

# The End of a Classical Ontology for Quantum Mechanics?

Peter W. Evans\*

School of Historical and Philosophical Inquiry, University of Queensland

## Abstract

In this paper, I argue that the Shrapnel–Costa no-go theorem undermines the last remaining viability of the view that the fundamental ontology of quantum mechanics is essentially classical: that is, the view that physical reality is underpinned by objectively real, counterfactually definite, uniquely spatiotemporally defined, local, dynamical entities with determinate valued properties, and where typically ‘quantum’ behaviour emerges as a function of our own in-principle ignorance of such entities. Call this view Einstein–Bell realism. One can show that the causally symmetric local hidden variable approach to interpreting quantum theory is the most natural interpretation that follows from Einstein–Bell realism, where causal symmetry plays a significant role in circumventing the nonclassical consequences of the traditional no-go theorems. However, Shrapnel and Costa argue that exotic causal structures, such as causal symmetry, are incapable of explaining quantum behaviour as arising as a result of noncontextual ontological properties of the world. This is particularly worrying for Einstein–Bell realism and classical ontology. In the first instance, the obvious consequence of the theorem is a straightforward rejection of Einstein–Bell realism. However, more than this, I argue that, even where there looks to be a possibility of accounting for contextual ontic variables within a causally symmetric framework, the cost of such an account undermines a key advantage of causal symmetry: that accepting causal symmetry is more economical than rejecting a classical ontology. Either way, it looks like we should give up on classical ontology.

## 1 Introduction

It should come as no surprise that many, perhaps even a good majority, of physicists after 1927 gave up on the view that the fundamental ontology of quantum mechanics is essentially classical: that is, the view that physical reality is underpinned by objectively real, counterfactually definite, uniquely spatiotemporally defined, local, dynamical entities with determinate valued properties, and where typically ‘quantum’ behaviour emerges as a function of our own in-principle ignorance of such entities. Let us call this position on the ontology of quantum theory Einstein–Bell realism. Despite the gloomy forecast for Einstein–Bell realists, it is well known that a class of responses to the canon of quantum no-go theorems, so-called causally

---

\*email: p.evans@uq.edu.au

symmetric local hidden variable approaches, plausibly rescues a large part of this classical picture of quantum theory.<sup>1</sup> Indeed, causally symmetric local hidden variable approaches arguably comprise the last refuge for Einstein–Bell realism, positioned as they are to navigate a classical ontology through Bell’s theorem, the Kochen–Specker theorem, and the PBR theorem.

Part of the appeal of causal symmetry in this context is that it circumvents one of the integral assumptions of these three no-go theorems—that the properties of some quantum system have definite values independently of the measurement context to which the system is to be subject. With respect to Bell’s theorem, this admits local hidden variables—or local *beables*, as Bell called them—and with respect to the Kochen–Specker theorem, this not only admits noncontextual hidden variables, but also provides a natural explanation for why quantum systems appear to be contextual (as contextuality arises from the specific epistemic constraints of causal symmetry). A more recent no-go theorem, due to Shrapnel and Costa (2018), undermines this case for noncontextual hidden variables. In short, the Shrapnel–Costa theorem removes the loophole open to ‘exotic causal structure’, and so implies that no ontological model, now including causally symmetric models, that satisfy the noncontextuality assumptions of the theorem can reproduce the statistical predictions of quantum mechanics. So in order to be a feasible model of such predictions, any causally symmetric ontology underpinning quantum behaviour must necessarily be contextual, along with the rest of the ontological models. Making matters worse is that the resulting form of the contextuality renders a natural explanation for this feature, as in the case of the Kochen–Specker theorem, much less plausible. In so far as causally symmetric local hidden variable approaches comprise the last refuge for Einstein–Bell realism, this contextuality is a concerning predicament for classical ontology. Indeed, I argue that this concerning predicament is as good as the end of a classical ontology for quantum mechanics.

The argument will proceed as follows. I begin in Section 2 by introducing the broad outline of causally symmetric local hidden variable approaches to the traditional no-go theorems, and I define Einstein–Bell realism. I focus in this section on contextuality, so introduce the ontological models framework and the operational formulation of the contextuality problem. In Section 3, I provide a brief outline of the history and development of the process matrix formalism and then go on to detail the Shrapnel–Costa theorem. I consider the three assumptions of the theorem that constrain ontological models that reproduce the statistical predictions of quantum mechanics— $\omega$ -mediation, instrument noncontextuality, and process noncontextuality—and briefly examine what a violation of each these assumptions implies. I consider in Section 4 what this means for Einstein–Bell realism and a classical ontology. I argue that, as a result of the Shrapnel–Costa theorem, the outlook is particularly worrying for Einstein–Bell realism and classical ontology. In the first instance, the obvious consequence of the theorem is a straightforward rejection of Einstein–Bell realism. However, more than this, I argue that, even where there looks to be a possibility of accounting for contextual ontic variables within a causally symmetric framework, the cost of such an account undermines a key advantage of causal symmetry: that accepting causal symmetry is more economical than rejecting a classical ontology.

---

<sup>1</sup>I shall not be providing an argument for or against any particular causally symmetric local hidden variable approach here. See (Friederich and Evans, 2019; Wharton and Argaman, 2020) for good reviews of such approaches. See also (Evans *et al.*, 2016), in which we argue in favour of what we call the Price–Wharton approach.

Either way, it looks like we should give up on classical ontology.

## 2 Causal Symmetry and Classical Ontology

### 2.1 The No-Go Theorems

Following the revolution of statistical mechanics in the late nineteenth century, one could surely be forgiven for expecting that the puzzles of quantum mechanics that emerged over the first few decades of the twentieth century would ultimately be explained by an underlying theory, comprised of hidden variables. This was a position for which Einstein, de Broglie, and others argued. It was not long before a no-go theorem jeopardising this position was proposed by von Neumann (1932), who purported to show that the probabilistic nature of quantum theory could not be a function of an underlying theory of hidden variables. Bohm (1952) showed by way of counterexample that there must have been something wrong with von Neumann’s theorem by proving that the predictions of his own hidden variable model were equivalent to the predictions of quantum theory. As a consequence of Bohm’s counterexample to von Neumann’s theorem, Bell developed a more precise no-go theorem—Bell’s theorem (Bell, 1966)—which clarifies that there can be no *local* hidden variable model that can match the predictions of quantum theory (and that satisfy the further reasonable assumptions of Bell’s theorem).

Bell’s comments on locality have attracted considerable attention (for instance, by Norsen (2011) and Wiseman and Cavalcanti (2017), to name just two), but this detail should not concern us here. The basis of Bell’s assumption is that the specification of local beables at spacelike separation must be independent conditioned on their causal pasts. There is, however, another assumption in Bell’s theorem that is significant for our current purposes. This further assumption is known as *measurement independence*, and it requires that any hidden variables underlying a quantum system must remain statistically independent of the choice of measurement settings to which that system is subject as part of the experimental procedure (Friederich and Evans, 2019). If one were to reject the assumption of measurement independence, one could maintain Bell’s assumption of locality by permitting beables to be correlated explicitly with the choice of measurement setting, even if that choice occurs in the future of the quantum system.<sup>2</sup> Such a correlation can then be accounted for by giving up the conventional assumption that causation must occur only in the forward temporal direction, such that the choice of measurement setting might causally influence the hidden variables underlying the quantum state in the backward temporal direction. Thus, the resulting causal symmetry denies one of the assumptions underpinning Bell’s theorem, measurement independence, without violating the assumption of locality. This move rescues the possibility of local hidden variables, so long as one admits symmetric causal influences, both forwards and backwards in time.

Very shortly after Bell’s theorem followed the Kochen–Specker theorem (Kochen and Specker, 1967), which will be particularly relevant for our discussion below. This theorem states that

---

<sup>2</sup>Another, metaphysically distinct, way to reject the assumption of measurement independence is to remove the freedom of agents to choose measurement settings arbitrarily, leading to superdeterministic local hidden variables approaches (Bell, 1981, 1990; Norsen, 2011). Causal symmetry is generally more favourable than superdeterminism on account of the fact that one can explicitly preserve the free choice of agents by rejecting the assumption of strictly forwards-in-time causality.

any hidden variable model that reproduces the predictions of quantum theory must be *contextual*. A model or theory is said to be contextual when the properties attributed to some system described by the theory are dependent upon the means of realising some value for that property, or *context* of measurement or observation of those properties, beyond the actual observation itself. Thus, we can think of a *noncontextual* ontological property as one that can be unambiguously distinguished experimentally. A relevant example of a context of measurement beyond the measurement itself would be the set of further properties, if any, that are measured in conjunction with the measurement. Since in classical mechanics there is no such dependence on measurement context for properties attributed to systems as a result of some measurement, classical mechanics is noncontextual. According to the Kochen–Specker theorem, any model of a quantum system that assumes that measurements deterministically uncover the values of pre-existing dynamical variables—such as beables—must be contextual. So, just as Bell’s theorem shows that there can be no *local* hidden variable model satisfying his assumptions that can match the predictions of quantum theory, the Kochen–Specker theorem shows that there can be no noncontextual deterministic local hidden variable model that can match the predictions of quantum theory.

