).

[32] M. Ringbauer, C. J. Wood, K. Modi, A. Gilchrist, A. G. White, and A. Fedrizzi, "Characterizing Quantum Dynamics with Initial System-Environment Correlations," *Phys. Rev. Lett.* **114**, 090402 (2015).

[33] F. A. Pollock, C. Rodríguez-Rosario, T. Frauenheim, M. Paternostro, and K. Modi, "Non-Markovian quantum processes: Complete framework and efficient characterization," *Phys. Rev. A* **97**, 012127 (2018).

[34] F. Costa and S. Shrapnel, "Quantum causal modelling," *New J. Phys.* **18**, 063032 (2016).

[35] J.-M. A. Allen, J. Barrett, D. C. Horsman, C. M. Lee, and R. W. Spekkens, "Quantum Common Causes and Quantum Causal Models," *Phys. Rev. X* **7**, 031021 (2017).

[36] S. Milz, F. A. Pollock, and K. Modi, "Reconstructing open quantum system dynamics with limited control," `arXiv:1610.02152 [quant-ph]`.

[37] S. Shrapnel, F. Costa, and G. Milburn, "Updating the Born rule," *New J. Phys.* **20** , 053010 (2018).

[38] N. Harrigan and R. Spekkens, "Einstein, Incompleteness, and the Epistemic View of Quantum States," *Found. Phys.* **40**, 125–157 (2010).

[39] M. S. Leifer, "Is the quantum state real? An extended review of ψ-ontology theorems," `arXiv:1409.1570 [quant-ph]`.

[40] R. W. Spekkens, "Negativity and Contextuality are Equivalent Notions of Nonclassicality," *Phys. Rev. Lett.* **101**, 020401 (2008).

[41] J. Pearl, *Causality.* Cambridge University Press, 2009.

[42] O. Oreshkov and C. Giarmatzi, "Causal and causally separable processes," *New J. Phys.* **18**, 093020 (2016).

[43] M. S. Morris, K. S. Thorne, and U. Yurtsever, "Wormholes, time machines, and the weak energy condition," *Phys. Rev. Lett.* **61**, 1446 (1988).

[44] S. Durand, "An amusing analogy: modelling quantum-type behaviours with wormhole-based time travel," *Journal of Optics B: Quantum and Semiclassical Optics* **4**, S351 (2002).

[45] Ä. Baumeler and S. Wolf, "The space of logically consistent classical processes without causal order," *New J. Phys.* **18**, 013036 (2016).

[46] Ä. Baumeler, F. Costa, T. C. Ralph, S. Wolf, and M. Zych, "Reversible time travel with freedom of choice," `arXiv:1703.00779 [quant-ph]`.

[47] Ä. Baumeler, A. Feix, and S. Wolf, "Maximal incompatibility of locally classical behavior and global causal order in multi-party scenarios," *Phys. Rev. A* **90**, 042106 (2014).

[48] C. Branciard, M. Araújo, A. Feix, F. Costa, and Č. Brukner, "The simplest causal inequalities and their violation," *New J. Phys.* **18**, 013008 (2016).

[49] J. Friedman, M. S. Morris, I. D. Novikov, F. Echeverria, G. Klinkhammer, K. S. Thorne, and U. Yurtsever, "Cauchy problem in spacetimes with closed timelike curves," *Phys. Rev. D* **42**, 1915–1930 (1990).

[50] F. Echeverria, G. Klinkhammer, and K. S. Thorne, "Billiard balls in wormhole spacetimes with closed timelike curves: classical theory," *Phys. Rev. D* **44**, 1077–1099 (1991).

[51] A. Lossev and I. D. Novikov, "The Jinn of the time machine: nontrivial self-consistent solutions," *Class. Quantum Grav.* **9**, 2309 (1992).

[52] I. D. Novikov, "Time machine and self-consistent evolution in problems with self-interaction," *Phys. Rev. D* **45**, 1989–1994 (1992).

[53] E. V. Mikheeva and I. D. Novikov, "Inelastic billiard ball in a spacetime with a time machine," *Phys. Rev. D* **47**, 1432–1436 (1993).

