Guiding our interpretation of quantum theory by principles of causation and inference

by

David Schmid

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Physics and Astronomy
in
Quantum Information

Waterloo, Ontario, Canada, 2021

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner: Stephen Bartlett
Professor, School of Physics, University of Sydney

Supervisor(s): Robert W. Spekkens
Senior faculty, Perimeter Institute for Theoretical Physics
Adjunct faculty, Dept. of Physics and Astronomy, University of Waterloo
Kevin Resch
Professor, Dept. of Physics and Astronomy

Internal Member: Robert Mann
Professor, Dept. of Physics and Astronomy, University of Waterloo

Internal-External Member: John Watrous
Professor, School of Computer Science, University of Waterloo

Other Member(s): Lucien Hardy
Senior Faculty, Perimeter Institute for Theoretical Physics
Adjunct faculty, Dept. of Physics and Astronomy, University of Waterloo

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

The work in this thesis was done in collaboration with various colleagues. Here is my rough attempt to attribute credit to my primary collaborators on each project in the thesis.

Chapter 2: The idea for this work was formulated by Rob Spekkens. Most of the results were proven by myself, with the guidance of Rob Spekkens. The writing is my own, with much feedback from Rob Spekkens.

Chapter 3: The initial ideas for this work were due to Rob Spekkens and myself, while the proof of the main result was largely due to Matthew Pusey and to previous work [260] done by myself, Matt Pusey, John Selby, and Rob Spekkens. The bulk of the writing is my own.

Chapter 4: The conceptual ideas reported in this excerpt of Ref. [322] were developed equally by Elie Wolfe, myself, Ana Belén Sainz, Ravi Kunjwal, and Rob Spekkens, with the initial idea originating with Rob Spekkens.

Chapter 5: The main idea of this work was my own, as a conceptual extension of related previous work [322] done by myself, Elie Wolfe, Ana Belén Sainz, Ravi Kunjwal, and Rob Spekkens. The results were developed by myself with guidance from Denis Rosset. The bulk of the writing is my own.

Chapter 6: The ideas, work, and actual writing that went into this project were contributed evenly by myself, John Selby, and Rob Spekkens. However, the deepest conceptual roots of the project can already be found throughout Rob Spekkens's work over the last decade or two.

Abstract

A key aim of quantum foundations is to characterize the sense in which nature goes beyond classical physics. Understanding nonclassicality is one of our best avenues towards finding a satisfactory interpretation of quantum theory. By determining which classical principles cannot be satisfied in any empirically adequate physical theory, we begin to see which principles *can* be preserved, which in turn gives us insight into the ontology of the world. These insights then guide us in determining which questions to ask and which experiments to perform next. Furthermore, it is these nonclassical aspects of nature that give rise to new technologies such as the speed-ups of quantum computation or the security of quantum key distribution.

The gold standard for establishing that a phenomenon is truly nonclassical is to prove that it violates the principle of local causality or the principle of noncontextuality. Much of this thesis reports on my research relating to these two principles. This research primarily involves (i) finding new justifications for our notions of nonclassicality; (ii) refining their fundamental definitions; (iii) quantifying and characterizing their various manifestations; and (iv) finding applications where nonclassical phenomena act as resources for information processing.

Ultimately, all of this work on nonclassicality is woven together into a novel framework for physical theories introduced by myself, John Selby, and Rob Spekkens. Its main advantage over preexisting frameworks is that it maintains a clear distinction between which elements of a given physical theory directly describe causal processes, and which refer only to one's inferences about causal processes. This clarifies a number of confusions in the literature which arose precisely because previous frameworks scrambled causal and inferential concepts. Furthermore, local causality and noncontextuality emerge in this framework as the assumptions that the causal and inferential structures (respectively) that are operationally observed must be respected in the underlying ontology. This work constitutes a first step in developing a new interpretation of quantum theory—the first interpretation designed to satisfy the spirit of both local causality and noncontextuality.

Acknowledgements

First, I extend my deepest thanks to Rob Spekkens. Working with Rob has been the most fulfilling experience I could ever imagine. Basking in the clarity, depth, and honesty of his thinking has been an absolute privilege, and his empathy and generosity made my time at PI a great pleasure. It wasn't until I read Rob's papers that I had any hope that we would ever understand quantum theory, and it wasn't until I started working with him that I started enjoying the day-to-day process of being a physicist.

Even when the daily process of doing research felt like banging my inadequate head against a brick wall, I was encouraged along my path by inspiring mentors. I owe a deep thanks to Paul Kwiat for his support and encouragement, and for believing in me before it made any sense to do so. The entire Kwiat group will forever be dear to my heart. I thank Aephraim Steinberg and Daniel James for selflessly encouraging and enabling me to chase my foundational dreams, even at the cost of shortening our enjoyable time together. I thank also the graduate students in Toronto for our good times together. I especially thank Nico and Kevin for convincing me that there is no prerequisite to being a theorist—that I just needed to sit down and start working.

Many friends impacted me during my time at Perimeter. Belén, Ravi, and Elie introduced me to PI and to quantum foundations, making me feel at home while sharing their wisdom and expertise. I thank all my PI friends, including John, Denis, Tom, Jamie, Tomáš, TC, Alvaro, and Nitica, for the ideas, the rock climbing, and the company. Special thanks to John Selby for voluntarily being trapped into working on exciting ideas (even when the pubs were already open). Special thanks to Aggie for her support and the great times we had teaching the quantum foundations PSI course together. Thanks also to Matt L., Matt P., Matty, Robin, Lidia, Lorenzo, and Josiah for good times and deep insights.

In addition to those already listed, I've benefited from the insights and encouragement of too many great researchers to mention. Among them, I thank especially Jon Barrett, Stephen Bartlett, Francesco Buscemi, Eric Cavalcanti, Giulio Chiribella, Fabio Costa, Lucien Hardy, Paolo Perinotti, Valerio Scarani, and Sally Shrapnel. I thank also the members of my committee for the time and energy spent on my behalf. I thank also the institutions that supported me, especially everyone at PI, and am deeply grateful for the generosity of the fellowships that kept me out of the fields for 10 years now.

Thanks to Dan, Jan, Yulia, Joel, Eric, Catie, Mae, and Linda. Each of you had a lasting impact on me that I cherish. Thanks to Haoxing (and Potato) for the great times, even when the times weren't great. To Mara—your shining example has been a central pillar in my life. Deepest thanks for your love and support. To Viola—you always keep everything colorful. Here's to a life covered in paint.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

Which features of quantum theory are impossible to explain classically?

The history of quantum foundations is littered with attempted answers to this question. Many of the attempted answers involve pointing out specific counterintuitive quantum phenomena that lack any obvious explanation in terms of, say, Newtonian mechanics or electrodynamics. Common suspects are, e.g., interference, entanglement, and no-cloning. For every argument of this sort, however, one can find a counterargument wherein a sufficiently creative researcher devises a relatively simple and classical explanation for the phenomena which were claimed to be mysterious. A more systematic approach to answering this question is to formulate a set of assumptions that are taken to encode one's notion of 'classicality', and then to derive from this a contradiction with the predictions of quantum theory. By now there are scores of such no-go theorems, and correspondingly many notions of classicality.

Of these, two no-go theorems stand far ahead of the pack, in terms of the compellingness of their assumptions, their breadth of applicability, and their utility in guiding us in developing new ideas and technologies. The first of these is Bell's proof that there is no local hidden variable model that can reproduce the predictions of quantum theory. The second, attributed to Kochen and Specker, establishes that there is no noncontextual hidden variable model which can reproduce the predictions of quantum theory. The precise characterization of the foundational assumptions and principles underlying these no-go theorems is an ongoing topic of study. For example, the set of assumptions required for Bell's no-go theorem has since been clarified and greatly extended in scope by the framework of classical causal modeling. Additionally, Kochen-Specker's notion of noncontextuality was significantly generalized and better motivated by Spekkens. Much of the work reported in this thesis also contributes to these clarifications and generalizations.

In order to convincingly argue that a phenomena is nonclassical, then, one should demonstrate that it cannot be explained by any local and noncontextual model. The natural next step is to better understand the range of scenarios in which quantum theory fails to respect them. Although the first proofs of nonlocality and contextuality involved fairly specific constructions without any practical relevance, it is by now standard to seek proofs of nonclassicality that exhibit connections to physical situations of independent interest. That such connections are regularly found is no coincidence, but rather is further evidence that these are good notions of nonclassicality. After all, it would be hard to convince ourselves that we had characterized the essence of quantum theory, if this essence had no practical ramifications for, say, information processing, computation, or building gadgets. In recent years, this recognition that nonclassicality can be viewed as a resource for a variety of practical tasks has led to a concerted effort to quantify and characterize its various manifestations.

The study of nonclassicality is of intrinsic foundational interest, and we have just noted that it has practical utility as well. There is a third motivation, however, which to me personally is the greatest of all. Understanding what makes quantum theory special (and in particular, nonclassical) provides insights that guide us in interpreting quantum theory. It is worth noting that the task of interpreting quantum theory is *not* merely of foundational interest. It is critical, for example, to the task of extending quantum theory to new domains, be those domains as fundamental as spacetime physics, as applied as causal inference, or as complex as artificial intelligence. The intuitive stories we build around the formal elements of a theory guide our research agendas, determining what questions get asked, what new experiments are carried out, and what new devices get built.

As an example of how our understanding of nonclassicality guides us in interpreting quantum theory, consider how researchers typically evaluate the various existing quantum interpretations. One tabulates the various foundational principles (e.g. of classicality or simplicity) that one wishes to uphold, and then determines which are respected or violated by each interpretation. For example, Bohmian mechanics preserves our intuitions of a deterministic world wherein particles have definite trajectories, but it violates both locality and noncontextuality. Many Worlds aims to preserve locality, but at the cost of an extravagant ontology and of giving up the idea that a measurement yields a single outcome. QBism also aims to preserve locality, but does so by giving up the idea that measurement outcomes are events independent of the observer who experienced them. Then, based on our judgement of which principles are most sacred, we determine which interpretations we prefer over which others.

Perhaps the most common view among researchers in quantum foundations today is that none of these interpretations is satisfactory. Furthermore, no interpretation has achieved even a sizeable minority consensus. As such, there is a real need for new interpretations. In my view, the reason existing interpretations are unappealing is because each abandons at least one or two of our foundational physical principles *wholesale*, in order to preserve the others. No matter how successful the interpretation is at resolving particular quantum mysteries, then, one is left with the fear that it is only for lack of creativity that we have given up on the abandoned principles. The only path I see to a broadly appealing interpretation is to (i) identify the most significant principles that appear to be violated by quantum theory, and then (ii) devise an interpretation that manages to preserve the spirit of *all* of these principles, by tweaking the precise definitions of the basic physical concepts used to define them. These modifications, of course, must be minimal enough so that the physical concepts do not lose their meaning (and so that the principles do not lose their appeal), but significant enough that one can avoid no-go theorems and actually reproduce the predictions of quantum theory.

This task may seem vague or difficult, and we certainly have not yet succeeded at it. However, I will by the end of this thesis present an explicit framework which formalizes these ideas and begins making progress on them. The key foundational principles that we aim to preserve using this framework are motivated directly by the Bell and Kochen-Specker no-go theorems, together with our refinements to the assumptions going into them. (These refinements are discussed in the following subsections; some of them are contributions of the work reported in this thesis.) In this framework, developed jointly with John Selby and Rob Spekkens, we consider modifications to the classical theory of causation and to the classical theory of inference, and we show how a synthesis of these revised theories has the potential to circumvent nonlocality and contextuality no-go theorems. Ultimately, we consider this to be a promising approach to reconstructing quantum theory in a manner that does the least possible violence to our most cherished physical principles.

1.0.1 How to read this thesis

This thesis consists of three parts. Part I consists of Chapters 2 and 3, which focus on the notion of *generalized noncontextuality*. Part II consists of Chapters 4 and 5, which focus on the notion of 'locality', which we argue is better understood as a notion of *causal compatibility*. Part III consists of Chapter 6, which presents a novel framework for realist theories, the framework of *causal-inferential theories*, which ties together all of my work on noncontextuality and nonlocality.

In the remainder of this introduction, I give a nontechnical introduction to the relevant two notions of classicality, with a focus on how the version of these concepts that I will consider are refined versions of the original notions introduced by Bell and Kochen and Specker. I also give a high-level summary of the contributions made by my research and reported in the following chapters. Finally, I summarize the basic aims of our framework for causal-inferential theories, as well as how this framework is motivated by and ties together my work on nonclassicality.

Formal definitions and complete details can be found in the respective chapters (but not in this introduction). The chapters in this thesis are taken almost verbatim from some of the papers published during my graduate studies.

1.0.2 Generalized noncontextuality

The term *noncontextuality* refers to a particular property that an ontological explanation of some given operational statistics may satisfy. Traditionally, a *noncontextual* hidden variable model is one wherein the fundamental (ontic) state of a system specifies what outcome will occur for any given measurement, independent of what other measurements are simultaneously carried out. This is the notion of Kochen-Specker noncontextuality [162]. As a simple example, imagine the two projective measurements $\{|0\rangle\langle 0|, |1\rangle\langle 1|, |2\rangle\langle 2|\}$ and $\{|+\rangle\langle +|, |-\rangle\langle -|, |2\rangle\langle 2|\}$, where $|\pm\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$. In a noncontextual model, every ontic state that dictates that the third outcome will occur in the first measurement must also dictate that the third outcome will occur in the second measurement, since the projector $|2\rangle\langle 2|$ corresponding to the two is the same. That is, whether or not a given outcome occurs cannot be a function of what specific measurement it is a part of.

Is Kochen-Specker noncontextuality even a plausible assumption, much less a sensible notion of classicality? As noted as far back as Bell [33], *"these different possibilities [that is, different measurements] require different experimental arrangements; there is no* a priori *reason to believe that the results... should be the same. The result of an observation may reasonably depend not only on the state of the system (including hidden variables), but also on the complete disposition of the apparatus"*. Nonetheless, many have tried to argue for the plausibility of Kochen-Specker's assumption of noncontextuality; see e.g. Ref. [205]. Perhaps the most convincing justification of the assumption was given by Spekkens in Ref. [281], which argued that noncontextuality is motivated by a principle of ontological economy: that a good theory should avoid introducing *'differences in the ontological explanations of empirical phenomena where there are no differences in the phenomena themselves'*.

More importantly, Ref. [281] argued that the same motivations that can be given for

Kochen-Specker noncontextuality can *also* be given for a more general notion which has come to be known as *generalized noncontextuality*. Essentially, an ontological representation of quantum theory is generalized-noncontextual when the representation of a given laboratory procedure depends only on the quantum description of that procedure (i.e. the density operator, channel, or positive-operator valued measure associated to it). In a noncontextual model, then, any two processes that lead to all the same predictions—no matter how distinct their physical implementation—must have representations that *also* lead one to make all the same predictions. Thus the indistinguishability of such distinct procedures is *explained* by the indistinguishability of their ontological behaviors.