However, just as we saw above that rejecting Bell’s assumption of measurement independence, and so permitting correlations that arise as a result of symmetric causal influences, rescues local hidden variables from Bell’s theorem, hypothesising such causal symmetry allows a similar rejection of the assumption of measurement independence for the Kochen–Specker theorem. Under the assumption that the correlation between the value of any pre-measurement hidden variables and the context of measurement can be explained by a causal influence directed from the latter to the former, we can take the measurement process to be *bringing about* the determinate valued properties that constitute the hidden variables, rather than uncovering independently existing variables (Friederich and Evans, 2019). This account has the added bonus of providing a ‘natural’ explanation for why quantum theory is contextual: the dependence of any measured values on the eventual context of measurement should be expected in a regime that admits symmetric causal influences. This is because the natural epistemic constraints that arise when describing a physical system whose complete description depends upon a future boundary condition result in an impoverished, statistical description of the system conditioned only upon the initial boundary condition (see, for instance, (Wharton, 2010)). Indeed, Wharton (2014, p.203) suggests that contextuality just is the failure of measurement independence.

So far this brief historical narrative is straightforward. However, since contextuality is at the heart of the Shrapnel–Costa theorem, we need to develop further layers to this narrative. The particular way of understanding contextuality adopted by Shrapnel and Costa has its origin in the work of Spekkens (2005). To get to grips with this treatment of contextuality, we need to say a few words about the ontological models framework.

## 2.2 Ontological Models

The ontological models framework was first introduced in Spekkens (2005), and developed further in Harrigan and Spekkens (2010) (see also (Leifer, 2014) and (Ringbauer, 2017) for good reviews), and provides a formalisation of our notion of classical ontology in an operational model of quantum processes (Friederich and Evans, 2019). As an operational model, the framework specifies a set of possible preparations, transformations, and measurements, and associated outcome probabilities, to describe the observed statistics over the possible outcomes. In addition, the framework also specifies an ontological model to account for the observed statistics in the following way.

The quantum state,  $\psi$ , is the description we give to the quantum system after the preparation procedure. We assume, however, that we can completely specify the properties of the quantum system as a result of preparation in terms of its actual *ontic* state,  $\lambda$ , which arises via a classical probability density,  $\mu_\psi(\lambda)$ , over the set of possible ontic states,  $\Lambda$ . We can then specify the outcome probabilities for each measurement procedure,  $M$ , conditional on  $\lambda$ , such that the outcomes,  $\{m\}$ , are independent of the preparation procedure. We say that  $\lambda$  *screens-off* the preparation from the measurement so that measurement outcomes only depend on the ontic state. Leifer and Pusey (2017) call this feature  $\lambda$ -*mediation*, as the ontic state completely mediates any correlations between preparation procedures and measurement outcomes. Importantly for our narrative here is that, explicitly according to the ontological models framework,  $\lambda$  does not causally depend on the future choice of measurement procedure,  $M$ ,

$$P(\lambda | M) = P(\lambda), \quad (1)$$

otherwise  $\lambda$  could not screen off the  $\{m\}$  from  $\psi$ . Finally, as a result of this setup, the eventual operational statistics given by the conditional outcome probabilities need to reproduce the quantum statistics.

Let us add to this a useful distinction from the ontological models framework between what Harrigan and Spekkens (2010) call  $\psi$ -ontic and  $\psi$ -epistemic interpretations of the wavefunction. In a  $\psi$ -ontic interpretation, each ontic state is consistent with a single quantum state. In a  $\psi$ -epistemic interpretation, multiple distinct quantum states can be consistent with a single ontic state. A further orthogonal distinction can be overlaid across this. A  $\psi$ -complete interpretation takes the quantum state to provide a complete description of objective reality (that is, there is a one-to-one correspondence between quantum states and ontic states), while a  $\psi$ -incomplete interpretation requires that the quantum state be supplemented with additional ontic degrees of freedom (such as, for instance, in typical  $\psi$ -ontic pilot-wave interpretations). Given this latter distinction,  $\psi$ -epistemic interpretations are naturally  $\psi$ -incomplete, and render the wavefunction a representation of the knowledge of the user of the formalism rather than a representation of objective reality. For  $\psi$ -epistemic interpretations in the ontological models framework, there exists an ontic state underlying the wavefunction.<sup>3</sup>

We can apply this framework to our analysis of Bell's theorem above. The ontological

---

<sup>3</sup>Beyond the ontological models framework, a  $\psi$ -epistemic wavefunction can be given an anti-realist or operationalist interpretation that makes no such claim for a deeper underlying objective reality.

models framework is explicit that  $\lambda$  does not causally depend on the measurement procedure (1), which is, of course, simply Bell’s assumption of measurement independence. As we saw above, by denying this assumption and accounting for the correlation between hidden variables, or the ontic state  $\lambda$ , and the measurement procedure by way of causal symmetry, we can rescue the assumption that any beables are local. As a result of this move, then, we can see that the admission of causal symmetry is a ready made violation of the assumption of a strict temporal and causal order for the ontological models framework (Ringbauer, 2017, p.94).

As a very brief aside, and before we continue with our explication of contextuality in quantum theory, we now have the tools to summarily deal with the causal symmetry response to the PBR theorem (Pusey *et al.*, 2012). The PBR theorem states, in short, that the ontic states of any interpretation of quantum mechanics that fits within the Bell framework and reproduces the Born rule must be in one-to-one correspondence with the quantum states. This, of course, is simply the definition of a  $\psi$ -ontic interpretation of the wavefunction. Thus, the PBR theorem appears to be ruling out  $\psi$ -epistemic approaches to the wavefunction. However, since the PBR theorem is framed in the ontological models framework, it is straightforward to note as before that it does not apply to causally symmetric approaches. Just as with Bell’s theorem, admitting causal symmetry rescues the possibility of an interpretation of the wavefunction as an epistemic representation of an underlying reality of local beables.

Let us return then to contextuality. Spekkens (2005) pioneered a new way of thinking of contextuality beyond the formulation of Kochen and Specker (1967). According to Spekkens’ characterisation, an ontological model of an operational theory is *noncontextual* when operationally equivalent experimental procedures have equivalent representations in the ontological model (Spekkens, 2005, p.1). Thus, contextuality on this account implies that operationally equivalent experimental preparation procedures may correspond to inequivalent ontic state representations (and where the state additionally depends on the context of measurement). Defining contextuality in this way, rather than in terms of the deterministic uncovering of the values of pre-existing dynamical variables, renders the assumption of noncontextuality as a principle of parsimony: no ontological difference without operational difference.<sup>4</sup> Despite this difference in flavour, Spekkens’ diagnosis of hidden variable approaches concludes with the same result. Assuming that there are ontic properties underlying the observed statistical behaviour of quantum systems, and that the associated ontic states are distinct just when there are corresponding operational differences, cannot account for the statistical behaviour entailed by quantum theory. The result is then that there is no noncontextual ontological model that can reproduce the observed statistics of quantum theory.

Just as we saw above with respect to Bell’s theorem and the PBR theorem, however, in so far as the no-go theorems can be characterised in the ontological models framework, the no-go theorems are undermined by the assumption of causal symmetry. Thus Spekkens’ more general understanding of contextuality in terms of the ontological models framework suffers the same response: causal symmetry is a categorical rejection of the assumption of a strict temporal

---

<sup>4</sup>Spekkens (2019) argues that this generalised noncontextuality is an embodiment of Leibniz’ Principle of the Identity of Indiscernibles, and labels this principle the *Leibnizian methodological principle*. Schmid *et al.* (2020) formalise this principle under the term *Leibnizianity*. We consider this principle again in Section 4.

and causal order that underpins the ontological models framework. Noting this is the main result of this section of the paper: in so far as there is a motivation to avoid the consequences of the above no-go theorems, particularly to maintain some semblance of a classical ontology, hypothesising causal symmetry is a strategy that is well placed to achieve this.

The no-go theorems are ordinarily taken to demonstrate that the underlying conceptual and ontological framework of quantum theory cannot be completely classical, and cannot be about local, noncontextual hidden variables. Any hope one might have for hanging on to such a classical ontology rests with causally symmetric local hidden variable approaches, as per our story here. Whilst such approaches are admittedly unorthodox, and even perhaps in a sense ‘nonclassical’ on account of causal symmetry, what is significant about such approaches is that they have the potential to rescue an objectively real, counterfactually definite, uniquely spatiotemporally defined, local, noncontextual (or where any contextuality is underpinned by noncontextual epistemic constraints), determinate valued ontology, where typically ‘quantum’ behaviour emerges as a function of our own in-principle ignorance of such entities (Friederich and Evans, 2019).

As I propose in Evans (2015) (and also in Friederich and Evans (2019)), the distinction that Quine (1951) draws between the ‘ontology’ and the ‘ideology’ of a theory is an apt distinction to understand the value of the above argument. According to Quine, the ideology of a theory is comprised of the set of ideas that can be expressed by a theory, and thus ideological economy is then a measure of the economy of primitive undefined statements employed to reproduce this ideology. Put in these terms, one of the consequences of the above narrative is that the ideology of causal symmetry is more economical than a rejection of classical ontology. The main goal of the current work is to make clear how the Shrapnel–Costa no-go theorem significantly jeopardises this position. In short, motivated by the ‘get out of jail free’ card that causal symmetry is able to play with respect to the no-go theorems, Shrapnel and Costa develop a no-go theorem that closes off the ‘measurement independence’ loophole to any approach that assumes an ‘exotic’ form of causality.