[54] M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information.* Cambridge University Press, 2000.

[55] E. Davies and J. Lewis, "An operational approach to quantum probability," *Comm. Math. Phys.* **17**, 239–260 (1970).

[56] E. Wigner, "On the Quantum Correction For Thermodynamic Equilibrium," *Phys. Rev.* **40**, 749–759 (1932).

[57] M. Scully and M. Zubairy, *Quantum Optics.* Cambridge University Press, 1997.

[58] E. G. Beltrametti and S. Bugajski, "A classical extension of quantum mechanics," *J. Phys. A: Math. Gen.* **28**, 3329 (1995).

[59] M. Araújo, A. Feix, M. Navascués, and Č. Brukner, "A purification postulate for quantum mechanics with indefinite causal order," *Quantum* **1**, 10 (2017).

A Wharton's retrocausal toy model

The core idea of the model is to represent a system across space-time, analogously to the representation of a system in space in thermodynamical equilibrium. Rather than being determined by dynamical evolution, the states at each point in space-time are known with some probability. This is similar to how macrostates can be considered as providing probability distributions for microstates.
In this model each event in space-time is represented as a site, labelled by the index j, within a lattice. At each site j we can have a particle in a state λ, whose possible values are assumed to be ± 1 for simplicity. The entire system across space-time is treated "all-at-once" in the same way one would treat a spatially extended system, where each site represents a different location in space. The system is then associated with a Hamiltonian $H = -\sum_{<i,j>} \lambda_i \lambda_j$, where the sum is taken over nearest-neighbours according the geometry of the lattice. All we know about the system is that it is in a thermal state, with inverse temperature β, thus the probability for a certain configuration $\vec{\lambda} := (\lambda_1, \lambda_2, \dots)$ is $p(\vec{\lambda}|\beta) \propto e^{-H(\vec{\lambda})\beta}$. If we learn the state of one of the sites, we need to update the thermal distribution by conditioning on the observed value. However, since the model is supposed to represent a space-time configuration, the sites we can observe at any given time are restricted.

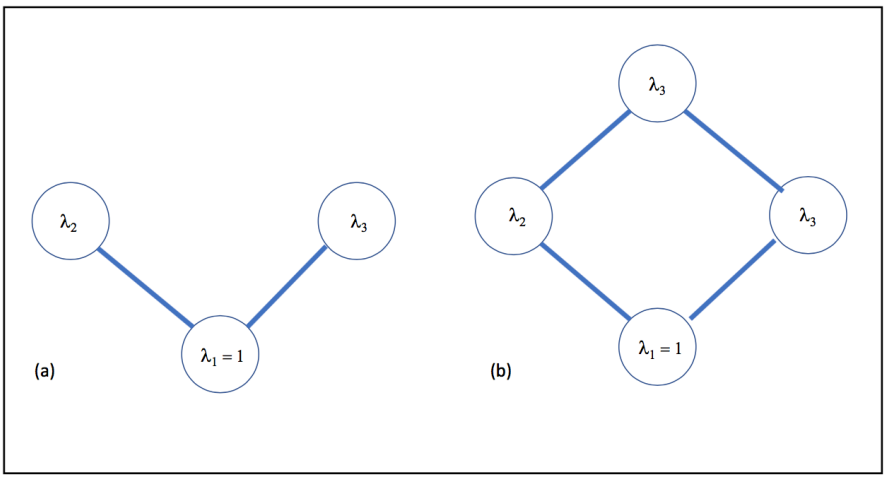

Figure 3: **Wharton's toy model** [16]. Each node j represents a location in space-time where a system can be found in a state λ_j, $j = 1, 2, \dots$. The state of the entire system is sampled from a thermal ensemble, defined by a Hamiltonian containing interactions between nodes connected by an edge, where each node is treated as a site in a spatially distributed lattice. (a) Observing the system at a given time reveals the state at one of the nodes, e.g. $\lambda_1 = 1$, upon which the probability assignment at the other nodes has to be updated. (b) The analogue of an interference experiment is represented by the insertion of an additional node in the future, which results in a different thermal state and thus in a different probability distribution for all states. An observer at an earlier time that ignores this possibility might interpret such a dependence from future actions as a form of contextuality.