The principle of generalized noncontextuality has an unparalleled pedigree as a notion of classicality. First, it can be motivated by an appeal to Leibniz's principle of the identity of indiscernibles [285], a principle for theory construction that underpins much of Einstein's work on special and general relativity. It has also been shown [260, 262, 283] that generalized noncontextuality is equivalent to the (independently motivated) notion of classicality that is unilaterally endorsed by the quantum optics community—namely, the existence of a positive Wigner representation. It is furthermore equivalent to a third (again independently motivated) notion of classicality that arises naturally within the framework of generalized probabilistic theories [23, 138] (GPTs)—namely, being a subtheory of a simplicial GPT [23] (that is, one whose extremal states are all perfectly distinguishable). Additionally, Ref. [241] shows that noncontextuality emerges in the classical limit considered within the Quantum Darwinian research program. Furthermore, other key indicators of nonclassicality, such as violations of local causality [33] or observations of anomalous weak values [228], have been proven to be instances of contextuality. Finally, contextuality has been proven to be a resource for information-processing [14, 60, 249, 286], computation [144, 234, 255], state discrimination [263], cloning [192], and metrology [191].

Part I of this thesis, namely Chapters 2 and 3, constitutes a collection of some of my research on generalized noncontextuality.

Chapter 2 presents Ref. [263], which demonstrates that contextuality is necessary for optimal quantum state discrimination. More precisely, for the minimum error state discrimination task that we consider, the optimal success rate achievable in quantum theory is higher than that achievable in any noncontextual model. We prove this result via an intuitive argument that has since been the basis for a proof that contextuality is also necessary for optimal cloning [192]. We also derive a noise-robust, experimentally testable noncontextuality inequality to witness this particular form of nonclassicality, using a more systematic argument which then formed the basis of the general algorithm discussed in the next chapter.

Chapter 3 presents Ref. [255], which proves that there is a unique noncontextual representation of every stabilizer subtheory in odd dimensions, and no noncontextual representation of any stabilizer subtheory in even dimensions. We prove that this unique noncontextual representation is equivalent to Gross's discrete Wigner representation [133], which is in turn equivalent to the Spekkens' toy model [282]. This constitutes a complete characterization of the (non)classicality of the most widely studied subtheory in quantum information. Leveraging the stabilizer subtheory's connection to universal quantum computation via state injection [43], we also prove that generalized contextuality is necessary for universal quantum computation in the state injection model, extending the analogous result of Ref. [154] for Kochen-Specker contextuality.

Five of my papers on generalized noncontextuality were not included in this thesis, to keep the length down. These can be found in Refs. [128, 260, 262, 264, 270].

Ref. [264] provides the first (and to date only known) systematic technique for deciding whether or not a set of data in a prepare-measure scenario admits of a noncontextual model. The set of noncontextual correlations is proven to always constitute a polytope, and we give an explicit method for deriving the set of all noncontextuality inequalities (defining the facets of this polytope) for any given prepare-and-measure scenario (with respect to any set of operational equivalences of interest) using linear programming.

Ref. [262] derives necessary and sufficient conditions on the geometry of a generalized probabilistic theory such that one's operational (prepare-and-measure) scenario admits of a noncontextual representation. In particular, we show that an operational theory admits of a noncontextual ontological model if and only if the associated GPT admits of *any* ontological model, or equivalently, if and only if the associated GPT embeds into a simplicial GPT. Since the standard notion of classicality in the GPT framework is simpliciality, this work is a strong piece of evidence that generalized noncontextuality is a good notion of classicality. It also suggests a second technique for deciding whether or not a set of data admits of a noncontextual model, although we have not yet worked out a generic algorithm to test for such simplex embeddings.

Ref. [260] proves that (under the assumption of tomographic locality [69, 138]) every generalized noncontextual representation has a very specific mathematical form; namely it is an exact frame representation [103], with precisely as many ontic states as the associated GPT dimension. We also extend many of the results of Ref. [262] from prepare-and-measure scenarios to arbitrary compositional scenarios. This is a critical step in strengthening the connections between contextuality and computation, since the latter is inherently about compositional circuits rather than simple prepare-and-measure scenarios. Indeed, the results of Ref. [260] were critical in proving the results of Ref. [255] (discussed in Chapter 3).

Ref. [270] establishes that one can find proofs of contextuality even in scenarios wherein all measurements are compatible—that is, can be simultaneously measured. This is in stark contrast to the well-known fact that Kochen-Specker contextuality may only be established in scenarios with incompatible measurements.

Ref. [128] proves that 'almost-quantum' correlations [209] are inconsistent with Specker's principle. This is one of only a few known principles that manages to distinguish quantum theory from any almost-quantum GPT [251].

1.0.3 Classicality of common-cause processes

Traditionally, Bell's theorem [33] is cast as a dilemma between abandoning local causality and abandoning realism. Proofs that Bell's assumptions cannot be satisfied are unilaterally considered strong evidence of nonclassicality, largely because such violations seem in tension with the relativistic speed limit. However, the precise definition of local causality and of realism have been the subject of much contention, within both the physics and philosophy of science communities. Indeed, it is not even clear that a notion of locality can be formulated without presupposing realism [213]. Furthermore, there are a host of different ways of parsing Bell's assumptions, and proponents of the different camps reach different conclusions about the lesson we should take from violations of Bell inequalities [319]. Although Bell's English-language definition of local causality is clear and well-motivated, his formalization of this definition as a particular set of conditional independences has been shown to be problematic [279] in light of the recent development of the framework of classical causal modeling [216, 290]. In particular, this framework demands that causal notions are fundamentally prior to probabilistic notions, and cannot be defined in terms of probabilities, e.g. in terms of conditional independences or probability raising.

In response to these problems, a growing line of work beginning with Wood and Spekkens [324] has aimed to recast Bell's theorem within the framework of causal models. In this approach, one derives *causal compatibility constraints* by assuming (i) a particular causal structure, encoded as a directed acyclic graph, and (ii) the assumption that observed correlations must be explained causally (within the framework of classical causal models, which assigns a random variable to each node in the graph, and assigns functional or stochastic dependences to the arrows in the graph). In the special case of Bell's theorem, the assumed causal structure of is that of Figure 1.1, and the causal compatibility constraints are Bell inequalities. Additionally, this new approach makes it clear how to extend the essence of Bell's reasoning to complex networks far beyond the standard Bell scenario, even those wherein space-time considerations (and hence 'locality') play no important role. For

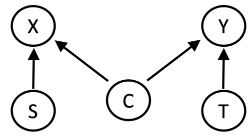

Figure 1.1: The natural causal structure in a Bell scenario.

any given causal structure, one can derive causal compatibility constraints and ask whether quantum resources are able to violate these constraints [56, 142]. If they can, then one obtains a quantum advantage for generating correlations within the assumed network. Such advantages often enable new forms of information-processing [86, 87, 97, 222].

Because quantum theory generates correlations that violate such causal compatibility constraints, it follows that the conjunction of assumptions (i) and (ii) cannot be satisfied. Either the causal structure is not the natural one, or one must give up on the idea that correlations can be explained causally within the framework of classical causal models. Assuming that a Bell scenario is described by a causal structure other than that in Fig. 1.1 requires advocating for radical possibilities such as superluminal signals or superdeterminism, each of which comes with a host of problems [261, 324]. Alternatively, giving up on the assumption that correlations can be explained causally is tantamount to giving up on scientific realism. To make progress without sacrificing either of these, one must seek causal explanations of observed correlations within a new *nonclassical* framework for causal modeling. For such an account to be meaningful, one must tweak only those assumptions of the classical causal modeling framework which one believes to encode 'classicality', while preserving (at least analogues of) the basic rules for carrying out causal reasoning. Although we are still lacking a fully fleshed out framework of this sort, recent progress has been made in defining quantum causal models [12, 29, 88]. In these, the causal relata are quantum systems rather than classical random variables, and the causal mechanisms are quantum processes rather than functional or stochastic dependences.

It is well known that entangled states enable us to violate Bell inequalities. In the language just introduced, this means that for the Bell causal structure (Fig. 1.1), quantum causal models are able to generate a strictly larger set of observed conditional probability distributions $P(XY|ST)$ than classical causal models. The essential component of such a quantum causal model is the entangled state corresponding to the common cause C in Fig. 1.1. This raises the natural question of how one can compare and quantify the resourcefulness of entangled states as nonclassical common-causes. Additionally, one can consider the set of common-cause realizable conditional probability distributions $P(XY|ST)$ as processes in their own right, and can then compare and quantify their nonclassicality properties.

To achieve this quantification, we follow the resource-theoretic approach of Ref. [80]. A resource theory is defined by a set of *free operations*. A valuable resource is one that cannot be constructed by composition of free operations, and one resource is more valuable than a second if it can be processed into the second by free operations. From this simple starting point, one can derive the full structure of a resource theory, including all quantitative measures of the notion of resourcefulness under consideration. In the special case wherein one wishes to study nonclassicality of common-cause processes (e.g. states, conditional probability distributions, channels, etc), my collaborators and I argue that the appropriate set of free operations is given by local operations and shared randomness (LOSR). Since shared randomness is just another word for 'classical common cause', it is clear that a process which is not LOSR-realizable constitutes a nonclassical common-cause resource. The most obvious examples of such resources are entangled quantum states and correlations that violate Bell inequalities, although we ultimately consider many other types of processes. Much of my research has aimed to quantify these different manifestations of nonclassicality.

Part II of this thesis, namely Chapters 4 and 5, constitutes a collection of some of my research on nonclassicality of common-cause processes.

Chapter 4 presents Section 2 of Ref. [322], which motivates our LOSR approach in more detail, and expands on the brief summary given just above. I omit the remainder of Ref. [322] from this thesis for the sake of length. In this omitted text, we develop the resource theory [80] of nonclassicality of common-cause boxes (what might colloquially be called the resource theory of 'nonlocality'). Therein, we provide a linear program for deciding which of two nonlocal boxes is more nonclassical, define two distinct monotones for nonlocal boxes, find simple closed form expressions for these in the simplest scenario, and prove a lower bound on the cardinality of any complete set of monotones. We also show that the information contained in the degrees of violation of facet-defining Bell inequalities is not sufficient for quantifying nonclassicality. We also prove a number of properties of the preorder over resources, including the fact that convexly extremal quantumly realizable correlations are all at the top of the order of quantumly realizable correlations.

Chapter 5 presents Ref. [258], which develops the resource theory of nonclassicality of common cause processes with arbitrary types of inputs and outputs (not just classical ones, as we considered in Ref. [322]). This resource theory subsumes a variety of quantum processes including quantum states [314], nonlocal boxes [44], steering assemblages [55, 320], channel steering assemblages [220], teleportages [57, 145], distributed measurements [37], measurement-device-independent steering channels [58], Bob-with-input steering channels [32], and generic no-signaling quantum channels [314]. We demonstrate that nonclassicality of common-cause is an umbrella notion of nonclassicality that unifies the study of all of these. Our framework naturally allows for conversions from any type of resource to any other, and for quantitative comparison of the nonclassicality of all resources across all types. We demonstrate the power of this framework by proving a number of abstract results about which types of resourcefulness can losslessly be encoded into resources of other types. As an application, we prove that resources of every type can have their nonclassicality characterized in a measurement-device-independent manner, greatly extending the scope of applicability of measurement-device-independent tests [41, 245].

Three of my papers on nonclassicality of common-cause processes were not included in this thesis, to keep the length down. These can be found in Refs. [256, 258, 322].

Ref. [322] was already discussed above.

Ref. [256] introduces a new branch of entanglement theory, based on local operations and shared randomness (rather than on local operations and classical communication). It also argues that this is in fact the relevant resource theory for understanding the interplay between entanglement and nonlocality. Aside from some *a priori* arguments (including those given in Chapter 5), we make this case by presenting three examples of the utility of the LOSR approach. Namely, we resolve the long-standing anomalies of nonlocality, we propose new definitions of genuine multipartite entanglement and nonlocality that are free from the pathological features exhibited by previous definitions, and we clarify and extend the notion of self-testing. We also prove some basic results regarding the notion of LOSR-entanglement.

Ref. [247] continues our work on the type-independent framework of Ref. [258] (in Chapter 5), but taking an algebraic rather than diagrammatic approach. Our main result is a hierarchy of semidefinite programs that can determine whether or not a given resource of any type is free or valuable. We also demonstrate how this hierarchy can be used to get explicit inequalities that witness the nonfreeness of any valuable resource, and to compute bounds on the value of LOSR monotones.

A fourth (somewhat) related paper that is not included in this thesis is Ref. [254], which considers the resource theory of local operations and shared entanglement (LOSE) [134] in the causal structure of Fig. 1.1. This is the appropriate set of free operations for studying postquantum common-cause resources. The type of questions we address in this work are very similar to those in Chapter 5, but with implications for our understanding of postquantumness rather than postclassicalness (that is, nonclassicality).

1.0.4 Causal-inferential theories

Finally, in Part III, namely Chapter 6, I present Ref. [261], our framework of causal-inferential theories, which weaves together the many threads of my research on nonclassicality.

This framework aims to sort out which elements of the quantum formalism refer to ontological concepts and which refer to epistemological concepts. The importance of settling this issue was famously noted by E.T. Jaynes [1]:

> [O]ur present [quantum mechanical] formalism is not purely epistemological; it is a peculiar mixture describing in part realities of Nature, in part incomplete human information about Nature — all scrambled up by Heisenberg and Bohr into an omelette that nobody has seen how to unscramble. Yet we think that the unscrambling is a prerequisite for any further advance in basic physical theory. For, if we cannot separate the subjective and objective aspects of the formalism, we cannot know what we are talking about; it is just that simple.

Taking the view that 'realities of Nature' are best understood as causal mechanisms acting on causal relata, we take the main aim of scientific realism to be providing a causal explanation of the world around us. Hence, the particular omelette of ontology and epistemology that we endeavor to unscramble is the one that results from the mixing up of the concepts of *epistemic inference* on the one hand, and of *causal influence* on the other. Such scramblings arise even in classical statistics, as in 'Simpson's paradox' [275] and 'Berkson's paradox' [39]. A satisfactory understanding of these phenomena was only found after the development of the mathematical framework of causal modeling [216, 290] that incorporated certain formal distinctions between inference and influence which are absent in the standard framework for statistical reasoning.

The difficulty of disentangling these is only compounded within the quantum formalism. One quantum example of such a scrambling arises when one studies quantum dynamics with initial system-environment correlations. In Ref. [257], my coauthors and I argue that maps that describe evolution are always completely positive, even in scenarios with initial correlations. We argue that previous work that claimed otherwise were led to confusion by their use of a definition of evolution maps which conflated inference and causation. This work is not reported in this thesis.

In Ref. [261], we aim to achieve an unscrambling of causation and inference akin to that achieved by the framework of causal modeling [216, 290]. However, we do so for *operational theories* (which characterize the scope of possible physical theories in a minimalist way, in terms of their operational predictions), and for *realist theories*, including ontological models, (which aim to provide causal explanations of the operational predictions of a given operational theory). Whether operational or realist, a theory in our framework is termed a *causal-inferential theory*, and is constructed out of two components:

- a causal theory, which describes physical systems in the world and the causal mechanisms that relate them, and
- an inferential theory, which describes an agent's beliefs about these systems and about the causal mechanisms that relate them, as well as how such beliefs are updated under the acquisition of new information.

The full causal-inferential theory is defined by the interplay between these two components, and allows one to describe a physical scenario in a manner that cleanly distinguishes causal and inferential aspects. Different causal-inferential theories can be obtained by varying the causal theory and/or varying the inferential theory.

One of the motivations for the standard framework for ontological models was to answer the question of whether the predictions of a given operational theory admit of an explanation in terms of an underlying ontology. The counterpart of this question in our new framework is whether the predictions of a given *operational* causal-inferential theory admit of an explanation in terms of an underlying *classical realist* causal-inferential theory, which represents systems as classical random variables and causal relations between systems as functional dependences. We refer to this as a *classical realist representation* of an operational causal-inferential theory.