Before we consider the Shrapnel–Costa theorem in Section 3, let us once more explicitly state our characterisation of classical ontology by adopting the framework which my collaborators and I have elsewhere called Einstein–Bell realism (Evans *et al.*, 2016), bringing together the above analysis of the no-go theorems.

### 2.3 Einstein–Bell Realism

Consider the following conditions from Evans *et al.* (2016), called the *Einstein–Bell conditions*, adapted here to nonrelativistic quantum theory:

- (i) Quantum mechanical probabilities are epistemic;
- (ii) Quantum mechanics is local;
- (iii) Quantum mechanics is consistent with the no-go theorems.

We can think of Einstein–Bell realism as the conjunction of these conditions with the assumption that there is an objective reality.

There is good reason for calling this view Einstein–Bell realism. It should be clear to see that the first condition epitomises the sentiment that Einstein expressed to Born in 1926 that “God does not play dice”.<sup>5</sup> This also has implications for any ensuing interpretation of the wavefunction. Consider the distinction above between  $\psi$ -ontic and  $\psi$ -epistemic interpretations of the wavefunction. Harrigan and Spekkens (2010) argue that Einstein, in his more sophisticated arguments for the incompleteness of quantum mechanics, advocated for a realist  $\psi$ -epistemic interpretation. Despite the fact that there are some  $\psi$ -ontic,  $\psi$ -incomplete interpretations—like pilot-wave interpretations—in which probabilities arise as a function of our ignorance over the full ontic degrees of freedom, and so are epistemic, the basic probabilistic concepts that occur in all  $\psi$ -epistemic interpretations are clearly also themselves epistemic. Thus,  $\psi$ -epistemic interpretations are natural bedfellows for epistemic notions of probability (Evans *et al.*, 2016, p.6).

The second assumption epitomises not only Einstein’s implicit appeal to consistency between quantum theory and relativity (as in, for instance, Einstein *et al.* (1935)), but also the exclusion of spacelike separated influences formalised in Bell’s own theorem (as in, for instance, Bell (1976)). For clarity, and following Evans *et al.* (2016, p.5), we can take the assumption here to mean that physical bodies and influences can only follow timelike or null spacetime trajectories, and so these trajectories must obey Lorentz covariance. The third assumption reflects that a plausible interpretation of quantum theory must obey the constraints established by the above no-go results, that is, from Bell’s theorem, the Kochen–Specker theorem, and the PBR theorem.

Employing the Einstein–Bell conditions, we argue in Evans *et al.* (2016) that causally symmetric local hidden variable approaches constitute a unique realist interpretation satisfying these constraints. Our logic there is straightforward. The first condition requires a  $\psi$ -incomplete interpretation, as epistemic probabilities cannot be accommodated by  $\psi$ -complete interpretations. The second and third conditions together exclude  $\psi$ -ontic interpretations, as such interpretations cannot be local according to the no-go theorems, implying that only  $\psi$ -epistemic interpretations can meet the Einstein–Bell conditions. Furthermore, a causally symmetric approach circumvents the results of Bell’s theorem, and so can be a local hidden variable theory, also circumvents the PBR theorem, and so permits a  $\psi$ -epistemic wavefunction, and also contains an explicit contextuality of the ontic state on the experimental procedure, so fits within the bounds given by the Kochen–Specker theorem. So we can see that the commitments of Einstein–Bell realism capture what we mean by local beables, or by noncontextual local hidden variables: a commitment to a local-realist ontology underlying a  $\psi$ -epistemic wavefunction. Causally symmetric hidden variable approaches are, according to the arguments in Evans *et al.* (2016), the only approaches that fit these constraints. As we will now see, however, these constraints have just gotten prohibitively tighter.

### 3 The Shrapnel–Costa No-Go Theorem

Shrapnel and Costa (2018) argue that exotic causal structures—such as causal symmetry—are incapable of explaining quantum behaviour arising as a result of noncontextual ontological

---

<sup>5</sup>To be clear, I take this as a statement of determinacy rather than determinism.



properties of the world. As we have just seen, one of the key underlying assumptions of the Kochen–Specker theorem is that quantum phenomena arise on a fixed background forwards-in-time causal structure, and so do not preclude the possibility of more exotic causal structures providing an ontologically classical (noncontextual) explanation. This leaves open the possibility of symmetric causal structure allowing noncontextual ontological properties to underpin quantum behaviour, and also providing a natural explanation for *why* typically quantum behaviour can arise from such a noncontextual ontology.

The Shrapnel–Costa theorem is stronger than the Kochen–Specker theorem, as it closes off the possibility of exotic causal structure providing just such a noncontextual explanation of quantum behaviour. In fact, it shows that *any* ontology underpinning quantum behaviour must be contextual; moreover, “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, which is responsible for the experimentally observed statistics” (Shrapnel and Costa, 2018, p.2). In order to take account of the possibility of exotic causal structure, the Shrapnel–Costa theorem first generalises the ontological models framework, and then employs the *process matrix formalism*, which is suited to describing processes with indefinite causal structure. To get to the heart of the Shrapnel–Costa theorem, then, let us begin by reviewing the process matrix formalism.

### 3.1 Process Matrix Formalism

The development of the process matrix formalism is punctuated by a number of independent redevelopments. The roots of the formalism stretch back to Accardi (1976) and Lindblad (1979) on the dynamics of open quantum systems and quantum stochastic processes in the 1970s. This early work is framed in the language of algebraic quantum field theory. The development of quantum information theory at the beginning of this century allowed for some of the ideas corresponding to those from the dynamics of open quantum systems to be recast into a modern form, whereby, for instance, ‘non-Markovian quantum stochastic processes’ has become ‘quantum channels with memory’. In the field of experimental quantum information, process matrices began to be popularised in the analysis of quantum process tomography of quantum optical systems (see, for instance, (O’Brien *et al.*, 2004; Riebe *et al.*, 2006; Kok *et al.*, 2007; Mohseni and Lidar, 2007; Riebe *et al.*, 2007))—whereby processes matrices were usually obtained experimentally from quantum state tomography (O’Brien *et al.*, 2004). However, much of the contemporary discussion regarding the process matrix formalism can be traced back to the seminal work of Kretschmann and Werner (2005), couched in the language of quantum information theory, and developed largely independently of the earlier work on open quantum systems.

In short, a process matrix—to be defined below—is a way of representing the state transformation denoted by the completely positive map,  $\mathcal{M}^A(\rho)$ , generated by the evolution of a quantum state,  $\rho$ , from an input state space,  $\mathcal{H}^{A_1}$ , to an output state space,  $\mathcal{H}^{A_2}$ , by some quantum operation,  $\mathcal{M}^A : \mathcal{L}(\mathcal{H}^{A_1}) \rightarrow \mathcal{L}(\mathcal{H}^{A_2})$ . A key step to developing the process matrix formalism is the realisation that the Choi–Jamiołkowski isomorphism (Jamiołkowski, 1972;

Choi, 1975), which establishes a correspondence between linear maps and linear operators, allows the linear map  $\mathcal{M}^A : \mathcal{L}(\mathcal{H}^{A_1}) \rightarrow \mathcal{L}(\mathcal{H}^{A_2})$  to be conveniently rewritten as a linear operator  $M^{A_1 A_2} \in \mathcal{L}(\mathcal{H}^{A_1} \otimes \mathcal{H}^{A_2})$ .<sup>6</sup> Following on from Kretschmann and Werner (2005), the ‘Pavia group’ employed this Choi–Jamiołkowski representation to describe transformations on a network of quantum gates, and developed a methodology for optimising such *quantum circuits* using what they call a *quantum comb*: a temporally ordered physical process with a quantum memory (Chiribella *et al.*, 2008a, 2009). (Incidentally, the quantum comb has been independently redeveloped no less than two more times (Modi, 2012; Luchnikov *et al.*, 2018).) They further demonstrated that simple quantum circuits can be physical implementations of what they call *quantum supermaps*, mapping an input quantum operation to an output quantum operation (Chiribella *et al.*, 2008b).

At about the same time, and independently, Gutoski and Watrous (2007) also employed the Choi–Jamiołkowski representation to represent what they call *quantum strategies*: a specification of the exchange and processing of quantum information in a quantum process. This development is largely equivalent to the quantum comb (along with the operator tensor formulation of quantum theory (Hardy, 2012), which contains close similarities). The Pavia group, however, went on in the ensuing years to develop and analyse the quantum control of temporal order in the *quantum switch*, where two operations are enacted in a quantum superposition of the two possible temporal orders (Chiribella, 2012; Colnaghi *et al.*, 2012; Chiribella *et al.*, 2013). The significance of the quantum switch here is that the process matrix formalism developed as the ideal formal system for describing such indefinite causal order.