Retrocausality is introduced by assuming that performing a measurement at any given time can result in the introduction of a new site, thus changing the geometry of the

system, Fig. 3. Assuming a thermal state with a given temperature, the two geometries result in different probability distributions for the microstates. If the system is interpreted as time oriented, and the influence of the future intervention is ignored, then one might be led to the conclusion that it is impossible to assign non-contextual states of reality to the system. The analogy is seen with a quantum interference experiment, where a measurement in the future is assumed to change the conditions that determine the state of the system in its past. If the influence from the future measurement is included, argues Wharton, then one might be able to recover an ontic interpretation of quantum mechanics, where the quantum state simply represents lack of information about the underlying state.
This model is interesting because causal influence is not mediated by an explicit mechanism, as opposed to ordinary dynamical systems including the time-travelling examples in the main text. Nonetheless, it is possible to fit this model into our general ontological framework, where the observed probabilities are mediated by an ontic process. Crucially, the model turns out to be both instrument and process non-contextual, showing that approaches of this type cannot reproduce the predictions of quantum theory.

Classical systems on an arbitrary geometry

We consider a more general version of Wharton's model, with arbitrary geometry, an arbitrary set of discrete values for the states, and arbitrary local interactions.
Consider a set $\mathcal{N}$ of $|\mathcal{N}| = N$ sites. Each site $j \in \mathcal{N}$ can contain a classical system whose state λ_j can take value in some set $\mathcal{S}_j$. The state of the entire system is thus described by a vector $\vec{\lambda} \equiv (\lambda_1, \dots, \lambda_N) \in \mathcal{S} := \bigtimes_{j\in\mathcal{N}} \mathcal{S}_j$.
A Hamiltonian function $H(\vec{\lambda})$ is defined on the system. We assume that this Hamiltonian is *local*, namely it is a sum of terms representing local interactions between sites. A subset of sites $e \subset \mathcal{N}$ contributing to an interaction term is called "hyperedge" and the set $\mathcal{E}$ of hyperedges defines a "hypergraph" over $\mathcal{N}$. The Hamiltonian can thus be decomposed as

$$H = \sum_{e\in\mathcal{E}} h_e, \tag{62}$$

where each term h_e is function on the space $\mathcal{L}_e := \bigtimes_{j\in e} \mathcal{S}_j$. By convention, we identify the state $\lambda_j = 0$ of system j with the "empty site", namely with no system in it. This implies that, for every hyperedge e containing j,

$$h_e\left(\dots, \lambda_j = 0, \dots\right) = 0. \tag{63}$$

In other words, each interaction term vanishes when one of the sites on which it acts is empty. In this way, "different geometries" corresponding to additional or missing sites, are simply represented as a particular choice of states in a fixed geometry.
In our terminology, each site j represents a (space-time) region and each state λ_j represents an event. We can interpret each event as "ontic"; however, since we assume that each ontic event can also be observed, ontic and operational events are identified.

No control. Before considering the possibility of interventions, it is useful to see how our framework applies to the simpler scenario with no interventions. In this case, "process" is synonymous with "state". Thus, a deterministic process is simply a specific microstate $\vec{\lambda}$, while a general probabilistic process is a probability distributions $P\left(\vec{\lambda}\right)$. For the case of a thermal state, where the only information we can access about the

environment is the inverse temperature β, the operational process is probabilistic, given by the Gibbs distribution

$$p(\vec{\lambda} \mid \beta) = \frac{e^{-\beta H(\vec{\lambda})}}{Z(\beta)}, \quad Z(\beta) = \sum_{\vec{\lambda} \in \mathcal{S}} e^{-\beta H(\vec{\lambda})}. \tag{64}$$

Since there are no irrelevant environment variables in this model, questions about contextuality do not arise: each value of the environment variable corresponds to just one process (i.e. to one probability distribution for the "events"). Therefore, at the formal level, we could identify the operational process with the ontic process; the resulting model would be process non-contextual by construction (instrument non-contextuality is even more trivial here, because there is no choice of instruments). A more natural ontological model is a deterministic one, where each ontic process (or ontic state) is identified with one microstate $\vec{\lambda}$. As required by the general formalism, the operational process provides a probability distribution over the possible ontic processes, and knowing the ontic process makes knowledge of the operational process redundant (the ontic process "screens off" the operational one), in agreement with the property of ω-mediation.