The key constraint we impose on such representations is that they preserve the causal and inferential structures encoded in one's operational theory. We show that these two forms of structure preservation are motivated (respectively) by the two foundational principles discussed above. For example, in the case of a Bell scenario, the assumption that the causal structure (in Fig. 1.1) is respected by the classical realist representation leads to Bell inequalities. In more general causal scenarios [63, 104, 106, 216, 323], preservation of causal structure implies that any predictions that can be reproduced by such a representation must satisfy causal compatibility constraints. Meanwhile, no-go theorems arising from generalized noncontextuality arise from the constraint that *inferentially equivalent* processes must have representations in the model that are also inferentially equivalent. That is, if two operational procedures lead to all the same predictions in any experimental in which they are embedded, then the representations of these must *also* lead to all the same predictions, in any ontological scenario in which they can be embedded.

The conventional realist responses to the standard no-go theorems are unsatisfactory in various ways, such as requiring superluminal causal influences, requiring fine-tuning, and running afoul of Leibniz's principle. In light of this, it has been suggested that a more satisfactory way out of these no-go theorems may be achieved by modifying the notion of classical realist representations. This has been described in past work as 'going beyond the standard ontological models framework', but here is understood as seeking a nonclassical generalization of the notion of a classical realist representation. Our process-theoretic [78] framework provides the formal means of achieving this because it allows the interpretation of causal and inferential concepts to be determined by the axioms of the process theories that describe them and hence to differ from the conventional, classical interpretations of these concepts. This is analogous to how, in non-Euclidean geometries, the concepts of point and line acquire novel meanings distinct from their conventional (Euclidian) ones. Success in such a research program consists in finding a *nonclassical* notion of a realist causal-inferential theory that can provide a noncontextual representation of operational quantum theory, while respecting relativistic causal assumptions. This is the sense in which we aim to preserve the spirit of locality and of noncontextuality. While we do not yet have such a nonclassical realist theory, we do propose natural constraints on the axioms describing a theory of causal influences, a theory of epistemic inferences, and their interactions.

Thus, the work reported in this thesis provides a significant step forward in this research program. On the one hand, it provides, for the first time, a concrete proposal for the mathematical form of the sought-after theory, and, on the other hand, it provides a set of ideas for the form of its axioms, thereby providing a road map for future research.

Part I.

Generalized Noncontextuality

Chapter 2

A contextual advantage for state discrimination

Abstract: Finding quantitative aspects of quantum phenomena which cannot be explained by any classical model has foundational importance for understanding the boundary between classical and quantum theory. It also has practical significance for identifying information processing tasks for which those phenomena provide a quantum advantage. Using the framework of generalized noncontextuality as our notion of classicality, we find one such nonclassical feature within the phenomenology of quantum minimum-error state discrimination. Namely, we identify quantitative limits on the success probability for minimum-error state discrimination in any experiment described by a noncontextual ontological model. These constraints constitute noncontextuality inequalities that are violated by quantum theory, and this violation implies a quantum advantage for state discrimination relative to noncontextual models. Furthermore, our noncontextuality inequalities are robust to noise and are operationally formulated, so that any experimental violation of the inequalities is a witness of contextuality, independently of the validity of quantum theory. Along the way, we introduce new methods for analyzing noncontextuality scenarios and demonstrate a tight connection between our minimum-error state discrimination scenario and a Bell scenario.

2.1 Introduction

Understanding the boundary between the quantum and the classical is of fundamental importance for understanding quantum theory. One successful metric for nonclassicality, violation of Bell's notion of local causality [33], defines a clear departure from classicality in relativistic theories, but is relevant only for experiments with space-like separated measurements. Another notion of classicality, which concerns context-independence, was proposed by Kochen-Specker [162] and Bell [34], and has since been significantly refined and generalized [281]. It is the generalized notion of noncontextuality from Ref. [281] that we study in this chapter, but we refer to it simply as "noncontextuality" hereafter. As a metric for nonclassicality, the failure of noncontextuality has a broader scope than the failure of local causality insofar as it does not require space-like separation. It has also been shown to subsume many other pre-existing notions of nonclassicality, such as the negativity of quasi-probability representations [283], the generation of anomalous weak values [228], and even the aforementioned violations of local causality [281].

The quantum-classical boundary is also of practical importance in identifying tasks that admit of a quantum advantage. For example, violations of Bell inequalities have been shown to be a resource for device-independent key distribution [27], certified randomness [9], and communication complexity [47]. The failure of noncontextuality has also been shown to be a resource, leading to advantages for cryptography [14, 60, 286] and computation [144, 154, 234].

We here analyze minimum-error state discrimination (MESD) from the point of view of noncontextuality. Quantum state discrimination is a task wherein one must guess which quantum state describes a given quantum system when the state of that system is drawn from a known set of possibilities with a known prior distribution, and the estimation is based on the outcome of a measurement of one's choosing. In the "minimum error" variety of state discrimination, the objective is to minimize the probability that the estimate is in error. We here focus on the simplest case of a set containing just two states having equal a priori probability.

Although it is common to assert that the impossibility of perfectly discriminating nonorthogonal quantum states is an intrinsically nonclassical effect, this claim does not meet the minimal standard that one should require of *any* claim that some operational feature of quantum theory cannot be explained classically: namely, that it be justified by a rigorous no-go theorem. Such a theorem articulates a principle of classicality that has implications for operational statistics, and then proves that these implications are inconsistent with some operational feature(s) of quantum theory. Because the principle of noncontextuality constrains operational statistics and also has very broad scope, it is a particularly useful notion of classicality. If one does take it as one's principle of classicality, then the impossibility of discriminating nonorthogonal pure quantum states *cannot* be considered a nonclassical effect because there are subtheories of quantum theory (containing a strict subset of the states, measurements and transformations of the full theory) [284] wherein this phenomenon arises and that admit of a noncontextual model. (Within such models, the phenomenon can be attributed to the fact that the probability distributions associated with such quantum states are overlapping[1].) It follows that one must look at more nuanced aspects of the phenomenology of quantum state discrimination to identify features that are truly nonclassical by these lights.

We identify one such strongly nonclassical aspect of minimum error state discrimination: the particular dependence of the probability of successful discrimination on the overlap of the quantum states. For a given overlap, the quantum probability of discrimination is larger than can be accounted for by a noncontextual model. After presenting this result as a no-go theorem—that no noncontextual model can reproduce certain features of quantum MESD—we reformulate the problem in a manner that makes no reference to quantum theory, and which does not rely on any theoretical idealizations such as noise-free measurements or preparations. Our entirely operational formulation allows us to derive inequalities that can experimentally witness a contextual advantage for state discrimination, in the presence of noise and independently of the validity of quantum theory.

Our result identifies a key feature of quantum state discrimination that cannot be understood in any noncontextual model, and hence that is strongly nonclassical. Because quantum state discrimination is a primitive in many important quantum information processing protocols [37, 89], this work constitutes a first step towards identifying contextuality as a resource for more tasks concerning communication, computation, and cryptography.

We also prove an isomorphism between our operational MESD scenario and a two-party Bell test in which one party performs one of a pair of binary-outcome measurements and the other performs one of three binary-outcome measurements. This is similar to the fact that the noncontextuality inequality delimiting the success rate for parity-oblivious multiplexing [286] is isomorphic to the CHSH inequality in the Bell scenario [286].

Finally, we introduce two powerful new technical tools. First, we generalize existing methods for simulating exact operational equivalences [202]. Namely, while Ref. [202] shows how one may find a set of procedures that respects certain operational equivalences exactly, we have further demonstrated that one can find procedures that respect operational equivalences and *simultaneously* obey useful auxiliary constraints, such as the symmetries native to our ideal MESD scenario. This tool may have more general applications in the comparison of experimental data with theoretical expectations. More importantly, we find our noncontextuality inequalities using a novel algorithm (presented in Appendix A.2) for deriving the full set of necessary and sufficient noncontextuality inequalities for *any* finite

[1] Such models are ψ-epistemic, in the terminology of Ref. [139].

prepare-and-measure scenario, with respect to any fixed operational equivalences[2].

2.2 Operational Theories and Ontological Models

An operational theory is a specification of sets of primitive laboratory operations (e.g., preparations and measurements) and a prescription for finding the probabilities $p(k|M,P)$ for each outcome k given any measurement M performed on any preparation P. Two preparations P and P' are termed *operationally equivalent* if they cannot be differentiated by the statistics of any measurement; we denote this operational equivalence by

$$P \simeq P'. \tag{2.1}$$

In this chapter, quantum theory is understood as an operational theory. In the quantum formalism, the density operator specifies the statistics for all measurements, so that two preparation procedures are operationally equivalent if and only if they are represented by the same density operator.

An ontological model of an operational theory has the following form. To every system, there is associated an ontic state space Λ, where each ontic state $\lambda \in \Lambda$ specifies all the physical properties of the system. Each preparation P of a system is presumed to sample the system's ontic state λ at random from a probability distribution, denoted $\mu_P(\lambda)$ and termed the *epistemic state* associated with P, where

$$\forall\lambda: \quad 0 \leq \mu_P(\lambda), \tag{2.2}$$

$$\int_\Lambda d\lambda\, \mu_P(\lambda) = 1. \tag{2.3}$$

Each measurement M on a system is presumed to have its outcome k sampled at random in a manner that depends on the ontic state λ. The term *effect* will be used to refer to the pair consisting of a measurement, M, together with one of its outcomes, k, and will be denoted by $k|M$. The probability of outcome k given measurement M, considered as a function of λ, will be termed the *response function* associated with $k|M$, and denoted

[2] A full description of this algorithm can be found in Ref. [264].

$\xi_{k|M}(\lambda)$, where

$$\forall\lambda, \forall k: \quad 0 \leq \xi_{k|M}(\lambda), \tag{2.4}$$

$$\forall\lambda: \quad \sum_k \xi_{k|M}(\lambda) = 1. \tag{2.5}$$

Finally, an ontological model of an operational theory must reproduce the latter's empirical predictions; that is,

$$p(k|M,P) = \int_\Lambda d\lambda\, \xi_{k|M}(\lambda)\mu_P(\lambda). \tag{2.6}$$

We are now in a position to describe the assumption of preparation noncontextuality defined in Ref. [281]. An ontological model is said to be *preparation noncontextual* if it assigns the same epistemic state to all operationally equivalent preparations [281]:

$$P \simeq P' \implies \mu_P(\lambda) = \mu_{P'}(\lambda). \tag{2.7}$$

In operational quantum theory, the principle of preparation noncontextuality is respected whenever any two preparations that are associated to the same density operator are represented by the same epistemic state. For instance, different ensembles of states that average to the same mixed state (and for which one discards the information about which element of the ensemble was prepared) are operationally equivalent, and must be assigned the same epistemic state in a preparation noncontextual model.

Although there is a corresponding notion of measurement noncontextuality (namely, that operationally equivalent outcomes of measurements are represented by the same response functions), we will not have use of it in this article.

A few terminological conventions will be useful. A measurement is said to be represented as *outcome-deterministic* in the ontological model if the associated response functions all take values in $\{0,1\}$. The support of an epistemic state is defined as the set of $\lambda \in \Lambda$ which are assigned nonzero probability by it, $\text{supp}[\mu_P(\lambda)] \equiv \{\lambda : \mu_P(\lambda) \neq 0\}$, while the support of a response function is defined as the set of $\lambda \in \Lambda$ for which the response function is nonzero, $\text{supp}[\xi_{k|M}(\lambda)] \equiv \{\lambda : \xi_{k|M}(\lambda) \neq 0\}$.

2.3 Quantum Minimum Error State Discrimination

We begin with the problem of discriminating two nonorthogonal pure quantum states $|\phi\rangle$ and $|\psi\rangle$. These two states span a 2-dimensional space, so we can represent them as points in an equatorial plane of the Bloch ball, as in Fig. 2.1.

First, we consider the operational signature of their nonorthogonality. A measurement of the ϕ basis, $B_\phi \equiv \{|\phi\rangle\langle\phi|, |\bar{\phi}\rangle\langle\bar{\phi}|\}$, perfectly distinguishes between state $|\phi\rangle$ and its complement; we denote the associated outcomes by ϕ and $\bar{\phi}$, respectively. A measurement of the ψ basis, $B_\psi \equiv \{|\psi\rangle\langle\psi|, |\bar{\psi}\rangle\langle\bar{\psi}|\}$, does the same for the state $|\psi\rangle$ and its complement, with associated outcomes ψ and $\bar{\psi}$. If one implements the ψ basis measurement on the

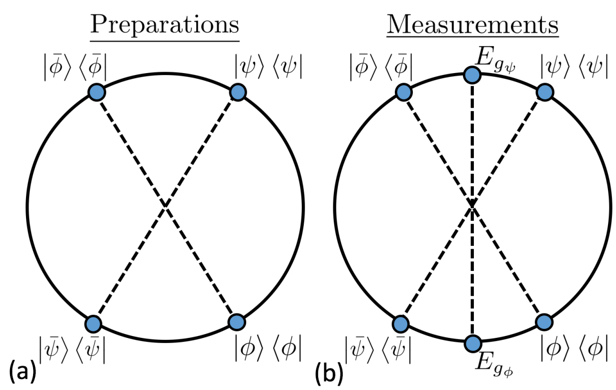

Figure 2.1: The quantum states and measurements in our scenario, depicted as Bloch vectors in an equatorial plane of the Bloch ball.

state ϕ, the probability of obtaining the ψ outcome is

$$c_q = \mathrm{Tr}[|\phi\rangle\langle\phi|\psi\rangle\langle\psi|] = |\langle\phi|\psi\rangle|^2, \tag{2.8}$$

Because one could think of this quantity as the probability that ϕ passes the test for ψ and thus is confusable with ψ, we henceforth call it the *confusability*. Note that if one implements the ϕ basis measurement on the state ψ, the probability of obtaining the ϕ outcome is also c_q.

If $|\phi\rangle$ and $|\psi\rangle$ have nonzero confusability (i.e., if they are not orthogonal), then no measurement can distinguish between the two without incurring a nonzero probability of error. We denote the discriminating measurement by $B_d \equiv \{E_{g_\phi}, E_{g_\psi}\}$, where the outcome for which one should guess ϕ (respectively ψ) is denoted g_ϕ (respectively g_ψ). Assuming equal prior probabilities of $|\phi\rangle$ and $|\psi\rangle$, the probability of guessing the state correctly with this measurement is

$$s_q \equiv \frac{1}{2}\mathrm{Tr}[E_{g_\phi}|\phi\rangle\langle\phi|] + \frac{1}{2}\mathrm{Tr}[E_{g_\psi}|\psi\rangle\langle\psi|]. \tag{2.9}$$

We assume that the discriminating measurement has the natural symmetry property $\mathrm{Tr}[E_{g_\phi}\,|\phi\rangle\,\langle\phi|] = \mathrm{Tr}[E_{g_\psi}\,|\psi\rangle\,\langle\psi|]$ so that

$$s_q = \mathrm{Tr}[E_{g_\phi}\,|\phi\rangle\,\langle\phi|] = \mathrm{Tr}[E_{g_\psi}\,|\psi\rangle\,\langle\psi|]. \tag{2.10}$$

The measurement scheme that yields the greatest probability of guessing correctly which of two nonorthogonal states was prepared is called the *minimum error* state discrimination (MESD) scheme. Since $|\phi\rangle$ and $|\psi\rangle$ are prepared with equal probability, the POVM $\{E_{g_\phi}, E_{g_\psi}\}$ achieving MESD is the one consisting of projectors onto the basis that straddles $|\phi\rangle$ and $|\psi\rangle$ in Hilbert space, which is depicted in the Bloch sphere in Fig. 2.1. This is called the *Helstrom measurement* [141]. It is well-known that the probability of guessing the state correctly using the Helstrom measurement is

$$s_q = \frac{1}{2}(1+\sqrt{1-|\,\langle\phi|\psi\rangle\,|^2}) = \frac{1}{2}(1+\sqrt{1-c_q}). \tag{2.11}$$

We have now described all of the preparations and measurements that usually appear in a discussion of the problem of discriminating two nonorthogonal quantum states, and some basic facts about the relations that hold among the operational quantities characterizing the discrimination problem (i.e., facts about the phenomenology of quantum state discrimination). However, these facts are insufficient for deriving a no-go theorem for noncontextuality. The reason is that the preparations and measurements described thus far do not exhibit any operational equivalences via which the assumption of noncontextuality could imply nontrivial constraints on the ontological model.