Hardy (2005) introduces the notion of indefinite causal structure in the context of quantum gravity, and Oeckl (2003), also motivated by quantum gravity, considers a “general boundary” formulation of quantum mechanics that does not assume *a priori* causal structure (and which overlaps significantly in the quantum context with the two-state vector formalism (Aharonov *et al.*, 1964)). However, it is Oreshkov *et al.* (2012) who analyse indefinite causal order in the context of quantum processes: they employ the Choi–Jamiołkowski representation, and a local causal direction, to derive a causal inequality violated by timelike and spacelike correlations with a global causal direction. It is this work on process matrices and indefinite causal order that has in part led to the development of quantum causal modelling (Leifer and Spekkens, 2013; Cavalcanti and Lal, 2014; Costa and Shrapnel, 2016; Allen *et al.*, 2017), the framework within which Shrapnel and Costa build their no-go theorem.

The utility of the process matrix formalism is that it provides a framework to describe communication tasks between parties that lack a definite causal order. Given a set of parties each residing in their own laboratory,  $A, B, \dots$ , one assumes that each party is able to act locally on some system that passes once through their laboratory, where local operations are described by ordinary quantum mechanics. No assumption is made concerning the relative spatiotemporal arrangement of the laboratories, nor that there is ultimately some causal structure within which the laboratories are positioned.

Following Oreshkov *et al.* (2012), the operations that any such party can perform are delineated by a *quantum instrument*,  $\mathcal{M}^A$ , which induces—by way of a unitary transformation and

---

<sup>6</sup>See (Adlam, 2020) for more on the interpretation of the Choi–Jamiołkowski isomorphism.

projective measurement—a transformation from input to output,  $\mathcal{M}_i^A$ , indexed by the *outcome*,  $i = 1, \dots, n$ .  $\mathcal{M}_i^A$  is a completely positive (CP) trace-non-increasing map. The set of all CP maps for some instrument—which when taken together,  $\{\mathcal{M}_i^A\}_{i=1}^n \equiv \mathcal{M}^A$ , exemplify the fact that with probability one there must be some outcome of the application of the instrument—is a CP and trace-preserving (CPTP) map.

For a set of parties, the set of all outcomes across the laboratories corresponds to a set of CP maps,  $\mathcal{M}_i^A, \mathcal{M}_j^B, \dots$ , and the complete list of probabilities  $P(\mathcal{M}_i^A, \mathcal{M}_j^B, \dots)$  for all  $i, j, \dots$  is what Oreshkov *et al.* (2012, p.3) call a *process*. The process thus exemplifies all the operational correlations between local laboratories. An important assumption that Oreshkov et al. make at this point is that each such joint probability is noncontextual, in the sense that the joint probability for any set of CP maps is not dependent on the local detail of any particular instrument  $\mathcal{M}^A$ . It is most obvious here in this characterisation of ‘process’ the way in which it can be seen as a generalisation of ‘state’: both can be seen as encoding a list of probabilities over a set of measurement scenarios.

Since, as above, the linear maps  $\mathcal{M}_i^A$  and  $\mathcal{M}_j^B$  can be represented, using the Choi–Jamiołkowski isomorphism, by the linear operators  $M_i^{A_1A_2}$  and  $M_j^{B_1B_2}$ , respectively, the joint probability for two measurement outcomes ( $i$  at  $A$  and  $j$  at  $B$ ) can be expressed as a function of the corresponding Choi–Jamiołkowski operators (Oreshkov *et al.*, 2012, p.4):

$$P(\mathcal{M}_i^A, \mathcal{M}_j^B) = \text{Tr} \left[ W^{A_1A_2B_1B_2} (M_i^{A_1A_2} \otimes M_j^{B_1B_2}) \right]. \quad (2)$$

Here,  $W^{A_1A_2B_1B_2}$  is known as the *process matrix*, under the condition that it is positive semi-definite ( $W^{A_1A_2B_1B_2} \geq 0$ )—which embodies the constraint that the instruments  $\mathcal{M}^A$  and  $\mathcal{M}^B$  are CPTP maps—and its trace is one (Wood, 2009).

The process matrix is readily understood as a generalisation of a density matrix—and so can be seen as a generalisation and extension of the notion of state—and the trace rule (2) as a generalisation of the Born rule (Oreshkov *et al.*, 2012). Further, more restrictive assumptions allow one to reduce the process matrix to the quantum state (when the output systems are one-dimensional), to reflect the characterisation of a quantum comb or network (when a definite causal order is fixed), or to represent quantum channels with memory (when only unidirectional signalling is possible).

With this brief survey of the process matrix formalism, let us turn our attention to the Shrapnel–Costa no-go theorem.

### 3.2 Causation Does not Explain Contextuality

Shrapnel and Costa begin by outlining a generalised operational framework to reflect the ontological models framework of Spekkens (2005). However, they are interested in replacing the operational notions of preparation, transformation, and measurement procedures with more temporally and causally neutral concepts. They replace these operations with the general notion of *local controllables*,  $\tilde{\mathcal{J}}^A$ , where each local controllable is indexed to a *local region*,  $A, B, \dots$ , and each choice of local controllable is labelled by an *outcome*,  $a, b, \dots$ . Furthermore, the physical features of the world external to the system, and independent of the choice of local

controllables, including “any global properties, initial states, connecting mechanisms, causal influence, or global dynamics”, responsible for correlating outcomes between local regions they call the *environment*,  $\tilde{W}$ . Importantly, any variable correlated with the choice of local controllable is necessarily considered an outcome and cannot be part of the environment (Shrapnel and Costa, 2018, p.5).

They define an *event*,  $M^A: =[(a, \tilde{\mathcal{J}}^A)]$ , as an operational equivalence class of pairs of outcome and local controllable, such that the joint probabilities over outcomes are equivalent for all possible outcomes and local controllables in all the other regions and for all environments (Shrapnel and Costa, 2018, p.7). In the process matrix framework (that is, in the quantum context) this is analogous to the role played by the CP maps  $M_i^A$ , which we can call here *quantum events*. They define an *instrument*,  $\mathcal{I}^A$ , similarly as an operational equivalence class of lists of possible (that is, with non-zero probability) events,  $M^A$ ; that is,  $\mathcal{I}^A: =\{M_i^A\}$ . This is analogous to the role played by the CPTP maps  $\mathcal{M}^A$  in the process matrix formulation, which recall are labelled *quantum instruments*.

Moreover, just as we noted above that the joint probability  $P(\mathcal{M}_i^A, \mathcal{M}_j^B, \dots)$  is noncontextual, in the sense that it is not dependent on the local detail of any particular instrument  $\mathcal{M}^A$ , the probability of some event  $M^A$ , so long as  $\mathcal{I}^A$  renders  $M^A$  possible, is independent of the particular instrument  $\mathcal{I}$ . Thus, as above, correlations between events across different regions are not a function the details of the instrument, which itself specifies events that do not happen, and so events screen-off instruments. Shrapnel and Costa call this *operational instrument equivalence* and note that it is equivalent to noncontextuality in the former sense (Shrapnel and Costa, 2018, p.13).

Finally, they define a process,  $W: =[\tilde{W}]$ , as an operational equivalence class of environments, such that the joint probabilities over outcomes given local controllables are equivalent across the equivalence class of environments. In the quantum context,  $W$  is the *process matrix*. This generalisation of outcomes, local controllables, and the environment into events, instruments, and the process serves to operationalise any joint probability distribution to allow the creation of an ontological model—along the lines of the ontological models framework—that underlies the distributions that we take to account for the observed statistics.

The major shift enacted by Shrapnel and Costa from the ontological models framework of Spekkens (2005) and Harrigan and Spekkens (2010) to ontological models in the process matrix framework is to move away from the idea that the ‘state’ encodes the ontology of some system towards the idea that more general properties of the environment are responsible for mediating correlations between the regions. As such, they replace the ontic state  $\lambda$  with the *ontic process*  $\omega$  (Shrapnel and Costa, 2018, p.8):

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.

Those aspects of reality that the ontic process captures are those parts of the environment that are not within the control of the experimenters, like initial conditions, causal influences, and global dynamics.

Shrapnel and Costa then make three natural assumptions that they take an ontological model in their framework to obey. Firstly, they replace the notion of  $\lambda$ -mediation (as per (Leifer and Pusey, 2017) above) with the notion of  $\omega$ -mediation, in which the ontic process  $\omega$  completely specifies the properties of the environment that mediate correlations between regions, and screens off outcomes produced by local controllables from the rest of the environment (Shrapnel and Costa, 2018, p.8):

$$P(a, b, \dots | \tilde{\mathcal{J}}^A, \tilde{\mathcal{J}}^B, \dots \tilde{W}) = \int d\omega P(a, b, \dots | \tilde{\mathcal{J}}^A, \tilde{\mathcal{J}}^B, \dots) P(\omega | \tilde{W}). \quad (3)$$

Secondly, they define the notion of *instrument noncontextuality* as a law of parsimony (much like (Spekkens, 2005)): operationally indistinguishable pairs  $(a, \tilde{\mathcal{J}}^A)$ ,  $(a', \tilde{\mathcal{J}}'^A)$  should remain ontologically indistinguishable. That is,  $\forall b, c, \dots \tilde{\mathcal{J}}^B, \tilde{\mathcal{J}}^C, \dots, \omega$  (Shrapnel and Costa, 2018, p.9):

$$P(a, b, \dots | \omega, \tilde{\mathcal{J}}^A, \tilde{\mathcal{J}}^B, \dots) = P(a', b, \dots | \omega, \tilde{\mathcal{J}}'^A, \tilde{\mathcal{J}}^B, \dots). \quad (4)$$

This allows them to define a probability distribution on the space of events, conditional on instruments and the ontic process,  $P(M^A, M^B, \dots | \omega, \tilde{\mathcal{J}}^A, \tilde{\mathcal{J}}^B, \dots)$ , in terms of a function that maps events to probabilities. As Shrapnel and Costa point out, instrument noncontextuality is formally identical to operational instrument equivalence, except for the fact that instrument noncontextuality includes the ontic process.