Local control. Local instruments are defined as subsets of events and represent the possibility of local control. Thus, in general, the possible sets of instruments at a site $j \in \mathcal{N}$ corresponds to a subset $\mathfrak{I}^j \subset \mathcal{S}_j$. As a simple case-study, we consider the scenario where the only control is inserting or removing a site, as in Wharton's example. Therefore, for each region $j \in \mathcal{N}$ there are two possible instruments:

$$\mathfrak{I}^j_0 := \{0\}, \qquad \mathfrak{I}^j_1 := \mathcal{S}_j \setminus \{0\}. \tag{65}$$

A prominent feature of this example is that instruments are disjoint sets, so there is never an event that belongs to two distinct instruments. This ensures the instrument non-contextuality of the model.
As in the no-control case, a deterministic process corresponds to a specification of all events, while a probabilistic process corresponds to a probability distribution for the possible events. The possibility of control means that the events now can depend on the instruments, so the process must encode this dependency. Thus, a deterministic process is given by a set of functions

$$\vec{\mathfrak{I}} \equiv \left(\mathfrak{I}^1, \dots, \mathfrak{I}^N\right) \mapsto \vec{\lambda}, \quad \lambda_j = \omega_j(\vec{\mathfrak{I}}),\ j \in \mathcal{N} \tag{66}$$

such that, for each $j \in \mathcal{N}$,

$$\omega_j\left(\dots, \mathfrak{I}^j = \mathfrak{I}^j_0, \dots\right) = 0, \qquad \omega_j\left(\dots, \mathfrak{I}^j = \mathfrak{I}^j_1, \dots\right) \neq 0. \tag{67}$$

Condition (67) simply says that if we choose to remove the system from site j ($\mathfrak{I}^j = \mathfrak{I}^j_0$), then there will be no system at site j ($\lambda_j = 0$), while if we choose to insert the system ($\mathfrak{I}^j = \mathfrak{I}^j_1$) then the system will be there, in one of its possible states ($\lambda_j \in \mathcal{S}_j \setminus \{0\}$).
For a probabilistic process, dependency on the instruments is encoded in a conditional probability distribution $p(\vec{\lambda} \mid \vec{\mathfrak{I}})$. The probabilistic version of the consistency condition (67) reads

$$\begin{aligned} P\left(\dots, \lambda_j \neq 0, \dots \mid \dots, \mathfrak{I}^j = \mathfrak{I}^j_0, \dots\right) &= 0, \\ P\left(\dots, \lambda_j = 0, \dots \mid \dots, \mathfrak{I}^j = \mathfrak{I}^j_1, \dots\right) &= 0. \end{aligned} \tag{68}$$

A more compact way to represent a (non-contextual) process is through a frame function, which can be defined piece-wise as:

$$f(\vec{\lambda}) := p(\vec{\lambda} \mid \vec{\mathfrak{I}}) \quad \text{for } \lambda_1 \in \mathfrak{I}^1, \ldots, \lambda_N \in \mathfrak{I}^N. \tag{69}$$

The consistency condition (68) is then expressed as

$$p(\vec{\lambda} \mid \vec{\mathfrak{I}}) = f(\vec{\lambda}) \prod_{j \in \mathcal{N}} \chi_{\mathfrak{I}^j}(\lambda_j). \tag{70}$$

Let us pause for a moment on this definition. The frame function is defined as a function $f : \mathcal{S} \to [0, 1]$. That is, it assigns a probability to each N-tuplet $\vec{\lambda} = (\lambda_1 \ldots, \lambda_N)$, without needing any additional information about the instruments. In the case we are considering, different instruments are non-overlapping sets of states. Therefore, if we know the state λ_j, we automatically know the instrument $\mathfrak{I}^j$ and requiring "independence from the instruments" is completely trivial: either the site is there, and the instrument is $\mathfrak{I}^j_1$, or the site is not there, and $\mathfrak{I}^j = \mathfrak{I}^j_1$. Once we know the state, there is nothing more the instrument can tell us. Technically, each value $\vec{\mathfrak{I}}$ defines a subset in the domain of f, and the value f takes in each of these subsets is given by the conditional probability (69). The non-overlapping of different instruments is crucial for this construction: if the same $\vec{\lambda}$ could belong to two different instruments, we would not know which value of $p(\vec{\lambda} \mid \vec{\mathfrak{I}})$ to use to define the frame function. For overlapping instruments, the existence of a frame function is equivalent to the assumption of instrument non-contextuality.