However, there is a simple solution: we also consider the problem of discriminating the pair of quantum states that are complementary to $|\phi\rangle$ and $|\psi\rangle$, namely, $|\bar{\phi}\rangle$ and $|\bar{\psi}\rangle$, also depicted in Fig. 2.1. By symmetry, the confusability of $|\bar{\phi}\rangle$ and $|\bar{\psi}\rangle$ is also equal to c_q, and the success rate for distinguishing $|\bar{\phi}\rangle$ and $|\bar{\psi}\rangle$ when they have equal prior probability is also equal to s_q (where the optimal measurement is again $\{E_{g_\phi}, E_{g_\psi}\}$, but now the outcomes g_ϕ and g_ψ signal one to guess preparations $|\bar{\psi}\rangle$ and $|\bar{\phi}\rangle$, respectively). So the $|\bar{\phi}\rangle$ vs. $|\bar{\psi}\rangle$ discrimination problem is a mirror image of the $|\phi\rangle$ vs. $|\psi\rangle$ discrimination problem, and consequently does not require specifying any additional facts about the phenomology of quantum state discrimination. However, the inclusion of $|\bar{\phi}\rangle$ and $|\bar{\psi}\rangle$ in our analysis provides us with a nontrivial operational equivalence relation among the preparations, namely,

$$\frac{1}{2}\,|\phi\rangle\,\langle\phi| + \frac{1}{2}\,|\bar{\phi}\rangle\,\langle\bar{\phi}| = \frac{1}{2}\,|\psi\rangle\,\langle\psi| + \frac{1}{2}\,|\bar{\psi}\rangle\,\langle\bar{\psi}| = \frac{\mathbb{1}}{2}. \tag{2.12}$$

We will show that this equivalence relation together with the phenomenology of quantum state discrimination described above is sufficient to derive a no-go theorem for noncontextuality.

The probability of a given measurement outcome occurring on a given preparation, for every possible pairing thereof, is summarized in Table 2.1. Here, the columns correspond to the distinct state-preparations and the rows correspond to the distinct effects (where one need only include a single effect for each binary-outcome measurement given that the probability for the other effect is fixed by normalization).

	$\lvert\phi\rangle$	$\lvert\psi\rangle$	$\lvert\bar{\phi}\rangle$	$\lvert\bar{\psi}\rangle$
$\lvert\phi\rangle\langle\phi\rvert$	1	c_q	0	$1-c_q$
$\lvert\psi\rangle\langle\psi\rvert$	c_q	1	$1-c_q$	0
E_{g_ϕ}	s_q	$1-s_q$	$1-s_q$	s_q

Table 2.1: Data table in the ideal quantum case.

2.4 Noncontextuality no-go theorem for MESD in Quantum Theory

The fact that the ontological model must reproduce the probabilities in Table 2.1 via Eq. (2.6) implies constraints on the epistemic states associated to the four preparations and the response functions associated to the three effects. For instance, to reproduce the first column of the table, one requires that

$$\int_\Lambda d\lambda\, \xi_{\phi|B_\phi}(\lambda)\mu_\phi(\lambda) = 1, \tag{2.13}$$

$$\int_\Lambda d\lambda\, \xi_{\psi|B_\psi}(\lambda)\mu_\phi(\lambda) = c_q, \tag{2.14}$$

$$\int_\Lambda d\lambda\, \xi_{g_\phi|B_d}(\lambda)\mu_\phi(\lambda) = s_q. \tag{2.15}$$

Given that convex mixtures of preparations are represented in an ontological model by the corresponding mixture of epistemic states (see Eq. (7) of [283] and the surrounding discussion), it follows that $\frac{1}{2}\lvert\phi\rangle\langle\phi\rvert + \frac{1}{2}\lvert\bar{\phi}\rangle\langle\bar{\phi}\rvert$ is represented by $\frac{1}{2}\mu_\phi(\lambda) + \frac{1}{2}\mu_{\bar{\phi}}(\lambda)$, and

$\frac{1}{2}|\psi\rangle\langle\psi|+\frac{1}{2}|\bar{\psi}\rangle\langle\bar{\psi}|$ is represented by $\frac{1}{2}\mu_\psi(\lambda)+\frac{1}{2}\mu_{\bar{\psi}}(\lambda)$. But because both of these mixtures of preparations are associated to the completely mixed state (Eq. (2.12)), they are operationally equivalent, and thus by the assumption of preparation noncontextuality, they are represented by the same epistemic state. It follows that

$$\frac{1}{2}\mu_\phi(\lambda)+\frac{1}{2}\mu_{\bar{\phi}}(\lambda)=\frac{1}{2}\mu_\psi(\lambda)+\frac{1}{2}\mu_{\bar{\psi}}(\lambda). \tag{2.16}$$

Any ontological model satisfying noncontextuality, and consequently Eq. (2.16), and reproducing the form of the data in Table 2.1, and consequently Eqs. (2.13)-(2.15) and their kin, can be shown to satisfy the following trade-off between s_q and c_q:

$$s_q \leq 1-\frac{c_q}{2}. \tag{2.17}$$

An intuitive proof is provided in Section 2.4.1, where we also discuss how this result is related to the results of Refs. [24, 179, 195]. (In Appendix A.1, we provide a proof using more general methods, which generalizes more easily to the noisy case discussed later, in Section 2.6.)

This tradeoff relation contradicts the one known to be optimal in quantum theory, Eq. (2.11). The optimal quantum tradeoff generally allows *higher* success rates for a given confusability than the noncontextual tradeoff. Therefore, we conclude that the phenomenology of minimum-error state discrimination in the noiseless quantum case is inconsistent with the principle of noncontextuality.

In Fig. 2.2, we plot the maximum success rate for MESD as a function of the confusability for both quantum theory (Eq. (2.11)) and for a noncontextual model (the tradeoff that saturates the inequality of Eq. (2.17)).

2.4.1 Intuitive proof of the noncontextual tradeoff

We now introduce some basic facts from classical probability theory, which we then leverage to prove Eq. (2.17).

Suppose that a classical variable λ has been sampled from one of two overlapping probability distributions, $p(\lambda|a)$ and $p(\lambda|b)$. Absent additional information, it is straightforward to see that in trying to guess which of the two distributions a given λ was drawn from, one cannot do better than guessing 'distribution a' for the values of λ for which $p(a|\lambda) > p(b|\lambda)$, and guessing 'distribution b' when the opposite is true. (Of course, it is irrelevant what one

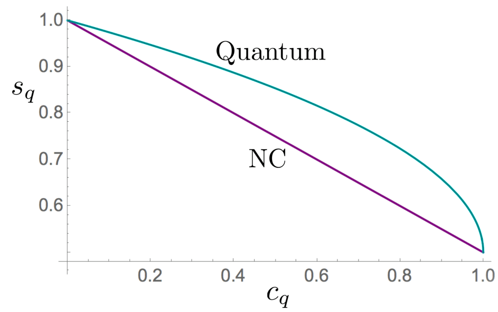

Figure 2.2: Optimal tradeoff for a noncontextual model (purple line) and for quantum theory (light blue curve).

guesses for the values of λ for which $p(a|\lambda) = p(b|\lambda)$.) In the special case we are considering, with equal prior probability $p(a) = p(b) = \frac{1}{2}$ for the two options, if we perform a Bayesian inversion, we find $p(\lambda|a) > p(\lambda|b)$ if and only if $p(a|\lambda) > p(b|\lambda)$, and hence one should guess 'distribution a' for the values of λ for which $p(\lambda|a) > p(\lambda|b)$, and guess 'distribution b' when the opposite is true.

The probability that the guess $g \in \{a, b\}$ was correct given a particular value of λ is simply $p(g|\lambda)$. Since we always guess the distribution a or b that has the higher likelihood of being correct, the probability that we are right in each run is simply $\max\{p(a|\lambda), p(b|\lambda)\}$. On average, then, the success probability r is

$$
\begin{aligned}
r &= \int_\Lambda d\lambda\ p(\lambda)\max\{p(a|\lambda), p(b|\lambda)\} && (2.18)\\
&= \int_\Lambda d\lambda\ p(\lambda)(1 - \min\{p(a|\lambda), p(b|\lambda))\} && (2.19)\\
&= 1 - \int_\Lambda d\lambda\ \min\{p(a|\lambda)p(\lambda), p(b|\lambda)p(\lambda)\} && (2.20)\\
&= 1 - \int_\Lambda d\lambda\ \min\{p(\lambda|a)p(a), p(\lambda|b)p(b)\} && (2.21)\\
&= 1 - \frac{1}{2}\int_\Lambda d\lambda\ \min\{p(\lambda|a), p(\lambda|b)\}, && (2.22)
\end{aligned}
$$

where the equality on line (2.19) uses the fact that $p(a|\lambda)+p(b|\lambda) = 1$ for all λ. The quantity $\int_\Lambda d\lambda\ \min\{p(\lambda|a), p(\lambda|b)\}$ is termed the *classical overlap* of the probability distributions $p(\lambda|a)$ and $p(\lambda|b)$.

In an MESD scenario, the task is to guess, in each particular run of the experiment, whether a system was prepared by state-preparation $|\phi\rangle$ or by state-preparation $|\psi\rangle$. If the experiment is described by an ontological model, then this task corresponds to guessing, from a single sample of the ontic state λ of the system, whether it was sampled from the distribution $\mu_\phi(\lambda)$ or from $\mu_\psi(\lambda)$. Given that we do not assume any operational equivalence relations among the measurements in the experiment, the assumption of measurement noncontextuality does not place any constraints on the ontological representation of the measurements. Therefore, in particular, the Helstrom measurement is *at best* represented in the ontological model by the set of response functions that yield the maximum probability of guessing which distribution the ontic state λ was sampled from. From our discussion concerning two overlapping classical probability distributions, it is clear that this corresponds to a measurement that returns the g_ϕ outcome whenever $\mu_\phi(\lambda) > \mu_\psi(\lambda)$ and the g_ψ outcome whenever $\mu_\phi(\lambda) < \mu_\psi(\lambda)$, and that the probability of guessing correctly based on the outcome of the Helstrom measurement is upper bounded as follows:[3]

$$s_q \leq 1 - \frac{1}{2}\int_\Lambda d\lambda \, \min\{\mu_\phi(\lambda), \mu_\psi(\lambda)\}. \tag{2.23}$$

We will now show that in a noncontextual model,

$$c_q = \int_\Lambda d\lambda \, \min\{\mu_\phi(\lambda), \mu_\psi(\lambda)\}, \tag{2.24}$$

so that substituting Eq. (2.24) into Eq. (2.23), we infer that $s_q \leq 1 - \frac{c_q}{2}$, the noncontextual bound on the trade-off between s_q and c_q described in Eq. (2.17).

First, in any preparation noncontextual model the response function $\xi_i(\lambda)$ for a projector onto pure state $|i\rangle$ satisfies

$$\xi_i(\lambda) = \begin{cases} 1, & \text{if } \lambda \in \text{supp}[\mu_i(\lambda)] \\ 0, & \text{otherwise.} \end{cases} \tag{2.25}$$

This *outcome determinism for sharp measurements* was first proven in Ref. [281]. It can be seen by considering the projector as part of some projective measurement M with effects $\{E_i = |i\rangle\langle i|\}$, and the corresponding basis of pure states $\{\rho_i = |i\rangle\langle i|\}$, so that $\text{Tr}[E_i\rho_j] = \delta_{i,j}$. Denoting the epistemic state of ρ_j as $\mu_j(\lambda)$ and the response function for

[3] A mathematically equivalent version of this upper bound was previously proven under different assumptions in Refs. [24, 239]. The former article considered the assumption that this inequality is saturated as a constraint on ontological models, which they termed "maximal ψ-episemicity". (Note that this constraint is different from the constraint considered in Ref. [179] even though it has the same name.)

E_i as $\xi_{i|M}(\lambda)$, this implies that $\int \mu_j(\lambda)\xi_{i|M}(\lambda)d\lambda = \delta_{i,j}$. Because $\mu_j(\lambda)$ is a normalized probability distribution, this implies that, for *any* ontological model,

$$\xi_{i|M}(\lambda) = \begin{cases} 1, & \text{if } \lambda \in \text{supp}[\mu_i(\lambda)] \\ 0, & \text{if } \lambda \in \text{supp}[\mu_{j\neq i}(\lambda)]. \end{cases} \tag{2.26}$$

Eq. (2.26) is not equivalent to Eq. (B.3), since there may exist ontic states that are not in the support of *any* of the $\mu_i(\lambda)$, and Eq. (2.26) does not constrain such ontic states in any way. In a preparation noncontextual model, however, we can furthermore show that there are no ontic states outside of the union of the supports of the set of basis states, $\cup_i \text{supp}[\mu_i(\lambda)]$, as follows. Every density operator ρ appears in *some* decomposition of the maximally mixed state $\frac{1}{d}\mathbb{1}$. By preparation noncontextuality, every such decomposition has the *same* distribution $\mu_{\frac{1}{d}\mathbb{1}}(\lambda)$ over ontic states. Thus, every ontic state in the support of the corresponding $\mu_\rho(\lambda)$ *also* appears in the support of $\mu_{\frac{1}{d}\mathbb{1}}(\lambda)$, so the full state space Λ is equivalent to $\text{supp}[\mu_{\frac{1}{d}\mathbb{1}}(\lambda)]$. Furthermore, for the basis of states $\{\rho_i\}$ above, $\frac{1}{d}\sum_i \rho_i = \frac{1}{d}\mathbb{1}$, so preparation noncontextuality implies that $\sum_i \frac{1}{d}\mu_i(\lambda) = \mu_{\frac{1}{d}\mathbb{1}}(\lambda)$, and therefore $\cup_i \text{supp}[\mu_i(\lambda)] = \text{supp}[\mu_{\frac{1}{d}\mathbb{1}}(\lambda)] = \Lambda$. Thus every ontic state λ must be in the support of exactly one of the ρ_i, and Eq. (2.26) can be strengthened to Eq. (B.3).

Recalling the expression for the confusability of quantum states $|\phi\rangle$ and $|\psi\rangle$ in an ontological model, $c_q = \int_\Lambda d\lambda\ \xi_{\phi|B_\phi}(\lambda)\mu_\psi(\lambda)$, Eq. (B.3) implies that for a preparation noncontextual model:

$$c_q = \int_{\text{supp}[\mu_\phi(\lambda)]} d\lambda\ \mu_\psi(\lambda). \tag{2.27}$$

By virtue of the symmetry of the problem, the analogous expression with the roles of ϕ and ψ reversed also holds. The fact that the expression for the ideal confusability $c_q = |\langle\phi|\psi\rangle|^2$ of ϕ and ψ in a preparation-noncontextual model is given by Eq. (2.27) was noted by Leifer and Maroney [179].