Thirdly, they define the notion of *process noncontextuality*: operationally indistinguishable  $\tilde{W}, \tilde{W}'$  should remain ontologically indistinguishable (Shrapnel and Costa, 2018, p.9).

$$P(\omega | \tilde{W}) = P(\omega | \tilde{W}'). \quad (5)$$

Again, this allows them to define a probability distribution on the space of ontic processes, in terms of a function that maps ontic processes to probabilities.

The Shrapnel–Costa no-go theorem is then that there can be no ontological model that satisfies  $\omega$ -mediation, instrument noncontextuality, and process noncontextuality. They argue as follows. As we have just noted, each of instrument and process noncontextuality defines a function that maps from the space of events,  $\{M^A, M^B, \dots\}$ , and ontic processes,  $\omega$ , respectively, to probabilities. However, the two noncontextuality assumptions force these functions to be ordinary positive probability distributions. Since quantum expectation values cannot be expressed in this way, no instrument and process noncontextual ontological model can reproduce the quantum statistical predictions.

### 3.3 Interpreting the Result

So what does this result mean, exactly? Well, to begin with, beyond pointing out the intended consequence of their theorem, Shrapnel and Costa do not speculate on further consequences in any great detail. The intended consequence is that, since preparations, transformations, and measurements have been replaced by local controllables, there is no further assumption in the no-go theorem that  $\omega$  is correlated with some controllables but independent of others. Recall that this is the form of the ‘loophole’ in the orthodox ontological models framework

through which we are able to thread causally symmetric local hidden variable approaches to defeat the nonclassical consequences of the no-go theorems from Section 2. The part of the theorem doing most of the heavy lifting on this point is  $\omega$ -mediation. By replacing  $\lambda$ -mediation with  $\omega$ -mediation, the relevant correlations are not simply a function of the quantum state, but the agent-independent rules or laws that we take the environment to contribute to the dynamical behaviour of a system, and the connection between local action and observed events. Where causal symmetry is ideally placed to circumvent  $\lambda$ -mediation, no such causal assumption can do so for  $\omega$ -mediation. Thus, this loophole is closed off in the Shrapnel–Costa theorem, rendering causally symmetric approaches just as contextual as the rest of the models captured by the ontological models framework (Friederich and Evans, 2019). So causally symmetric local hidden variable approaches, on account of being ontological models, must violate one of the assumptions of the Shrapnel–Costa theorem to hope to match the statistical predictions of quantum mechanics.<sup>7</sup> In so far as this sets a challenge to causally symmetric approaches, the discussion in the next section explores the possibility of meeting this challenge. To this end, while Shrapnel and Costa note only briefly the consequences of violating each of their assumptions, let us consider precisely what such violations entail.

The consequences of violations of the assumption of  $\omega$ -mediation are not limited to transgressions against a realist attitude towards the ontology of quantum systems. As we noted just above, the assumption of  $\omega$ -mediation insinuates that there are observer-independent aspects of the world, such as boundary conditions and global dynamics, that are ‘there’ to be discovered by the experimental procedure. Violations of this assumption do not just offend realist attitudes towards the state, then, but would require a radical rethink of the nature of scientific inquiry and our role as observers in that process.

Shrapnel and Costa have notable things to say about the assumption of instrument contextuality. Firstly, it is interesting to note the possibilities that they consider in which ontological models that satisfy instrument noncontextuality could be interpreted, from an ordinarily time-oriented perspective, as contextual (Shrapnel and Costa, 2018, p.11–13). This is the same phenomenon at play as the one employed by causally symmetric approaches with respect to measurement noncontextuality. As we saw above in Section 2.1, the ‘added bonus’ of causally symmetric approaches to contextuality in the Kochen–Specker theorem was that it provided a natural explanation for this contextuality in terms of epistemic constraints that arise from the ordinary temporal orientation of observers. However, the nonextendibility result in the case of measurement noncontextuality—that no noncontextual extension of quantum theory can provide more accurate predictions of outcomes (Chen and Montina, 2011)—holds in the case of instrument noncontextuality, too. That is, no instrument noncontextual hidden variable can provide more information than is contained in the process matrix (Shrapnel and Costa, 2018, p.17). This result rules out nontrivial hidden variable extensions such that obtaining greater predictive power for quantum theory can only be achieved by the addition of *contextual* variables.

Violations of the assumption of instrument noncontextuality would imply that correlations

---

<sup>7</sup>Superdeterministic hidden variable models, also on account of being ontological models, fare no better at meeting this challenge.

between events across different regions depend upon the details of the quantum instrument—in particular, on the CP maps that are not employed as part of the choice of local controllable—and so on events that do not in fact happen. Interestingly, this would imply that events do not screen-off instruments, and so lends weight to the idea that contextuality is a species of fine-tuning (Cavalcanti, 2018). This flavour of noncontextuality has strong similarities to preparation and measurement noncontextuality, and thus the ontological consequences of violations of these have received considerable attention already (see, for instance, (Harrigan and Rudolph, 2007; Leifer, 2014)). Since these do not represent a particularly novel type of consequence for the Shrapnel–Costa theorem, I will not labour these consequences here.

The consequences of violations of the assumption of process noncontextuality are certainly novel. Process contextuality implies that operationally equivalent arrangements of an experiment do not necessarily lead to equivalent ontic descriptions of that experiment. What is significant about this is that included in this ontic description are (Shrapnel and Costa, 2018, p.18):

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.

Thus, it seems that process contextuality would have a devastating effect on our ability to conduct orthodox scientific inquiry. The implication here is that operationally equivalent experimental arrangements may be realised by inequivalent ontic states, global properties, causal mechanisms, laws, or boundary conditions—indeed, any part of the environment that is not within the control of the experimenters. It is difficult to imagine exactly what this means, ontologically speaking. However, I speculate in the next section that one solution to a previously identified problem with a particular causally symmetric approach may provide a suggestion as to what process contextuality could amount.

## 4 The End of a Classical Ontology?

We are now in a position to assess the prospects for causally symmetric local hidden variable approaches to quantum theory in the face of the Shrapnel–Costa theorem and, in so far as such approaches are the only approaches that fit the constraints of Einstein–Bell realism, with them the prospects for Einstein–Bell realism and a classical ontology. To some, the loss of the possibility of a classical ontology is surely no news at all; as I noted in the opening sentence to this work, many physicists gave up on the view that the fundamental ontology of quantum mechanics is essentially classical after 1927, and certainly after the Bohr–Einstein debates ran their course beyond 1935 (Einstein *et al.*, 1935; Bohr, 1935).

Thus the first obvious option in response to the Shrapnel–Costa theorem is to accept that any rescue effort for Einstein–Bell realism is now determinately impossible. If one wanted, one could still maintain that the wavefunction be interpreted as  $\psi$ -epistemic, so long as there were no classical ontology underlying the wavefunction description (as in some versions of QBism (Fuchs *et al.*, 2014)). Alternatively, one could maintain a realist attitude towards the quantum formalism by adopting a  $\psi$ -ontic interpretation, with or without additional ontic degrees of

freedom. Perhaps the clearest avenue here is a  $\psi$ -ontic wavefunction without additional ontic structure (such as the many-worlds interpretation (Wallace, 2012)), as one must keep in mind that any additional ontic degrees of freedom would need to be nonlocal and/or contextual.

However, is there a way to rescue Einstein–Bell realism? I see only two remaining options. The first is to undermine the Shrapnel–Costa theorem, possibly by rejecting one or more underlying assumptions. It is a difficult task to identify the most appropriate foundations to challenge, but perhaps rethinking the nature of causality, inference, and/or probability might be a fruitful place to begin searching.<sup>8</sup>

Schmid *et al.* (2020) propose a generalisation of causality and inference—what they call a *causal-inferential theory*—with the express intention of avoiding the consequences of the traditional no-go theorems that rule out local or noncontextual realist approaches. They make the suggestion that previous such generalisations that can be employed to provide diagnoses of the true consequences of the no-go theorems “scramble” the ontological and epistemological aspects of this problem. The previous generalisations they mention include the ontological models framework from above (Spekkens, 2005), quantum generalisations of propositional logic (‘quantum logic’ (Jauch and Piron, 1963; Gudder, 1969; Mittelstaedt, 1976)), operational probability theories (which originate from the Pavia group (Chiribella *et al.*, 2010)), generalised probability theories (beyond quantum logic (Hardy, 2001)), and a nonclassical generalisation of Bayesian inference (Leifer and Spekkens, 2013). What is significant about this proposal is the way that Schmid *et al.* are explicitly attempting to “salvage” the notions of locality and noncontextuality, and so salvage, at least in part, a classical ontology. A rough outline of their main conjecture is as follows.