Processes for the thermal state Once again, the only environment variable is the inverse temperature β, which thus parametrises the operational processes. Given the above discussion, it should be clear that, for each β, we can write a conditional probability in the form (70), where the frame function is defined as in Eq. (69), with probabilities provided by the Gibbs distribution (64). Explicitly,

$$f_\beta(\vec{\lambda}) = \frac{e^{-\beta H(\vec{\lambda})}}{Z(\beta \mid \vec{\mathfrak{I}})} \quad \text{for } \vec{\lambda} \in \vec{\mathfrak{I}}, \tag{71}$$

$$\text{where } Z(\beta \mid \vec{\mathfrak{I}}) := \sum_{\vec{\lambda} \in \vec{\mathfrak{I}}} e^{-\beta H(\vec{\lambda})}. \tag{72}$$

Let us stress that, from the perspective of our framework, we might as well stop here: we already have a model that is both instrument and process non-contextual. The point of our theorem is to see if it is possible to write a given operational model in terms of an underlying non-contextual model; if that is possible, the operational model cannot reproduce the predictions of quantum mechanics. In this case we already have a non-contextual model, so we know it cannot reproduce quantum mechanics. Note that the theorem does not rely on any interpretation we might assign to ontic processes, events etc; it is simply a statement about properties that an ontological model can or cannot have.

For the sake of completeness, and since we would more naturally associate ontology with determinism, we can write explicitly how a deterministic process model looks in the present case study. Recall that a deterministic process is a (multi-valued) function $\vec{\omega} \equiv (\omega_1 \ldots \omega_N)$ from the instruments to the events. For notational convenience, we can identify the two possible instruments $\mathfrak{I}^j_{x_j}$ at each site j with their label $x_j \in \{0, 1\}$. Therefore, a choice of instruments is given by N binary variables $\vec{x} \equiv (x_1, \ldots, x_N)$ and a

process is identified with 2^N N-tuples $\{\vec{a}_{\vec{x}}\}_{\vec{x}\in\{0,1\}^N}$, where $\vec{a}_{\vec{x}} := \vec{\omega}(\vec{\mathfrak{I}}_{\vec{x}})$. A deterministic ontological model is thus defined by a conditional probability distribution

$$p\left(\vec{\omega} \mid \beta\right) \equiv P\left(\{\vec{a}_{\vec{x}}\}_{\vec{x}\in\{0,1\}^N} \mid \beta\right) \tag{73}$$

that reproduces the operational probabilities via ω-mediation:

$$\sum_{\{\vec{a}_{\vec{x}}\}_{\vec{x}\in\{0,1\}^N}} p\left(\vec{\lambda} \mid \vec{\mathfrak{I}}, \{\vec{a}_{\vec{x}}\}_{\vec{x}\in\{0,1\}^N}\right) p(\{\vec{a}_{\vec{x}}\}_{\vec{x}\in\{0,1\}^N} \mid \beta) = p\left(\vec{\lambda} \mid \vec{\mathfrak{I}}, \beta\right), \tag{74}$$

where the sum should be understood as

$$\sum_{\{\vec{a}_{\vec{x}}\}_{\vec{x}\in\{0,1\}^N}} \equiv \sum_{a_0^1\in\mathfrak{I}_0^1} \sum_{a_1^1\in\mathfrak{I}_1^1} \cdots \sum_{a_0^N\in\mathfrak{I}_0^N} \sum_{a_1^N\in\mathfrak{I}_1^N} \tag{75}$$