The second implication of preparation noncontextuality which we require to prove Eq. (2.24) is that for each of the four quantum states $\Psi \in \{\phi, \psi, \bar{\phi}, \bar{\psi}\}$, $\mu_\Psi(\lambda) = 2\mu_{\frac{\mathbb{1}}{2}}(\lambda)$ for all $\lambda \in \text{supp}[\mu_\Psi(\lambda)]$, where $\mu_{\frac{\mathbb{1}}{2}}(\lambda)$ is the distribution associated with the maximally mixed state $\frac{\mathbb{1}}{2}$. This was also first proven in Ref. [281], and follows immediately from preparation noncontextuality, $\frac{1}{2}\mu_\phi(\lambda) + \frac{1}{2}\mu_{\bar{\phi}}(\lambda) = \frac{1}{2}\mu_\phi(\lambda) + \frac{1}{2}\mu_{\bar{\phi}}(\lambda) = \mu_{\frac{\mathbb{1}}{2}}(\lambda)$, and the fact that an ontic state can be in the support of at most one state from a set of orthogonal states; that is, $\mu_\phi(\lambda)\mu_{\bar{\phi}}(\lambda) = 0$ and $\mu_\psi(\lambda)\mu_{\bar{\psi}}(\lambda) = 0$.

Hence for all $\lambda \in \mathrm{supp}[\mu_\phi(\lambda)] \cap \mathrm{supp}[\mu_\psi(\lambda)]$, we have $\mu_\phi(\lambda) = \mu_\psi(\lambda) = 2\mu_{\frac{1}{2}}(\lambda)$. It follows that $\min\{\mu_\phi(\lambda), \mu_\psi(\lambda)\} = \mu_\phi(\lambda) = \mu_\psi(\lambda)$ for all $\lambda \in \mathrm{supp}[\mu_\phi(\lambda)] \cap \mathrm{supp}[\mu_\psi(\lambda)]$, and is equal to 0 everywhere else, and consequently

$$\begin{aligned}
&\int_{\mathrm{supp}[\mu_\phi(\lambda)]} d\lambda\ \mu_\psi(\lambda) \\
&\quad = \int_{\mathrm{supp}[\mu_\phi(\lambda)] \cap \mathrm{supp}[\mu_\psi(\lambda)]} d\lambda\ \mu_\psi(\lambda) \\
&\quad = \int_{\mathrm{supp}[\mu_\phi(\lambda)] \cap \mathrm{supp}[\mu_\psi(\lambda)]} d\lambda\ \min\{\mu_\phi(\lambda), \mu_\psi(\lambda)\} \\
&\quad = \int_\Lambda d\lambda\ \min\{\mu_\phi(\lambda), \mu_\psi(\lambda)\},
\end{aligned} \tag{2.28}$$

Finally, Eq. (2.27) and Eq. (2.28) together imply Eq. (2.24), which is what we sought to prove.

2.4.2 Graphical summary of the proof

The intuitive proof is best summarized graphically, by contrasting a preparation-contextual ontological model, Fig. 2.3, with a preparation noncontextual ontological model, Fig. 2.4. For visual simplicity, we have chosen a continuous, 1-dimensional, bounded ontic state space. We arrange the state space into a circle, so that each point on the circle is a unique ontic state, and epistemic states are represented as probability distributions on the surface of the circle (where the probability density corresponds to the radial height). In each figure, we show the epistemic states for the four preparations and for the two mixed preparations, the classical overlap for two epistemic states, a representative response function, and the confusability generated by that response function. We then show that in the contextual model, the classical overlap and confusability can differ, while in the noncontextual model, they must be identical.

In the ontological model of an MESD scenario shown in Fig. 2.3, the distributions $\frac{1}{2}\mu_\phi(\lambda) + \frac{1}{2}\mu_{\bar{\phi}}(\lambda)$ and $\frac{1}{2}\mu_\psi(\lambda) + \frac{1}{2}\mu_{\bar{\psi}}(\lambda)$ are not identical; hence, this model is preparation-contextual. The classical overlap $\int_\Lambda d\lambda\ \min\{\mu_\phi(\lambda), \mu_\psi(\lambda)\}$ is equal to the area of the shaded region in (g). The response function $\xi_{\phi|B_\phi}(\lambda)$ must have value 0 on the support of $\mu_{\bar{\phi}}(\lambda)$ and value 1 on the support of $\mu_\phi(\lambda)$, as pictured in (h); however, in the region outside both of these supports, its value is arbitrary, as indicated schematically. Given the response function pictured, the confusability $c_q = \int_\Lambda d\lambda \xi_{\phi|B_\phi}(\lambda)\mu_\psi(\lambda)$ equals the area of the shaded region in (i). One can clearly see that the classical overlap and the confusability need not be the same in a preparation-contextual model.

In the ontological model of an MESD scenario shown in Fig. 2.4, the distributions $\frac{1}{2}\mu_\phi(\lambda) + \frac{1}{2}\mu_{\bar{\phi}}(\lambda)$ and $\frac{1}{2}\mu_\psi(\lambda) + \frac{1}{2}\mu_{\bar{\psi}}(\lambda)$ *are* identical; hence, this model is preparation-noncontextual. Furthermore, these two distributions are equal to the unique distribution $\mu_{\mathbb{1}/2}(\lambda)$ (whose support must span the entire ontic state space), and the epistemic states $\mu_\phi(\lambda)$, $\mu_{\bar{\phi}}(\lambda)$, $\mu_\psi(\lambda)$, and $\mu_{\bar{\psi}}(\lambda)$ must both be equal on their support to $2\mu_{\mathbb{1}/2}(\lambda)$. Thus, in a preparation-noncontextual model, the classical overlap is given simply by the integral of $2\mu_{\mathbb{1}/2}(\lambda)$ in the region of common support, as shown by the shaded region in (g). Furthermore, preparation noncontextuality implies that the response function $\xi_{\phi|B_\phi}(\lambda)$ is 1 on the support of $\mu_\phi(\lambda)$ and 0 on *all* other ontic states, as shown in (h). Given this form for the response function, the confusability $c_q = \int_\Lambda d\lambda \xi_{\phi|B_\phi}(\lambda)\mu_\psi(\lambda)$ is given by the area of the shaded region in (i). Clearly, the classical overlap and the confusability are identical in a preparation-noncontextual model.

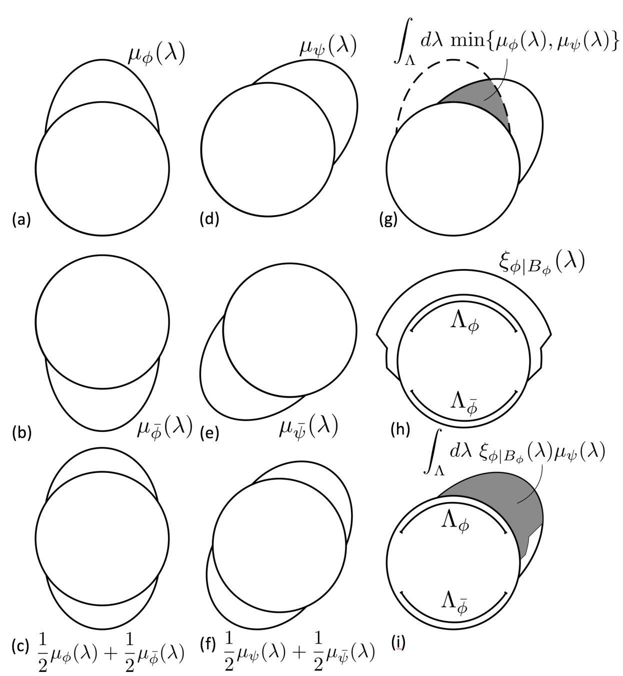

Figure 2.3: In a contextual model of an MESD scenario: (a)-(f) Epistemic states; (g) Classical overlap between $\mu_\phi(\lambda)$ and $\mu_\psi(\lambda)$; (h) Response function $\xi_{\phi|B_\phi}(\lambda)$, with indication of $\Lambda_\phi \equiv \text{supp}[\mu_\phi(\lambda)]$ and $\Lambda_{\bar{\phi}} \equiv \text{supp}[\mu_{\bar{\phi}}(\lambda)]$; (i) Confusability defined by $\xi_{\phi|B_\phi}(\lambda)$, also with indication of Λ_ϕ and $\Lambda_{\bar{\phi}}$.

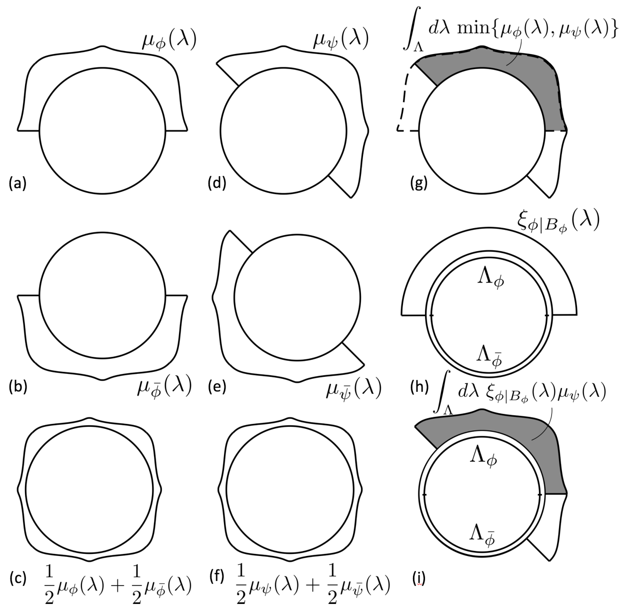

Figure 2.4: In a noncontextual model of an MESD scenario: (a)-(f) Epistemic states; (g) Classical overlap between $\mu_\phi(\lambda)$ and $\mu_\psi(\lambda)$; (h) Response function $\xi_{\phi|B_\phi}(\lambda)$, with indication of $\Lambda_\phi \equiv \text{supp}[\mu_\phi(\lambda)]$ and $\Lambda_{\bar\phi} \equiv \text{supp}[\mu_{\bar\phi}(\lambda)]$; (i) Confusability defined by $\xi_{\phi|B_\phi}(\lambda)$, also with indication of Λ_ϕ and $\Lambda_{\bar\phi}$..

2.4.3 Relation to previous work

Leifer and Maroney [179] consider the assumption that Eq. (2.27) should hold for *every* possible pair of quantum states ϕ and ψ as a constraint on ontological models that is worthy of investigation in its own right. They term ontological models that satisfy this assumption *maximally ψ-epistemic*. As we noted in Sec. 2.4.1 (and as demonstrated in their article), this assumption follows from preparation noncontextuality (and hence from universal noncontextuality). However, Leifer and Maroney investigate the consequences of making the assumption of maximal ψ-epistemicity without also assuming other consequences of universal noncontextuality, in particular, without assuming other consequences of preparation noncontextuality.

They establish their no-go theorem for maximal ψ-epistemicity (and hence for universal noncontextuality) by demonstrating that maximal ψ-epistemicity implies the Kochen-Specker notion of noncontextuality (which is measurement noncontextuality together with the assumption of outcome determinism for sharp measurements), and then relying on the fact that quantum theory does not admit of a Kochen-Specker noncontextual model (the Kochen-Specker theorem).

Both this chapter and their article explore senses in which a pair of quantum states may be said to be "indistinguishable", and to what extent some operational counterpart of this indistinguishability can be explained in an ontological model satisfying certain properties. But there are key differences. As we've noted, the property of ontological models that we focus on is different: we consider the assumption of universal noncontextuality rather than just maximal ψ-epistemicity.[4] The more important difference between our work and that of Leifer and Maroney, however, is in how we operationalize the notion of indistinguishability.

To explain the difference, it is useful to highlight two distinct facts about a pair of nonorthogonal pure quantum states (i.e., a pair $|\psi\rangle$ and $|\phi\rangle$ for which $|\langle\psi|\phi\rangle|^2 > 0$): (i) they are not perfectly *discriminable*, which is to say that there is no quantum measurement that achieves zero error in the discrimination task, formalized as $s_q > 0$, and (ii) they are *confusable*, which is to say that the ideal quantum measurement that tests for being in the state $|\phi\rangle$ has a nonzero probability of being passed by the state $|\psi\rangle$, and similarly for $|\phi\rangle$ and $|\psi\rangle$ interchanged, formalized as $c_q > 0$.

The determination of the maximum probability of discrimination for a given confusability, that is, the optimal tradeoff relation that holds between s_q and c_q, is one of the central results in the field of quantum state estimation. Our work seeks to determine constraints on this tradeoff relation from assumptions about the ontological model.

Leifer and Maroney, by contrast, do not consider this tradeoff relation, nor the expression for the discriminability of quantum states. Rather, they address (and answer in the negative) the question of whether the degree of confusability of nonorthogonal pure quantum states can be given a particular expression in the ontological model, namely, that of Eq. (2.27), which asserts that the test associated to the state $|\phi\rangle$ is a test for whether the ontic state λ is inside the ontic support of the distribution representing $|\phi\rangle$.[5] While the expression for the confusability of two quantum states is a feature of their indistinguishability, it is not one that has previously been of interest in the field of quantum state estimation.

[4]Whereas we believe that the assumption of universal noncontextuality is well motivated (namely, by Leibniz's principle of the identity of indiscernibles), it is unclear to us whether any motivation can be given for maximal ψ-epistemicity that is not simultaneously a motivation for universal noncontextuality. Therefore, unlike Ref. [195], we remain unconvinced that the assumption that Eq. (2.27) holds for every pair of quantum states is interesting in its own right.

[5]This is the sort of explanation one obtains in the toy theory model of the single qubit stabilizer subtheory of quantum theory [282] or the Kochen-Specker model of a single qubit [162]. Note that this is not the only way to explain the degree of confusability; the response function for $|\phi\rangle$ might be nontrivial outside the ontic support of the distribution representing $|\phi\rangle$ and even indeterministic in that region, and if so, one can have a nonzero confusability even though μ_ϕ and μ_ψ have disjoint ontic supports.

Thus, whereas Leifer and Maroney show the impossibility of a particular ontological expression for the confusability from a known no-go result for Kochen-Specker noncontextuality (the Kochen-Specker theorem), we begin with the native phenomenology of minimum-error state discrimination (the quantum tradeoff between s_q and c_q), and we derive a novel no-go result for universal noncontextuality from it.

The form of the tradeoff relation between discriminability and confusability has relevance for quantum information processing tasks that make use of minimum error state discrimination. For instance, it is used in Ref. [287] to derive the tradeoff relation between concealment and bindingness in quantum bit commitment protocols [189, 200], and such protocols can be used as subroutines in protocols for other tasks, such as strong coin flipping [13, 287]. It has also used in the analysis of quantum protocols for the task of oblivious transfer [61]. Our results may be useful, therefore, in determining whether or not the failure of universal noncontextuality is a resource for such tasks.