Recall that in Section 2.2 we were introduced to the term *Leibnizianity* that Schmid *et al.* employ to characterise generalised noncontextuality. Their formalisation simply captures the proclamation that there should be no ontological difference without operational difference. They claim that the traditional noncontextuality no-go theorems constrain logical space, as we noted above in Section 2.1, to either contextual realist approaches—and so to violations of Leibnizianity—or to approaches that forego some aspect of realism about the quantum state. The key claim of Schmid *et al.* is that they identify that it is indeed possible to have a Leibnizian realist interpretation of quantum theory, but only so long as the causal and inferential components are inherently nonclassical. While they do not provide any great amount of detail, they note that quantum causal modelling (Leifer and Spekkens, 2013; Cavalcanti and Lal, 2014; Costa and Shrapnel, 2016; Allen *et al.*, 2017) is likely to play a significant role in the development of any such realist interpretation, as well as more abstract characterisations of probability theory (Fritz, 2020). While the sentiment of this project is laudable, I make merely two cautionary comments in passing. As a rejoinder to the involvement of quantum causal modelling, it seems as though the Shrapnel–Costa theorem, which itself is built on the principles of quantum causal modelling, will make it difficult for any such generalisation on its own to salvage noncontextuality. On the involvement of abstract characterisations of

---

<sup>8</sup>Hofer-Szabó (2020) considers the possibility of giving up on Spekkens’ definition of generalised noncontextuality, which also underpins the Shrapnel–Costa theorem, and simply permitting unproblematically the possibility of a quantum system responding differently to different measurements represented by the same operator.



probability theory, any generalisation of probability theory that salvages noncontextuality that is non-Kolmogorovian faces the contention—brought by Feintzeig and Fletcher (2017)—that it simply does not offer any clear advantage, in terms of, say, guiding rational action, over contextual theories with ordinary Kolmogorovian probability.

The second remaining option for rescuing Einstein–Bell realism is arguably the more interesting for proponents of causal symmetry. This option would entail biting the bullet on contextuality, and either (i) considering the contextuality merely apparent and finding some natural explanation for it in terms of noncontextual ontic structure and some epistemic constraint, or (ii) accepting that quantum theory is underpinned by contextual ontic structure and then accounting for how the world might conspire to render this contextuality operationally undetectable in a classical setting (or, perhaps even both (i) and (ii)). The former, of course, is the method by which causally symmetric approaches meet the challenge of the Kochen–Specker theorem and the ontological models framework. The natural explanation there is that an observer’s ordinary temporal orientation constrains epistemic access to the future causes that influence the properties of some quantum system’s beables, and so make these ultimately noncontextual hidden variables seem contextual from the observer’s perspective. The ideal response for proponents of causal symmetry in the current scenario would be to identify a corresponding epistemic, or otherwise, constraint that shows how apparent contextuality can arise from actual noncontextuality. While I do not have a definitive argument to demonstrate the plausibility of such an account, I would like to make a speculative suggestion for the sort of constraint and corresponding consequences that might be required for this task.

However, before I do, it is worth pushing the dialectic of this work towards the following inconvenient conclusion. If the primary motivation for adopting a causally symmetric framework is to rescue Einstein–Bell realism, then we have just seen that the Shrapnel–Costa theorem renders this task either impossible, or at best beholden to the possibility of some further account explaining how, say, apparent contextuality arises from some noncontextual footing. However, and importantly, even if such an account could be found, it still may not be enough to rescue Einstein–Bell realism. Whether it does or not hangs on how ‘natural’ the account is. As we saw in Section 2.2, one of the strengths of causally symmetric approaches that rescue Einstein–Bell realism from the traditional no-go theorems is that the ideology of causal symmetry is more economical than a rejection of classical ontology. However, it is difficult to see how any account that introduces potentially artificial constraints or complex mechanisms can be proposed without significantly reducing the ideological economy of causal symmetry, jeopardising the very grounds upon which one might consider the approach more virtuous (Friederich and Evans, 2019).

Thus, in response to the Shrapnel–Costa theorem, the outlook is particularly worrying for Einstein–Bell realism and classical ontology. Not only is the obvious consequence of the theorem a straightforward rejection of Einstein–Bell realism, but even where there looks to be a possibility of accounting for contextual ontic variables within a causally symmetric framework, the cost of such an account could result in a dramatic decrease in the ideological economy of causal symmetry, and so render the rejection of classical ontology favourable on grounds of scientific virtue. Either way, it looks very difficult to maintain a classical ontology

for quantum theory.

Despite this pessimistic prognosis, I will finish this part of the dialectic with a speculative suggestion of where I see the most promising kind of possibilities for Einstein–Bell realism in the face the Shrapnel–Costa theorem. In Evans *et al.* (2016) we considered some peculiar problems faced by what we call the Price–Wharton approach to causally symmetric local hidden variables in the context of a universal  $\psi$ -epistemic wavefunction (which we refer to as  $\Psi$ -epistemic quantum cosmology). The problem is the following. The Price–Wharton approach poses a two-time boundary value problem for the kinematics of the underlying hidden variables. Additionally, as a result of the traditional no-go theorems, initial data on a Cauchy surface forming a well-posed Cauchy problem and comprised of local, noncontextual classical variables cannot possibly recover the statistical predictions of quantum theory. This is despite the fact that the initial data about the quantum state on the same Cauchy surface perfectly well comprise a well-posed Cauchy problem, due to the parabolicity of the Schrödinger equation. That the ontic state variables on the initial Cauchy surface do not comprise a well-posed Cauchy problem should be expected, though, as a two-time boundary problem requires additional information from the final boundary to obtain a complete determination between the boundaries. So sufficient initial data on a Cauchy surface to form a well-posed Cauchy problem would lead to an overdetermination of the ontic state at the final boundary. One way around this that we consider is that the laws governing the ontic state variables cannot be parabolic or hyperbolic PDEs.

However, there is an alternative, more speculative possibility that we briefly consider in this context, albeit one that is set to one side in Evans *et al.* (2016). The ontic state variables do solve a Cauchy problem, but are also determined as a two-time boundary problem (so as to circumvent the no-go theorems). The tension generated by this overdetermination is addressed by a constraint on the free action of an agent controlling the nature of the data on the initial and final boundary. We put it as follows (Evans *et al.*, 2016, p.7):

the tension that would need to be resolved is between: (i) the solution of a Cauchy problem from freely, arbitrarily and (ideally) completely specifiable initial data; and (ii) the symmetric expectation that the final boundary be equally freely, arbitrarily and completely specifiable. One way to escape this tension would be to remove the freedom to completely specify data on the final boundary: an agent controlling the final boundary would just happen to ‘choose’ a measurement that concords with the deterministic evolution of the ontic state. However, this would break the symmetry between the final and initial boundaries and would also remove the element of control that renders the Price–Wharton picture causally symmetric. To retain the symmetry would thus require some as-yet-unspecified principled constraint that limits an agent’s ability to freely, arbitrarily and completely specify both initial and final boundary data. This constraint must be such that the aspects of the ontic state on the *initial* Cauchy surface that are a consequence of the choices specified at the *final* boundary are not epistemically accessible before the final boundary is specified—and vice versa.

The significance of this possibility for the dialectic around the Shrapnel–Costa theorem is the following. Consider the nature of this proposed constraint on an agent’s ability to control the boundary data. One way for this constraint to work would be to limit agential control over the complete specifiability of data on both the initial and final boundary (in the same way as one can perform quantum measurements, but cannot control the precise outcome). Another

way for this constraint to work would be to limit, say, the time intervals at which an agent can specify data on initial and final boundaries. Without such a temporal constraint, we should expect the data on the initial and final boundaries in a two-time boundary value account to be generally related by elliptic PDEs. However, one could imagine, albeit in a highly speculative fashion, that if the distance between the boundaries were constrained to certain discrete time intervals—perhaps frequent enough to be practically undetectable to the agent—the initial and final ontic state variables could be related consistently by some hyperbolic PDE.<sup>9</sup> This would apparently resemble a well-posed Cauchy problem over the ontic state variables on the initial boundary, but those variables would be at least partly determined by the ontic state on the final boundary. There would be overdetermination in such a case, but the overdetermination would be fine-tuned so as not to create a contradiction. An agent would then be similarly constrained to intervene on a system at specific times on pain of generating a contradiction in the hidden variables that describe the kinematical properties of some quantum system between the boundaries.