and the "ontic" probabilities are given by

$$P\left(\vec{\lambda} \mid \vec{\mathfrak{I}}, \{\vec{a}_{\vec{x}}\}_{\vec{x}\in\{0,1\}^N}\right) = \prod_{j\in\mathcal{N}} \chi_{\mathfrak{I}^j}(\lambda_j) \sum_{\vec{x}\in\{0,1\}^N} \delta_{\vec{\lambda}\,\vec{a}_{\vec{x}}}, \tag{76}$$

$$\delta_{\vec{\lambda}\,\vec{a}_{\vec{x}}} := \prod_{j\in\mathcal{N}} \delta_{\lambda_j\,(a_{\vec{x}})_j}. \tag{77}$$

We now show that the conditional probabilities for the ontic process in our thermal model are given by

$$P\left(\{\vec{a}_{\vec{x}}\}_{\vec{x}\in\{0,1\}^N} \mid \beta\right) = \prod_{\vec{x}\in\{0,1\}^N} f_\beta(\vec{a}_{\vec{x}}), \tag{78}$$

where the operational frame function f_β is given by expression (71). To see that the conditional probabilities (78) provide an ontological model for the original thermal-state model, one can verify that, by putting together expressions (78) and (76) into Eq. (74), one indeed obtains the operational probabilities (70). Explicitly,

$$\begin{aligned}
&\sum_{\{\vec{a}_{\vec{x}}\}_{\vec{x}\in\{0,1\}^N}} p\left(\vec{\lambda} \mid \vec{\mathfrak{I}}, \{\vec{a}_{\vec{x}}\}_{\vec{x}\in\{0,1\}^N}\right) p(\{\vec{a}_{\vec{x}}\}_{\vec{x}\in\{0,1\}^N} \mid \beta) \\
&\quad= \prod_{j\in\mathcal{N}} \chi_{\mathfrak{I}^j}(\lambda_j) \sum_{\{\vec{a}_{\vec{x}}\}_{\vec{x}\in\{0,1\}^N}} \sum_{\vec{x}'\in\{0,1\}^N} \delta_{\vec{\lambda}\,\vec{a}_{\vec{x}'}} \prod_{\vec{x}\in\{0,1\}^N} f_\beta(\vec{a}_{\vec{x}}) \\
&\quad= \prod_{j\in\mathcal{N}} \chi_{\mathfrak{I}^j}(\lambda_j) \sum_{\vec{x}'\in\{0,1\}^N} \sum_{\vec{a}_{\vec{x}'}\in\vec{\mathfrak{I}}_{\vec{x}'}} \delta_{\vec{\lambda}\,\vec{a}_{\vec{x}'}} \left[\sum_{\{\vec{a}_{\vec{x}}\}_{\vec{x}\neq\vec{x}'}} \prod_{\vec{x}\in\{0,1\}^N} f_\beta(\vec{a}_{\vec{x}})\right] \\
&\quad= \prod_{j\in\mathcal{N}} \chi_{\mathfrak{I}^j}(\lambda_j) \sum_{\vec{x}'\in\{0,1\}^N} \sum_{\vec{a}_{\vec{x}'}\in\vec{\mathfrak{I}}_{\vec{x}'}} \delta_{\vec{\lambda}\,\vec{a}_{\vec{x}'}} f_\beta(\vec{a}_{\vec{x}'}) \left[\prod_{\vec{x}\neq\vec{x}'} \sum_{\vec{a}_{\vec{x}}\in\vec{\mathfrak{I}}_{\vec{x}}} f_\beta(\vec{a}_{\vec{x}})\right] \\
&\quad= \prod_{j\in\mathcal{N}} \chi_{\mathfrak{I}^j}(\lambda_j) f_\beta(\vec{\lambda}),
\end{aligned}$$

where we used the normalisation of the frame function, $\sum_{\vec{\lambda}\in\vec{\mathfrak{I}}_{\vec{x}}} f_\beta(\vec{\lambda}) = 1$ for every collection of instruments $\vec{\mathfrak{I}}_{\vec{x}} \equiv \left(\mathfrak{I}_{x_1}^1, \ldots, \mathfrak{I}_{x_N}^N\right)$.