Note that because MESD for two pure quantum states is a phenomenon occuring in a two-dimensional Hilbert space (the subspace spanned by the two states) while the Kochen-Specker theorem can only be proven in Hilbert spaces of dimension three or greater, there is no possibility of leveraging facts about Kochen-Specker-uncolourable sets to infer anything about which aspects of MESD resist explanation within a universally noncontextual model.[6]

A final crucial advantage of our approach over that of Ref. [179] is that it can be used to derive noncontextuality inequalities that are noise-robust and hence experimentally testable, as we will show in the next section. Noise-robustness is critical if one hopes to leverage contextuality as a resource in real (hence noisy) implementations of information-processing protocols.

2.5 Dealing with noise

It is important to recognize that the inequality of Eq. (2.17) is not experimentally testable. To clarify this point, we first draw a distinction between *noncontextuality no-go results* and *noncontextuality inequalities*. A *noncontextuality no-go result* is a proof that no noncontextual model can reproduce certain predictions of quantum theory; as such, a no-go result can contain idealizations (such as perfect correlations) which are justified by

[6] The reason there is no possibility of proving the Kochen-Specker theorem with projective measurements in dimension 2 is that no projector appears in more than a single context [139, 162]. By contrast, it is known that there are proofs of the failure of preparation noncontextuality that hold even in 2-dimensional Hilbert spaces [281], and the proof we have presented here is of this type.

quantum theory but which never hold in real experiments. In some cases (as above), a no-go result may derive an inequality on the way to deriving a logical contradiction, but such an inequality may not qualify as a proper noncontextuality inequality. In our usage, a *noncontextuality inequality* makes no reference to the quantum formalism and must not invoke idealized assumptions in its derivation. We give such an inequality for MESD in Section 2.6.

The distinction between no-go results and robust inequalities has historical precedent. In his 1964 paper [33], in deriving an inequality that could be shown to be violated by quantum correlations, Bell assumed an experiment wherein certain pairs of measurements had perfectly correlated outcomes. Such perfect correlations hold for ideal quantum states and measurements, but are never observed in nature. Hence, Bell's 1964 result is a no-go result, with consequences for the interpretation of quantum theory, but the inequality he derives en route to this contradiction does not provide a means of experimentally testing the principle of local causality. In 1969, Clauser, Horne, Shimony, and Holte [75] derived an inequality without assuming these idealizations. Because their inequality makes no reference to perfect correlations or to any other feature of quantum theory, its violation rules out all locally causal ontological models, independently of the validity of quantum theory. Only inequalities of this type are termed "Bell inequalities" in modern usage (so that the inequality in Bell's 1964 paper is not a "Bell inequality").

Similarly, Eq. (2.17) is not a proper noncontextuality inequality because it relies upon the idealization of perfect correlations between which of the states $|\phi\rangle$ or $|\bar{\phi}\rangle$ was prepared and which of the outcomes will occur in the measurement of the B_ϕ basis (and similarly for ψ and $\bar{\psi}$). To get a noncontextuality inequality, we must allow these correlations to be imperfect. Thus, in Table 2.1, the entries that take the values 0 and 1 must instead be presumed to take the values ϵ and $1-\epsilon$ respectively, such that ϵ becomes a parameter in our noncontextuality inequality which quantifies the degree of imperfection of the correlations. We then show that quantum mechanics still allows higher success rates for a given confusability than any noncontextual model, even when $\epsilon \neq 0$.

Before proving this, we first rephrase the scenario as a totally operational prepare-and-measure experiment, with no reference to the quantum formalism (despite the suggestive notation below). This is a necessary first step for deriving any proper noncontextuality inequality.

2.5.1 Operationalizing MESD

We imagine an experiment involving four preparations $\{P_\phi, P_\psi, P_{\bar{\phi}}, P_{\bar{\psi}}\}$ and three binary-outcome measurements, $\{M_\phi, M_\psi, M_d\}$, with outcome sets denoted $\{\phi, \bar{\phi}\}$, $\{\psi, \bar{\psi}\}$, and $\{g_\phi, g_\psi\}$, respectively. An arbitrary data table for such an experiment would contain 12 independent parameters, specifying the probability of the first outcome of each measurement when acting on each preparation (the probability of obtaining the second outcome being fixed by normalization).

However, we wish to study the scenario in which preparations P_ϕ, P_ψ, $P_{\bar{\phi}}$, and $P_{\bar{\psi}}$ satisfy the following relation: the procedure $P_{\frac{1}{2}\phi+\frac{1}{2}\bar{\phi}}$ defined by sampling from preparations P_ϕ and $P_{\bar{\phi}}$ uniformly at random (and then forgetting which preparation occurred) is indistinguishable from the similarly defined procedure $P_{\frac{1}{2}\psi+\frac{1}{2}\bar{\psi}}$. We denote this operational equivalence by

$$P_{\frac{1}{2}\phi+\frac{1}{2}\bar{\phi}} \simeq P_{\frac{1}{2}\psi+\frac{1}{2}\bar{\psi}}. \tag{2.29}$$

This implies that only 3 of the parameters in each row are independent, so only 9 independent parameters remain.

Previously the operational equivalence of Eq. (2.29) was guaranteed by quantum theory (Eq. (2.12)), but now we wish to justify it experimentally. In order to do so, one must show that the statistics for $P_{\frac{1}{2}\phi+\frac{1}{2}\bar{\phi}}$ and for $P_{\frac{1}{2}\psi+\frac{1}{2}\bar{\psi}}$ are identical for all measurements. Because the statistics of a tomographically complete set of measurements allows one to predict the statistics for *all* measurements, it suffices to verify this identity for such a tomographically complete set. Accumulating evidence that a given set of measurements is indeed tomographically complete represents the most difficult challenge for an experimental test of noncontextuality (See Refs. [202, 203] for a more detailed discussion.).

Note that in a realistic experiment, the four preparations that are realized, called the primary preparations, will not satisfy Eq. (2.29) perfectly. However, this problem can be solved by post-processing these into "secondary preparations" that are chosen to enforce this equivalence [202, 229], as discussed in Section 2.7.

For this 9-parameter problem, the algorithm we describe in Appendix A.2 gives the full set of necessary and sufficient noncontextuality inequalities, which we list in Appendix A.4. For now, however, we consider a special case with just three parameters, which captures the essence of minimum error state discrimination. Namely, we assume symmetries that parallel those in the ideal quantum case:

$$\begin{aligned} s \equiv p(g_\phi|M_d, P_\phi) &= 1 - p(g_\phi|M_d, P_\psi) \\ = p(g_{\bar{\phi}}|M_d, P_{\bar{\phi}}) &= 1 - p(g_{\bar{\phi}}|M_d, P_{\bar{\psi}}), \end{aligned} \tag{2.30}$$

$$\begin{aligned} c \equiv p(\phi|M_\phi, P_\psi) &= p(\psi|M_\psi, P_\phi), \\ = p(\bar{\phi}|M_{\bar{\phi}}, P_{\bar{\psi}}) &= p(\bar{\psi}|M_{\bar{\psi}}, P_{\bar{\phi}}) \end{aligned} \tag{2.31}$$

and

$$\begin{aligned} 1 - \epsilon \equiv p(\psi|M_\psi, P_\psi) &= p(\phi|M_\phi, P_\phi) \\ = p(\bar{\psi}|M_{\bar{\psi}}, P_{\bar{\psi}}) &= p(\bar{\phi}|M_{\bar{\phi}}, P_{\bar{\phi}}). \end{aligned} \tag{2.32}$$

We have denoted the three free parameters that remain after imposing the symmetries by s, c, and $1-\epsilon$, paralleling their ideal quantum counterparts, s_q, c_q, and 1, respectively. Just like the operational equivalence, these symmetries will never hold exactly for the primary procedures, but we can enforce them while choosing secondary procedures, as discussed in Section 2.7.

The notation $P_\phi, P_\psi, P_{\bar{\phi}}$, $P_{\bar{\psi}}$, M_ϕ, M_ψ, and M_d will henceforth be used to denote the secondary procedures, for which the operational equivalence and symmetries are exact.

The resulting data table, Table 2.2, is similar to the ideal scenario of Table 2.1, but contains the noise parameter ϵ $(1-\epsilon)$ in place of the probability 0 (1).

	P_ϕ	P_ψ	$P_{\bar{\phi}}$	$P_{\bar{\psi}}$
$\phi\|M_\phi$	$1-\epsilon$	c	ϵ	$1-c$
$\psi\|M_\psi$	c	$1-\epsilon$	$1-c$	ϵ
$g_\phi\|M_d$	s	$1-s$	$1-s$	s

Table 2.2: Data table for our operational scenario.

Note that for each row, the average of the entries in the P_ϕ and $P_{\bar{\phi}}$ columns is $\frac{1}{2}$ (and similarly for P_ψ and $P_{\bar{\psi}}$). Here, this follows from the assumed symmetries, not from the operational equivalence (which specifies that the average of the entries for P_ϕ and $P_{\bar{\phi}}$ is the same as the average of the entries for P_ψ and $P_{\bar{\psi}}$, but not necessarily $\frac{1}{2}$); in Table 2.1, the same averaging property is implied by the operational equivalence of each of the two mixtures to the maximally mixed quantum state in Eq. (2.12) (and redundantly implied by these symmetries).

Finally, we assume that the measurements and outcomes are labeled in the natural way; e.g., the outcome of M_ϕ that is more likely to occur given the preparation P_ϕ is ϕ rather than $\bar{\phi}$, etc. Then, the data satisfies the constraint that

$$\epsilon \leq c \leq 1-\epsilon. \tag{2.33}$$

2.6 Noncontextuality inequalities for MESD

The operational equivalence relation of Eq. (2.29) together with the assumption of preparation noncontextuality implies via Eq. (2.7) that

$$\frac{1}{2}\mu_{P_\phi}(\lambda)+\frac{1}{2}\mu_{P_{\bar{\phi}}}(\lambda)=\frac{1}{2}\mu_{P_\psi}(\lambda)+\frac{1}{2}\mu_{P_{\bar{\psi}}}(\lambda), \tag{2.34}$$

where we have again used the fact that convex mixtures of preparations are represented in an ontological model by the corresponding mixture of epistemic states. The fact that the ontological model must reproduce Table 2.2 implies constraints analogous to Eqs. (2.13)-(2.15) and their kin.

As we prove in Appendix A.2, the tradeoff between s, c, and ϵ in any noncontextual model of our operational scenario must satisfy

$$s \leq 1-\frac{c-\epsilon}{2}. \tag{2.35}$$

In Appendix A.3, we show that quantum theory allows a tradeoff of

$$s=\frac{1}{2}(1+\sqrt{1-\epsilon+2\sqrt{\epsilon(1-\epsilon)c(1-c)}+c(2\epsilon-1)}). \tag{2.36}$$

Thus quantum theory predicts a higher state discrimination success rate for any given c and ϵ than a noncontextual model allows. One easily verifies that Eq. (2.35) reduces to Eq. (2.17) in the limit of $\epsilon\to 0$, and that Eq. (2.36) reduces to Eq. (2.11) in the same limit. It is an open question whether Eq. (2.36) is the optimal tradeoff that quantum theory allows. We conjecture that it is optimal for pairs of states in a 2-dimensional Hilbert space.

The noncontextual and quantum tradeoffs are shown in Fig. 2.5. The purple surface represents the triples (s, c, ϵ) saturating the inequality of Eq. (2.35), while the light blue surface represents the triples (s, c, ϵ) corresponding to the quantum success rate of Eq. (2.36).

If an experiment generates data having the form of Table 2.2 and satisfying Eq. (2.29), and it is found to lie above the purple shaded surface, then one has experimental evidence for the failure of noncontextuality. This evidence is independent of the validity of quantum theory, and signals a contextual advantage for state discrimination, even when one's preparations and measurements are imperfect.

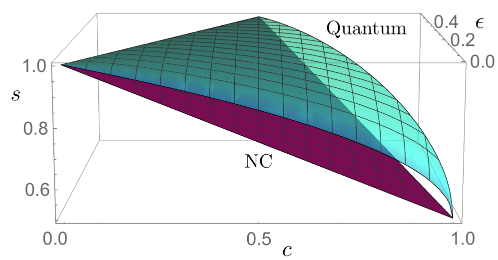

Figure 2.5: Maximum success rate achievable in a noncontextual model (purple surface), and quantumly-acheivable success rate (light blue surface).

2.6.1 Understanding the quantum and noncontextual bounds

For both quantum and noncontextual models, we adopt the natural labeling convention described above Eq. (2.33), so that all operational data necessarily satisfies $\epsilon \le c \le 1-\epsilon$. In the $c-\epsilon$ plane of Fig. 2.5, these constraints describe a triangular wedge that points into the page.

In the plane with $\epsilon = 0$, Section A.1 provides an intuitive explanation for the tradeoff relation.

In the plane with $\epsilon = c$, we can see that for both quantum and noncontextual models, the preparations can be perfectly distinguishable, $s = 1$. This follows from the fact that the value of ϵ quantifies the noise in M_ϕ and M_ψ, and when c is no larger than ϵ we can attribute *all* of the confusability to this noise. Explicitly, one can construct a quantum model where preparation P_ϕ is represented by $|0\rangle\langle 0|$ and P_ψ is represented by $|1\rangle\langle 1|$ and where effect $E_{\phi|M_\phi}$ is represented by $(1-\epsilon)|0\rangle\langle 0| + \epsilon|1\rangle\langle 1|$ and $E_{\psi|M_\psi}$ is represented by $\epsilon|0\rangle\langle 0| + (1-\epsilon)|1\rangle\langle 1|$, which implies that $c = \epsilon$, while $s = 1$ for the Helstrom measurement

$\{|0\rangle\langle 0|, |1\rangle\langle 1|\}$. Furthermore, since these states and effects are all diagonal in the same basis, we can take the eigenvalues of these to define the conditional probabilities of a noncontextual model which achieves $c = \epsilon$ and $s = 1$.

Whenever $c > \epsilon$, however, the noise in M_ϕ and M_ψ cannot explain all of the confusability, and therefore some of this confusability must be explained by the lack of perfect distinguishability of the preparations; that is, in a quantum model, the preparations must be represented by nonorthogonal states, while in a noncontextual model, they must be represented by overlapping probability distributions. Thus, the maximum value of s falls away from 1 as we move away from the $\epsilon = c$ plane. In a noncontextual model, it falls off linearly, interpolating between its value for $\epsilon = c$ and its value for $\epsilon = 0$. The quantum bound falls off more slowly.

2.6.2 Robustness to depolarizing noise

We can get a sense for the robustness of our noncontextuality inequalities by considering a specific noise model in quantum theory. Imagine that one's attempts to implement the ideal quantum preparations and measurements are thwarted by a depolarizing channel which has the same noise parameter v for all states and effects:

$$\mathcal{D}_v(\rho) = (1-v)\rho + v\frac{\mathbb{1}}{2} \tag{2.37}$$

$$\mathcal{D}_v(E_k) = (1-v)E_k + v\frac{\mathbb{1}}{2}. \tag{2.38}$$

The resulting states and effects are shown in Fig. 2.6 for some fixed v. One can graphically see that this uniform depolarization map generates a new set of states and measurements that satisfy the symmetries and operational equivalence we require. However, if the noise is too large, our noncontextuality inequality will not be violated, as we now show.