This speculative example is interesting because of its potential realisation of process contextuality. According to process contextuality, the parts of the physical system that are contextual could include ontic states, global properties, causal mechanisms, dynamical laws, and boundary conditions. This implies that these features of the ontic process might depend upon the context of measurement. This seems to be a feature of the above example: the ontic state variables on the boundaries, and whether they are modelled by, say, elliptical or hyperbolic dynamical laws depends upon, say, the precise time intervals across which the agent is able to intervene on the system. In addition, this account looks highly fine-tuned, which also appears to be a feature of process contextuality. Could this then be a case of process contextuality?

Moreover, the imagined (and admittedly unspecified) constraint on the agent that limits the agent’s capability “to freely, arbitrarily, and completely specify both initial and final boundary data” to specific time intervals should ultimately, with the right detail, render such a causally symmetric local hidden variable approach process noncontextual, as the constraint removes the above-mentioned contextuality (but admittedly for the price of fine-tuning). This, by my lights, appears to qualify as biting the bullet on contextuality as per the Shrapnel–Costa theorem, but also as having such process contextuality be merely apparent. This is because the apparent contextuality would simply be a feature of our continuous-time model of the causally symmetric hidden variables, which themselves are determined on a discrete-time basis due to the constraint. Provided that the imagined constraint could be given a ‘natural’ explanation (which I have certainly not attempted here), the process contextuality could be accounted for in terms of noncontextual ontic structure and some agent-centred constraint. However, there must be considerable doubt concerning whether such a constraint could provide a natural explanation, at the very least because it explicitly introduces fine-tuning into the explanation

---

<sup>9</sup>One could think of this as roughly analogous to turning a ‘thick’ sandwich problem into a ‘thin’ sandwich problem in geometrodynamics (Baierlein *et al.*, 1962; Wheeler, 1964; Misner *et al.*, 1973). This problem in geometrodynamics is motivated by the corresponding problem in electrodynamics, where the free specification of data on two boundaries with periodic boundary conditions can be made consistent by stipulating that the interval of time between the two boundaries is an integer multiple of a half period of the harmonic oscillation of the field (Wheeler, 1964, p.355).

of how apparent contextuality arises from a noncontextual ontic state. So long as one thinks that a Leibnizian principle of parsimony is a good guide to scientific methodology, fine-tuning of this sort is not a virtue of a theory—and it is moreover worth noting that, due to the nonextendibility result above, the addition of noncontextual ontic structure cannot provide more information than is contained in the quantum description.

I do not mention this speculative example here to argue in its favour, by any means. My purpose here is to provide a demonstration of the sort of argumentation that *would* be required as part of what is arguably the only reasonable path forward for the Einstein–Bell realist. What is required is an account of, say, apparent process contextuality that has these sorts of features. This moreover demonstrates the interrelation between the naturalness of the constraint and the ideological economy of the model: unless the constraint makes use of uncontroversial features of time, space, agency, and so on, it will face a hard task keeping the ideology economical.

## 5 Final Thoughts

Is this the end of a classical ontology for quantum mechanics? Yes, I think this has to be the end. Perhaps there is hope that locality and noncontextuality can be salvaged by an appropriate generalisation of probability theory and the characterisation of inference. However, the path ahead for the causal symmetry program looks less hopeful. While there is an open logical possibility that a causally symmetric hidden variable approach might be able to provide a natural explanation of how a noncontextual ontic state might appear to be contextual on account of some low-cost constraint, this seems at the present moment highly unlikely. What is more, the nonextendibility result implies that any noncontextual ontic state cannot contribute to the predictive power of quantum mechanics—any such contribution can only be made by additional instrument-contextual variables. So any such logical possibility is unlikely to render accepting causal symmetry more economical than rejecting a classical ontology, and will necessarily fail to contribute to increased predictive power for quantum mechanics. In so far as this unlikely logical possibility is the last refuge for Einstein–Bell realism, it looks like we should give up on Einstein–Bell realism and, with it, classical ontology.

**Acknowledgments:** This work is based on a talk I gave at the BPS conference in 2018. Some of this talk, including parts of the logical structure and the broad recommendation, is dispersed amongst the detail of Friederich and Evans (2019). For useful discussion and comments I would like to thank Fabio Costa, Simon Friederich, Sean Gryb, Sally Shrapnel, Karim Thébault, Ken Wharton, and the audience at the BPS 2018 conference held at the University of Oxford. This work is supported by the University of Queensland and the Australian Government through the Australian Research Council (DE170100808).

## References

- Accardi, L. (1976). Nonrelativistic Quantum Mechanics as a Noncommutative Markov Process. *Advances in Mathematics* **20**(3): 329–366. doi:10.1016/0001-8708(76)90201-2.
- Adlam, E. (2020). The Operational Choi–Jamiołkowski Isomorphism. *Entropy* **22**(9): 1063. doi:10.3390/e22091063.

- Aharonov, Y., Bergmann, P. G. and Lebowitz, J. L. (1964). Time Symmetry in the Quantum Process of Measurement. *Physical Review* **134**: B1410–B1416. doi:10.1103/PhysRev.134.B1410.
- Allen, J.-M. A., Barrett, J., Horsman, D. C., Lee, C. M. and Spekkens, R. W. (2017). Quantum Common Causes and Quantum Causal Models. *Physical Review X* **7**: 031021. doi:10.1103/PhysRevX.7.031021.
- Baierlein, R. F., Sharp, D. H. and Wheeler, J. A. (1962). Three-Dimensional Geometry as Carrier of Information about Time. *Physical Review* **126**: 1864–1865. doi:10.1103/PhysRev.126.1864.
- Bell, J. S. (1966). On the Problem of Hidden Variables in Quantum Mechanics. *Reviews of Modern Physics* **38**: 447–452. doi:10.1103/RevModPhys.38.447.
- (1976). The Theory of Local Beables. *Epistemological Letters* **9**: 11–24.
- (1981). Bertlemann’s socks and the nature of reality. *Journal de Physique Colloques* **42**(C2): 41–62. doi:10.1051/jphyscol:1981202.
- (1990). La nouvelle cuisine. In A. Sarlemijn and P. Kroes (eds.), *Between Science and Technology*, Elsevier, Amsterdam, pp. 97–115. doi:10.1016/B978-0-444-88659-0.50010-7.
- Bohm, D. (1952). A Suggested Interpretation of the Quantum Theory in Terms of “Hidden” Variables. I. *Physical Review* **85**(2): 166–179. doi:10.1103/PhysRev.85.166.
- Bohr, N. (1935). Can Quantum-Mechanical Description of Physical Reality be Considered Complete? *Physical Review* **48**: 696–702. doi:10.1103/PhysRev.48.696.
- Cavalcanti, E. G. (2018). Classical Causal Models for Bell and Kochen-Specker Inequality Violations Require Fine-Tuning. *Physical Review X* **8**: 021018. doi:10.1103/PhysRevX.8.021018.
- Cavalcanti, E. G. and Lal, R. (2014). On modifications of Reichenbach’s principle of common cause in light of Bell’s theorem. *Journal of Physics A: Mathematical and Theoretical* **47**(42): 424018. doi:10.1088/1751-8113/47/42/424018.
- Chen, Z. and Montina, A. (2011). Measurement contextuality is implied by macroscopic realism. *Physical Review A* **83**: 042110. doi:10.1103/PhysRevA.83.042110.
- Chiribella, G. (2012). Perfect discrimination of no-signalling channels via quantum superposition of causal structures. *Physical Review A* **86**: 040301. doi:10.1103/PhysRevA.86.040301.
- Chiribella, G., D’Ariano, G. M. and Perinotti, P. (2008a). Quantum Circuit Architecture. *Physical Review Letters* **101**: 060401. doi:10.1103/PhysRevLett.101.060401.
- (2008b). Transforming quantum operations: Quantum supermaps. *Europhysics Letters* **83**(3): 30004. doi:10.1209/0295-5075/83/30004.
- (2009). Theoretical framework for quantum networks. *Physical Review A* **80**: 022339. doi:10.1103/PhysRevA.80.022339.
- (2010). Probabilistic theories with purification. *Physical Review A* **81**: 062348. doi:10.1103/PhysRevA.81.062348.