This noisy model generates a data table of the form of Table 2.2 with

$$s = \frac{1}{2} + (1-v)^2((\frac{1}{2}(1+\sqrt{1-c_q})) - \frac{1}{2}), \tag{2.39}$$

$$c = \frac{1}{2} + (1-v)^2(c_q - \frac{1}{2}), \tag{2.40}$$

$$\epsilon = \frac{1}{2}(1-(1-v)^2). \tag{2.41}$$

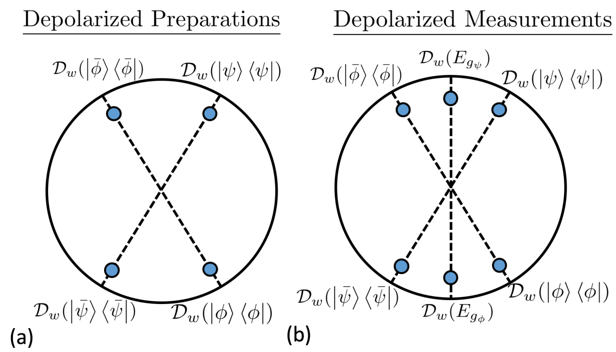

Figure 2.6: The images of the ideal quantum states and effects under a depolarization map for some fixed value of v.

As always, $c_q = |\langle\phi|\psi\rangle|^2$.

The maximum level of noise v that still violates our noncontextuality inequality, Eq. (2.35), is easily calculated as a function of the Bloch sphere angle θ between the two states (defined by $\cos^2(\frac{\theta}{2}) = |\langle\phi|\psi\rangle|^2$), by substituting Eqs. (2.39)-(2.41) into Eq. (2.35):

$$v = 1 - \frac{1}{c_q + \sqrt{1-c_q}} = 1 - \frac{1}{\cos^2(\frac{\theta}{2}) + \sin(\frac{\theta}{2})}. \tag{2.42}$$

Eq. (2.42) is plotted in Fig. 2.7. For $\theta = 0$ or $\theta = \pi$, the noncontextual bound equals the ideal quantum bound, and hence no experiment can violate our noncontextuality inequality at these extremal angles. For all other θ, an experiment with depolarizing noise such that $v \leq 1 - \frac{1}{\cos^2(\frac{\theta}{2})+\sin(\frac{\theta}{2})}$ can violate the inequality. The maximum tolerance to noise ($v = 0.2$) occurs when $\theta = \frac{\pi}{3}$.

2.7 Enforcing symmetries and operational equivalences

In Section 2.5.1, we predicated our noncontextuality inequalities on the exact operational equivalence of Eq. (2.29) and exact operational symmetries of Eq. (2.30)-(2.32), yet we claimed that these idealizations *can* in fact be realized in realistic, noisy experiments. Of course, no experimental data will *directly* satisfy either of these requirements; rather, one performs a post-processing of the data, as originally outlined in [202].

For pedagogical clarity, we will discuss this data processing under the assumption that the operational theory is quantum theory. Note, however, that our comments can easily be

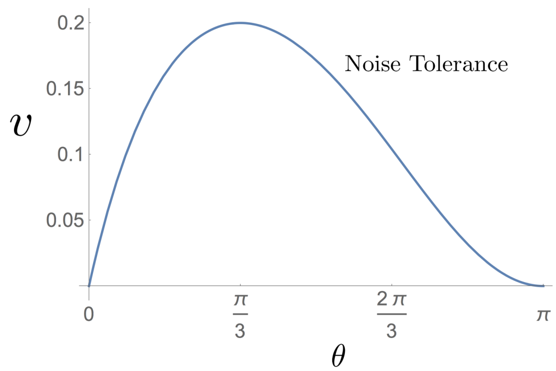

Figure 2.7: The maximum value of the parameter v for the depolarizing noise model that allows a violation of our noncontextuality inequality, as a function of the Bloch sphere angle θ between the two states.

generalized to the framework of generalized probabilistic theories (defined in Refs. [23, 136]), as demonstrated in Refs. [202] and [203]. Indeed, the analysis *must* be performed in this framework if one hopes to directly test the hypothesis of noncontextuality against one's experimental data (i.e., without assuming the validity of quantum theory).

For any set $\mathfrak{P}$ of noisy preparations that has been performed experimentally, one can simulate perfectly the statistics of all other preparations in the convex hull of $\mathfrak{P}$, viewed as points in the quantum state space (here, simply a plane of the Bloch sphere). Similarly, for any set $\mathfrak{E}$ of noisy measurement effects, one can perfectly simulate the statistics of all other effects in the convex hull of $\mathfrak{E}$, viewed as points in the space of valid quantum effects. In [202], this fact was leveraged to simulate exact operational equivalences for a set of "secondary preparations" from data on a set of "primary preparations" that failed to satisfy the operational equivalences exactly. Here, we leverage this trick to simulate preparations and measurements that simultaneously satisfy our operational equivalence *as well as* the symmetries. We now argue that this can always be done, although if the primary preparations or measurements are too noisy, the resulting simulated data will not violate our inequalities.

As we showed explicitly in Section 2.6.2, even a partially depolarized set of states and effects can violate our inequality. Hence, one need only realize experimental sets $\mathfrak{P}$ and $\mathfrak{E}$ that contain in their convex hull the images of our ideal states and effects under the depolarization map $\mathcal{D}_v$ with $v \leq 1 - \frac{1}{\cos^2(\frac{\theta}{2})+\sin(\frac{\theta}{2})}$. Then, one can post-process the data obtained from $\mathfrak{P}$ and $\mathfrak{E}$ to obtain a physically meaningful set of data that satisfies the operational equivalence and symmetries that we assumed in the main text, and our inequality will still be violated. Geometrically, this simply means that the primary preparations must have a convex hull containing the image of the ideal states under a depolarizing map with $v \leq 1 - \frac{1}{\cos^2(\frac{\theta}{2})+\sin(\frac{\theta}{2})}$, as pictured in Fig. 2.8 (and similarly for the measurements, also pictured).

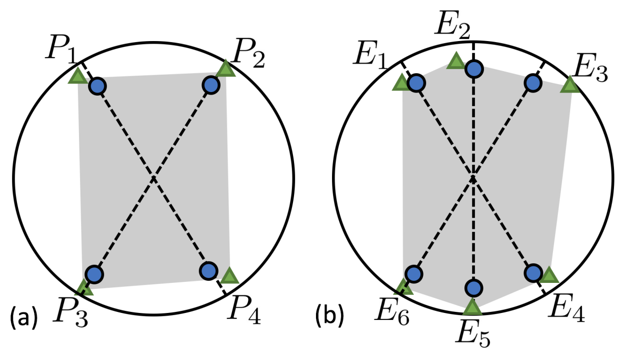

Figure 2.8: (a) If one can perform the four primary preparations P_1 to P_4 (shown as green triangles), then one can simulate any preparation in their convex hull (shown as a light grey shaded region). In particular, one can simulate secondary procedures that are depolarized versions of the ideal preparations (shown as blue circles like those in Fig. 2.6). (b) Similarly for the measurements.

In fact, there are other noisy sets of preparations and measurements besides the depolarized versions of the corresponding ideals which satisfy the operational equivalence and symmetries needed for the noncontextuality inequality to apply. A simple example is states and measurements that are depolarized versions of the ideals that are also rotated in the plane by the same angle. By doing such a rotation, one may be able to simulate a set of states and effects with less depolarization, which then leads to larger violations. In general, there are many sets of states and effects that satisfy our operational equivalence and symmetries. Given a set of primary procedures that one has performed and characterized, finding the states and measurements satisfying our constraints thatmaximize the violation of our inequality is a straightforward linear program [202].

Leveraging the convex structure of operational theories in order to define secondary laboratory procedures which respect certain theoretical idealizations is a powerful tool that we expect to have broad applicability. To date, this method has been proposed to identify operational procedures which respect exact operational equivalences. What we have just shown is that the method also allows one to enforce natural symmetries which greatly simplify the problem at hand (as evidenced by comparing Eq. (2.35) to the set of inequalities in Appendix A.4). Of course, this tool does not allow one to define laboratory procedures that satisfy *any* desired idealizations; for example, one could never generate a pure state or a sharp measurement effect by convexly mixing the noisy procedures actually performed in the lab. We expect future work to continue expanding the range of practical applicability of the technique of secondary procedures.

2.8 Isomorphism between MESD and a Bell scenario

Any noncontextuality scenario that makes no assumptions of measurement noncontextuality, and for which there is a *single* mixed preparation whose various ensemble decompositions generate *all* of the operational equivalences of interest, is isomorphic to a related Bell scenario [184]. Both of these conditions hold for our MESD scenario, since we do not consider any operational equivalences among the measurements, and the operational equivalences among the preparations are generated by decompositions of a single mixed preparation (e.g. the maximally mixed state in the ideal case). The operational Bell scenario related to our MESD scenario is one with two parties, whom we denote by S and M (for reasons that will become apparent), where S has 2 binary measurements, denoted S_1 and S_2, and M has 3 binary measurements, denoted M_1, M_2, and M_3. The outcomes (which we denote s_i for S_i and m_j for M_j) take values in the set $\{-1, +1\}$.

For such a scenario, the set of constraints defining the local set of correlations is given by positivity inequalities, $p(s_i m_j | S_i M_j) \geq 0$, the normalization condition $\sum_{s_i m_j} p(s_i m_j | S_i M_j) \geq 0$, and the CHSH inequalities [75] (applied to any of the 3 possible pairings of 2 measurement settings on S with 2 measurement settings on M) [2]. As we will show, the bound on our MESD success rate follows under our assumed symmetries from the CHSH inequality

$$\langle s_1 m_1 \rangle + \langle s_1 m_3 \rangle + \langle s_2 m_1 \rangle - \langle s_2 m_3 \rangle \leq 2 \tag{2.43}$$

where

$$\begin{aligned} \langle s_i m_j \rangle &= \sum_{s_i m_j} s_i m_j p(s_i m_j | S_i M_j) \\ &= 2p(s_i = m_j | S_i M_j) - 1. \end{aligned} \tag{2.44}$$

The connection between this Bell scenario and our MESD scenario is most easily seen in the ideal quantum realization. Imagine that the two parties share a maximally entangled state $|\Phi^+\rangle_{SM} = \frac{1}{\sqrt{2}}(|00\rangle_{SM} + |11\rangle_{SM})$ (with $|0\rangle$ and $|1\rangle$ defined so that $|\phi\rangle$ and $|\psi\rangle$ have real coefficients when written in this basis), and imagine that their measurements correspond to the quantum measurements from the main text, as follows:

$$\begin{aligned} S_1 &= \{|\phi\rangle_S \langle\phi|, |\bar{\phi}\rangle_S \langle\bar{\phi}|\} \\ S_2 &= \{|\psi\rangle_S \langle\psi|, |\bar{\psi}\rangle_S \langle\bar{\psi}|\} \\ M_1 &= \{|\phi\rangle_M \langle\phi|, |\bar{\phi}\rangle_M \langle\bar{\phi}|\} \\ M_2 &= \{|\psi\rangle_M \langle\psi|, |\bar{\psi}\rangle_M \langle\bar{\psi}|\} \\ M_3 &= \{E_M^{g_\phi}, E_M^{g_\psi}\}. \end{aligned} \tag{2.45}$$

We take the $+1$ outcome for each measurement to correspond to the first quantum effect for that measurement. This ideal quantum realization of this Bell scenario is conceptually transformed into our ideal quantum realization of the MESD scenario by viewing a measurement by party S to be a remote preparation (via quantum steering) for party M. For example, outcome $+1$ for S_1 remotely prepares the state $|\phi\rangle_M$ (which is why we have chosen the notation S, for 'source'). Similarly, outcome -1 for measurement S_2 prepares the state $|\bar{\psi}\rangle_M$, and so on.

Thus, one can verify that in the ideal quantum realization, s_q and c_q become (in our new notation, and assuming the symmetries in Eqs. (2.30)-(2.32))

$$\begin{aligned} s_q &= p(s_1 = m_3|S_1M_3) = 1 - p(s_2 = m_3|S_2M_3) \\ c_q &= p(s_1 = m_2|S_1M_2) = p(s_2 = m_1|S_2M_1), \end{aligned} \tag{2.46}$$

while the fact that paired preparations and measurements are perfectly correlated in the ideal quantum realization corresponds to

$$0 = 1 - p(s_1 = m_1|S_1M_1) = 1 - p(s_2 = m_2|S_2M_2). \tag{2.47}$$

Furthermore, the no-signaling condition in the Bell scenario implies the operational equivalence of our MESD scenario. If party S performs measurement S_1, the updated state on M will be either $|\phi\rangle$ or $|\bar{\phi}\rangle$ with equal likelihood, and if party S performs measurement S_2, the updated state on M will be either $|\psi\rangle$ or $|\bar{\psi}\rangle$ with equal likelihood. In quantum theory, the no-signaling condition implies that the average density operator prepared on M is the same for either choice of measurement by S, which is precisely the operational equivalence of Eq. (2.12).

Using Eq. (2.44), we can write Eq. (2.46) and Eq. (2.47) in terms of expectation values:

$$\begin{aligned} s_q &= \frac{1}{2}(1 + \langle s_1 m_3 \rangle) = \frac{1}{2}(1 - \langle s_2 m_3 \rangle) \\ c_q &= \frac{1}{2}(1 + \langle s_1 m_2 \rangle) = \frac{1}{2}(1 + \langle s_2 m_1 \rangle) \\ 0 &= \frac{1}{2}(1 - \langle s_1 m_1 \rangle) = \frac{1}{2}(1 - \langle s_2 m_2 \rangle). \end{aligned} \tag{2.48}$$

Rewriting Eq. (2.43) in terms of s_q and c_q instead of expectation values, one obtains

$$s_q \leq 1 - \frac{c_q}{2}, \tag{2.49}$$

recovering Eq. (2.17), our bound for the success rate in state discrimination.

Because both the Bell scenario and our MESD scenario are operationally defined, one can also make the translation without assuming the ideal quantum realizations. In a realistic operational scenario, ϵ will be nonzero, and one obtains

$$\begin{aligned} &s = p(s_1 = m_3 | S_1 M_3) = 1 - p(s_2 = m_3 | S_2 M_3) \\ &c = p(s_1 = m_2 | S_1 M_2) = p(s_2 = m_1 | S_2 M_1) \\ &\epsilon = 1 - p(s_1 = m_1 | S_1 M_1) = 1 - p(s_2 = m_2 | S_2 M_2). \end{aligned} \tag{2.50}$$

Rewriting Eq. (2.43) in terms of s, c, and ϵ instead of expectation values, one obtains

$$s \leq 1 - \frac{c - \epsilon}{2}, \tag{2.51}$$

recovering Eq. (2.35), our bound for the success rate in state discrimination.

Due to the redundancies induced by our assumed symmetries, Eq. (2.35) follows also from the CHSH inequality

$$\langle s_1 m_2 \rangle + \langle s_1 m_3 \rangle + \langle s_2 m_2 \rangle - \langle s_2 m_3 \rangle \leq 2, \tag{2.52}$$

by the same logic. More generally, if we do not assume any symmetries, then there are no redundant inequalities. If we furthermore do not assume the natural labeling constraint (Eq. (2.33)), then the full polytope of local correlations for this Bell scenario [2] (and described just above Eq. (2.43)) is isomorphic to the full polytope of noncontextual correlations for our MESD scenario.

2.9 Future directions

We have identified a quantitative feature of minimum-error state discrimination in quantum theory that fails to admit of a noncontextual model. We have derived noncontextuality inequalities that delimit the tradeoff between success rate, error rate, and confusability in state discrimination, independently of the validity of quantum theory.

Our results show that contextuality is a resource for state discrimination, even in realistic, noisy experiments. This suggests many directions for future research. One important question is how our results translate into advantages for quantum information processing tasks which have MESD as a sub-routine. Because many such tasks (e.g., key distribution) consider consecutive measurements on the system, this research program would require further analysis regarding the consequences of noncontextuality for experiments involving sequential measurements [181, 228, 231].