- Chiribella, G., D’Ariano, G. M., Perinotti, P. and Valiron, B. (2013). Quantum computations without definite causal structure. *Physical Review A* **88**: 022318. doi:10.1103/PhysRevA.88.022318.
- Choi, M.-D. (1975). Completely positive linear maps on complex matrices. *Linear Algebra and its Applications* **10**(3): 285–290. doi:10.1016/0024-3795(75)90075-0.
- Colnaghi, T., D’Ariano, G. M., Facchini, S. and Perinotti, P. (2012). Quantum computation with programmable connections between gates. *Physics Letters A* **376**(45): 2940–2943. doi:10.1016/j.physleta.2012.08.028.
- Costa, F. and Shrapnel, S. (2016). Quantum causal modelling. *New Journal of Physics* **18**(6): 063032. doi:10.1088/1367-2630/18/6/063032.
- Einstein, A., Podolsky, B. and Rosen, N. (1935). Can Quantum-Mechanical Description of Physical Reality Be Considered Complete? *Physical Review* **47**: 777–780. doi:10.1103/PhysRev.47.777.
- Evans, P. W. (2015). Retrocausality at no extra cost. *Synthese* **192**(4): 1139–1155. doi:10.1007/s11229-014-0605-0.
- Evans, P. W., Gryb, S. and Thébault, K. P. (2016).  $\Psi$ -epistemic quantum cosmology? *Studies in History and Philosophy of Modern Physics* **56**: 1–12. doi:10.1016/j.shpsb.2016.10.005.
- Feintzeig, B. H. and Fletcher, S. C. (2017). On Noncontextual, Non-Kolmogorovian Hidden Variable Theories. *Foundations of Physics* **47**: 294–315. doi:10.1007/s10701-017-0061-z.
- Friederich, S. and Evans, P. W. (2019). Retrocausality in Quantum Mechanics. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Stanford University.
- Fritz, T. (2020). A synthetic approach to Markov kernels, conditional independence and theorems on sufficient statistics. *Advances in Mathematics* **370**: 107239. doi:10.1016/j.aim.2020.107239.
- Fuchs, C. A., Mermin, N. D. and Schack, R. (2014). An introduction to QBism with an application to the locality of quantum mechanics. *American Journal of Physics* **82**(8): 749–754. doi:10.1119/1.4874855.
- Gudder, S. (1969). On the quantum logic approach to quantum mechanics. *Communications in Mathematical Physics* **12**: 1–15. doi:10.1007/BF01646431.
- Gutoski, G. and Watrous, J. (2007). Toward a General Theory of Quantum Games. In *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing*. Association for Computing Machinery, New York, NY, USA, p. 565–574. doi:10.1145/1250790.1250873.
- Hardy, L. (2001). Quantum Theory From Five Reasonable Axioms. arXiv:quant-ph/0101012.
- (2005). Probability Theories with Dynamic Causal Structure: A New Framework for Quantum Gravity. arXiv:gr-qc/0509120.
- (2012). The operator tensor formulation of quantum theory. *Philosophical Transactions of the Royal Society A* **370**: 3385–3417. doi:10.1098/rsta.2011.0326.
- Harrigan, N. and Rudolph, T. (2007). Ontological models and the interpretation of contextuality. arXiv:0709.4266 [quant-ph].

- Harrigan, N. and Spekkens, R. W. (2010). Einstein, Incompleteness, and the Epistemic View of Quantum States. *Foundations of Physics* **40**(2): 125–157. doi:10.1007/s10701-009-9347-0.
- Hofer-Szabó, G. (2020). Commutativity, comeasurability, and contextuality in the Kochen-Specker arguments. arXiv:2004.14053 [quant-ph].
- Jamiołkowski, A. (1972). Linear transformations which preserve trace and positive semidefiniteness of operators. *Reports on Mathematical Physics* **3**(4): 275–278. doi:10.1016/0034-4877(72)90011-0.
- Jauch, J. M. and Piron, C. (1963). Can hidden variables be excluded in quantum mechanics? *Helvetica Physica Acta* **36**: 827–837.
- Kochen, S. and Specker, E. P. (1967). The Problem of Hidden Variables in Quantum Mechanics. *Journal of Mathematics and Mechanics* **17**: 59–87.
- Kok, P., Munro, W. J., Nemoto, K., Ralph, T. C., Dowling, J. P. and Milburn, G. J. (2007). Linear optical quantum computing with photonic qubits. *Reviews of Modern Physics* **79**: 135–174. doi:10.1103/RevModPhys.79.135.
- Kretschmann, D. and Werner, R. F. (2005). Quantum channels with memory. *Physical Review A* **72**: 062323. doi:10.1103/PhysRevA.72.062323.
- Leifer, M. S. (2014). Is the Quantum State Real? An Extended Review of  $\Psi$ -ontology Theorems. *Quanta* **3**(1): 67–155. doi:10.12743/quanta.v3i1.22.
- Leifer, M. S. and Pusey, M. F. (2017). Is a time symmetric interpretation of quantum theory possible without retrocausality? *Proceedings of the Royal Society of London A* **473**(2202). doi:10.1098/rspa.2016.0607.
- Leifer, M. S. and Spekkens, R. W. (2013). Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference. *Physical Review A* **88**(5): 052130. doi:10.1103/PhysRevA.88.052130.
- Lindblad, G. (1979). Non-Markovian Quantum Stochastic Processes and Their Entropy. *Communications in Mathematical Physics* **65**(3): 281–294. doi:10.1007/BF01197883.
- Luchnikov, I. A., Vintskevich, S. V. and Filippov, S. N. (2018). Dimension truncation for open quantum systems in terms of tensor networks. arXiv:1801.07418 [quant-ph].
- Misner, C. W., Thorne, K. S. and Wheeler, J. A. (1973). *Gravitation*. H. W. Freeman and Company, San Francisco.
- Mittelstaedt, P. (1976). *Quantum Logic*, D. Reidel Publishing Company, Dordrecht, pp. 501–514. doi:10.1007/978-94-010-1449-6\_28.
- Modi, K. (2012). Operational approach to open dynamics and quantifying initial correlations. *Scientific Reports* **2**(1): 581. doi:10.1038/srep00581.
- Mohseni, M. and Lidar, D. A. (2007). Direct characterization of quantum dynamics: General theory. *Physical Review A* **75**: 062331. doi:10.1103/PhysRevA.75.062331.
- Norsen, T. (2011). John S. Bell’s concept of local causality. *American Journal of Physics* **79**(12): 1261–1275. doi:10.1119/1.3630940.

- O’Brien, J. L., Pryde, G. J., Gilchrist, A., James, D. F. V., Langford, N. K., Ralph, T. C. and White, A. G. (2004). Quantum Process Tomography of a Controlled-NOT Gate. *Physical Review Letters* **93**: 080502. doi:10.1103/PhysRevLett.93.080502.
- Oeckl, R. (2003). A “general boundary” formulation for quantum mechanics and quantum gravity. *Physics Letters B* **575**(3): 318–324. doi:10.1016/j.physletb.2003.08.043.
- Oreshkov, O., Costa, F. and Brukner, v. (2012). Quantum correlations with no causal order. *Nature Communications* **3**: 1092. doi:10.1038/ncomms2076.
- Pusey, M. F., Barrett, J. and Rudolph, T. (2012). On the reality of the quantum state. *Nature Physics* **8**(6): 475–478. doi:10.1038/nphys2309.
- Quine, W. V. O. (1951). Ontology and Ideology. *Philosophical Studies* **2**: 11–15. doi:10.1007/BF02198233.
- Riebe, M., Chwalla, M., Benhelm, J., Häffner, H., Hänsel, W., Roos, C. F. and Blatt, R. (2007). Quantum teleportation with atoms: quantum process tomography. *New Journal of Physics* **9**(7): 211–211. doi:10.1088/1367-2630/9/7/211.
- Riebe, M., Kim, K., Schindler, P., Monz, T., Schmidt, P. O., Körber, T. K., Hänsel, W., Häffner, H., Roos, C. F. and Blatt, R. (2006). Process Tomography of Ion Trap Quantum Gates. *Physical Review Letters* **97**: 220407. doi:10.1103/PhysRevLett.97.220407.
- Ringbauer, M. (2017). *Exploring Quantum Foundations with Single Photons*. Springer Theses. Springer International Publishing, Cham.
- Schmid, D., Selby, J. H. and Spekkens, R. W. (2020). Unscrambling the omelette of causation and inference: The framework of causal-inferential theories. arXiv:2009.03297 [quant-ph].
- Shrapnel, S. and Costa, F. (2018). Causation does not explain contextuality. *Quantum* **2**: 63. doi:10.22331/q-2018-05-18-63.
- Spekkens, R. W. (2005). Contextuality for preparations, transformations, and unsharp measurements. *Physical Review A* **71**: 052108. doi:10.1103/PhysRevA.71.052108.
- (2019). The ontological identity of empirical indiscernibles: Leibniz’s methodological principle and its significance in the work of Einstein. arXiv:1909.04628 [physics.hist-ph].
- von Neumann, J. (1932). *Mathematische Grundlagen der Quantenmechanik*. Springer, Berlin.
- Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford University Press, Oxford.
- Wharton, K. B. (2010). Time-Symmetric Boundary Conditions and Quantum Foundations. *Symmetry* **2**(1): 272–283. doi:10.3390/sym2010272.
- (2014). Quantum states as ordinary information. *Information* **5**(1): 190–208. doi:10.3390/info5010190.
- Wharton, K. B. and Argaman, N. (2020). Colloquium: Bell’s theorem and locally mediated reformulations of quantum mechanics. *Reviews of Modern Physics* **92**: 021002. doi:10.1103/RevModPhys.92.021002.
- Wheeler, J. A. (1964). *Geometrodynamics and the Issue of the Final State*, Gordon and Breach Science Publishers, New York, pp. 315–520.



Wiseman, H. M. and Cavalcanti, E. G. (2017). *Causarum Investigatio and the Two Bell's Theorems of John Bell*, Springer International Publishing, Cham, pp. 119–142. doi:10.1007/978-3-319-38987-5\_6.

Wood, C. J. (2009). *Non-Completely Positive Maps: Properties and Applications*. Ph.D. thesis, Macquarie University, Sydney, New South Wales.