It would also be interesting to generalize these results to other types of state discrimination, such as unambiguous state discrimination. Indeed, one can easily derive a relevant no-go theorem. The challenge is to define an operational notion of "unambiguous" given that no measurement yields truly unambiguous knowledge in the presence of noise. Once this challenge is met, it should be straightforward to apply the general algorithm we have introduced in this article in order to derive the noncontextuality inequalities for this scenario. Understanding the relation between noncontextuality and other kinds of state discrimination should translate into new kinds of quantum advantages for information processing tasks.

Chapter 3

The stabilizer subtheory has a unique noncontextual model

Abstract: We prove that there is a unique nonnegative and diagram-preserving quasiprobability representation of the stabilizer subtheory in odd dimensions, namely Gross' discrete Wigner function. This representation is equivalent to Spekkens' epistemically restricted toy theory, which is consequently singled out as the unique noncontextual ontological model for the stabilizer subtheory. Strikingly, the principle of noncontextuality is powerful enough (at least in this setting) to single out one particular classical realist interpretation. Our result explains the practical utility of Gross' representation, eg why (in the setting of the stabilizer subtheory) negativity in this particular representation implies generalized contextuality, and hence sheds light on why negativity of this particular representation is a resource for quantum computational speedup.

3.1 Introduction

Quantum computers have the potential to outperform classical computers at many tasks. One of the major outstanding problems in quantum computing is to understand what physical resources are necessary and sufficient for universal quantum computation. These resources may depend on one's model of computation [93, 161, 312], and in some cases it seems that neither entanglement nor even coherence is required in significant quantities [93].

The primary obstacle to building a quantum computer is that implementing low-noise gates is difficult in practice. While there are no gate sets that are easy to implement and also universal [99], the entire stabilizer subtheory [130, 131] can in fact be implemented in a fault-tolerant manner relatively easily. To promote the stabilizer subtheory to universal quantum computation, one must supplement it with additional nonstabilizer (or 'magic') processes. Because these nonstabilizer resources do not have a straightforward fault-tolerant implementation, they are typically noisy. To get around this problem, Bravyi and Kitaev [43] introduced the magic state distillation scheme, whereby fault-tolerant stabilizer operations are used to distill pure resource states out of the initially noisy resources. However, not every nonstabilizer resource can be distilled in this fashion to generate a state that promotes the stabilizer subtheory to universal quantum computation. It is a major open question to determine which states are in fact necessary and sufficient for this purpose.

Quasiprobability representations are a central tool for making progress on these and related problems. For finite-dimensional quantum systems, a number of quasiprobability representations have been studied. For example, Gibbons, Hoffman, and Wootters (GHW) identified a family of representations on a discrete phase space [123], and Gross then singled out one of these with a higher degree of symmetry [133], by virtue of satisfying a property known as Clifford covariance. All of these have been used to study quantum computation [95, 118, 154, 194, 235, 304, 309, 310].

Gross's representation in particular has been the most useful in understanding the resources required for computation. For instance, Ref. [309] extended the Gottesman-Knill theorem [131] by devising an explicit simulation protocol for quantum circuits composed of Clifford gates supplemented with arbitrary states and measurements that have nonnegative Gross's representation. Ref. [309] also proved that every state that is useful for magic state distillation necessarily has negativity in its Gross's representation. In Ref. [154], this result was leveraged to prove that every state that promotes the stabilizer subtheory to universal quantum computation via magic state distillation must also exhibit Kochen-Specker contextuality [162]. In recognition that negativity in Gross's representation is a resource for quantum computation in this sense, Ref. [310] introduced an entire resource theory [81] of Gross's negativity.

From a foundational perspective, it is surprising that any *particular* quasiprobability representation plays such a central role. As argued in Ref. [283], negativity of any one quasiprobability representation is not sufficient to establish nonclassicality in general scenarios. So how can it be that Gross's representation plays such an important role, e.g. that negativity in it is associated with a strong form of nonclassicality, namely computational speedups? Although Gross's representation is uniquely singled out from the family of GHW representations by Clifford covariance, it has previously been unclear what this property has to do with nonclassicality (not to mention why one would restrict one's attention to the family of GHW representations in the first place).

In this chapter, we resolve this mystery by showing that the *only* nonnegative and diagram-preserving [260] quasiprobability representation of the stabilizer subtheory in any odd dimension is Gross's. We also prove that in all even dimensions, there is *no* nonnegative and diagram-preserving quasiprobability representation of the stabilizer subtheory. This implies that the stabilizer subtheory exhibits generalized contextuality in all even dimensions.

In the setting of the full stabilizer subtheory, our result for odd dimensions proves that negativity *of this particular* quasiprobability representation is a rigorous signature of nonclassicality, i.e., the failure of generalized noncontextuality. *Generalized noncontextuality* is a principled [281, 285], useful [14, 60, 144, 168, 192, 228, 234, 248, 249, 263, 286, 325], and operational [166, 169, 202, 264] notion of classicality. If one's process has negativity in Gross's representation, then our result establishes that there is no nonnegative representation of the full stabilizer subtheory together with that process. Since nonnegative quasiprobability representations are in one-to-one correspondence with generalized noncontextual ontological models [260, 262, 283], this means that there is no noncontextual representation for the scenario, and hence no classical explanation of it.[1]

Given the known links between resources for quantum computation and negativity in Gross's representation, together with our result connecting such negativity to the failure of generalized noncontextuality, one can then derive connections between resources for quantum computation and generalized noncontextuality.

We illustrate this by proving two such results. First, we give an analogue of the celebrated result in Ref. [154]: namely, we prove that generalized contextuality is necessary for universal quantum computation. Second, we prove that a sufficient condition for any unitary to promote the stabilizer subtheory to universal quantum computation is that it have negativity in Gross's representation. This is in analogy with the fact that a sufficient condition for any pure state to promote the stabilizer subtheory to universal quantum computation via magic state distillation is that it have negativity in Gross's representation [15, 309].

Finally, we note that our main result demonstrates that the principle of generalized noncontextuality is a much stronger principle than was previously recognized, at least in some settings. This is exemplified by the fact that for stabilizer theories in odd dimensions, it does not merely provide constraints on ontological representations, it *completely fixes* the ontological representation. This offers some hope that if the notion of a generalized noncontextual model can be relaxed in such a way [261] that lifts the obstructions to

[1] Note that Ref. [261] introduced a more refined framework for studying ontological models and noncontextuality, and argued that better terminology for these are 'classical realist representations' and 'Leibnizianity', respectively. We do not use this framework or terminology here only so that our results are easier to parse for readers who have not read Ref. [261].

modelling the entirety of quantum theory, such a model of the full theory might also be unique. In our view such a uniqueness result would offer a compelling reason to take the identified ontology seriously.

3.2 The stabilizer subtheory

The stabilizer subtheory is one of the most important subtheories of quantum theory in the field of quantum information, playing an important role in quantum computing [3, 43, 130, 131, 154, 313], quantum error correction [130, 131, 225, 240, 298], and quantum foundations [52, 53, 185, 186, 226, 284].

The stabilizer subtheory is built around the Clifford group, whose elements will be referred to as Clifford unitaries. To define these, we first introduce the *Weyl operators* (also called generalized Pauli operators). Consider a d-dimensional quantum system, and define the computational basis $\{|0\rangle, \ldots, |d-1\rangle\}$ in its Hilbert space $\mathcal{H}$. Each basis element is labelled by an element of $\mathbb{Z}_d$ [2], which we refer to as the configuration space. Writing $\omega = \exp(\frac{2\pi i}{d})$, we define the translation operator X and boost operator Z via

$$X|x\rangle = |x+1\rangle \tag{3.1}$$
$$Z|x\rangle = \omega^x |x\rangle . \tag{3.2}$$

Note that here and throughout, all arithmetic is within $\mathbb{Z}_d$. These can be viewed as discrete position and momentum translation operators, respectively, for a particle on a ring. From these, the single-system Weyl operators are defined as

$$W_{p,q} = Z^p X^q, \tag{3.3}$$

where $p, q \in \mathbb{Z}_d$. Note that these are often defined with an additional phase factor $\omega^{\gamma_{pq}}$; however, the choice of this phase is irrelevant for the definition of the stabilizer subtheory, so we set γ_{pq} to zero. (We highlight this irrelevance by introducing the stabilizer subtheory using superoperators, for which any choice of phase cancels.)

The Weyl operators are unitaries whose associated superoperators, $\mathcal{W}_{p,q}(\cdot) := W_{p,q}(\cdot)W^\dagger_{p,q}$, form a group with composition law

$$\mathcal{W}_{p,q}\mathcal{W}_{p',q'} = \mathcal{W}_{p+p',q+q'}, \tag{3.4}$$

[2] When d is prime, $\mathbb{Z}_d$ has the structure of a finite algebraic field. For non-prime d, things are somewhat more complicated [133], but the results in this work still hold.

and inverse

$$\mathcal{W}_{p,q}^{-1} = \mathcal{W}_{p,q}^{\dagger} = \mathcal{W}_{-p,-q}. \tag{3.5}$$

(Note that the Weyl operators themselves do not form a group as the above equations only hold up to a particular phase factor.)

The Clifford unitaries are defined as unitaries which—up to a phase—map Weyl operators to other Weyl operators under conjugation. Equivalently, their associated superoperators map Weyl superoperators to other Weyl superoperators under conjugation. That is, $\mathcal{U}$ is a Clifford unitary superoperator if for every p, q, one has

$$\mathcal{U}\mathcal{W}_{p,q}\mathcal{U}^{\dagger} = \mathcal{W}_{p',q'}. \tag{3.6}$$

Let us now define the *phase space* $V := \mathbb{Z}_d \times \mathbb{Z}_d$, which is a module[3] equipped with the symplectic product $[\cdot,\cdot] : V \times V \to \mathbb{Z}_d$ given by

$$\left[\begin{pmatrix} p \\ q \end{pmatrix}, \begin{pmatrix} p' \\ q' \end{pmatrix}\right] := pq' - qp'. \tag{3.7}$$

Note that each Weyl operator is labeled by a phase space point $(p, q) = a \in V$. A function $f : V \to V$ is said to be linear if $f(\lambda a + b) = \lambda f(a) + f(b)$, for $\lambda \in \mathbb{Z}_d$, $a, b \in V$. A function $S : V \to V$ is called symplectic if it is linear and preserves the symplectic product, i.e. $[S\cdot, S\cdot] = [\cdot,\cdot]$. A transformation of the form $S \cdot +a$ where S is symplectic and $a \in V$ is called a symplectic affine transformation. Note that the symplectic functions form a group, and that the symplectic affine transformations also form a group.

As shown in Ref. [133], every Clifford superoperator is of the form $\mathcal{W}_a\mathcal{M}_S$, where $\mathcal{W}_a$ is a Weyl superoperator labelled by $a \in V$, $S : V \to V$ is a symplectic function, $\mathcal{M}$ is a unitary superoperator representation of the symplectic group (i.e. $\mathcal{M}_S\mathcal{M}_T = \mathcal{M}_{ST}$), and where $\mathcal{M}_S\mathcal{W}_v\mathcal{M}_S^{\dagger} = \mathcal{W}_{Sv}$ for any symplectic function S and for all $v \in V$.

Hence, each Clifford operation can be indexed by a phase space vector a and a symplectic map S, and so we will denote them by $\mathcal{C}_{a,S} := \mathcal{W}_a\mathcal{M}_S$. Clearly, a Weyl operator $\mathcal{W}_{p,q}$ is a Clifford unitary $\mathcal{C}_{a,S}$, where $a = (p, q)$ and $S = \mathbb{1}$. Furthermore, the mapping $S \cdot +a \mapsto \mathcal{W}_a\mathcal{M}_S$ is a representation of the group of symplectic affine transformations [133].

The Clifford unitary superoperators form a group, often termed the Clifford group, with composition rule

$$\mathcal{C}_{a,S}\mathcal{C}_{b,T} = \mathcal{C}_{Sb+a,ST}. \tag{3.8}$$

[3] If d is a prime power $d = p^k$, then this is moreover a finite vector space.

The inverse of a Clifford unitary superoperator is

$$\mathcal{C}_{a,S}^{-1} = \mathcal{C}_{a,S}^{\dagger} = \mathcal{C}_{-S^{-1}a,S^{-1}}. \tag{3.9}$$

It is therefore clear that the Clifford superoperator group in dimension d and the symplectic affine group for $\mathbb{Z}_d \times \mathbb{Z}_d$ are isomorphic groups.

For a fixed dimension, the Clifford group is generated by the superoperators associated with the generalized Hadamard gate H and the generalized phase gate P [102], defined respectively by

$$H\,|x\rangle = \frac{1}{\sqrt{d}} \sum_{k\in\mathbb{Z}_d} \omega^{xk}\,|k\rangle\,, \tag{3.10}$$

$$P\,|x\rangle = \omega^{\frac{1}{2}x(x+d)}\,|x\rangle\,. \tag{3.11}$$

The stabilizer subtheory for a single system in dimension d is defined as the set of processes that can be generated by sequential composition of: i) pure states uniquely identified by being the simultaneous eigenstates of a given set of Weyl operators, ii) projective measurements in the spectral decomposition of the Weyl operators[4], and iii) Clifford unitary superoperators on the associated Hilbert space, as well as convex mixtures of such processes.

This construction is easily generalized to allow for parallel composition, that is, for systems made up of n qu*d*its[5], by defining the multiparticle Weyl operators as tensor products of those defined above, and defining the multiparticle Clifford operators as unitary superoperators that preserve the multiparticle Weyl operators under conjugation; see Ref. [133] for more details. An important feature is that in general the stabilizer subtheory defined by parallel composition of n qu*d*its is not the same as the stabilizer subtheory defined by a single d^n dimensional system—for instance, the latter generally has far fewer states [133]. Therefore, for a given dimension D there may be multiple different stabilizer theories that could be associated with it, depending on whether one views it as a single monolithic system of dimension D (which Gross calls the single-particle view), or views it as some tensor product of multiple qu*d*its (which Gross calls a multi-particle view).

[4] Note that although the Weyl operators are not Hermitian operators, they *are* normal operators, and hence have a spectral decomposition, which implies one can carry out a projective measurement in the eigenbasis of each.

[5] To the authors' knowledge, parallel composition of systems of different dimensions has not been considered in the literature.

3.3 Quasiprobability representations

A *quasiprobability representation* [103, 260] is akin to a mathematical representation of quantum[6] processes as stochastic processes on a sample space, except that the representation may take negative values. For the reasons laid out in Refs. [260, 261], we are only interested in quasiprobability representations that satisfy the assumption of *diagram preservation* [260, 261]—namely, that the representation of a composite process is equal to the composition of the representations of its component processes. This assumption is satisfied by most of the useful quasiprobability representations considered in the literature, including the standard (continuous-dimensional) Wigner function and Gross's representation.

The arguments of Ref. [260] imply that every diagram-preserving quasiprobability representation of a full dimensional subtheory[7] of quantum theory can be written as an exact frame representation [103], constructed as follows. One first associates with each system a basis $\{F_\lambda\}_\lambda$ for the real vector space $\mathsf{Herm}(\mathcal{H})$ of Hermitian operators on the associated Hilbert space $\mathcal{H}$, where

$$\mathsf{tr}[F_\lambda] = 1, . \tag{3.12}$$

Every basis has a unique dual basis, $\{D_\lambda\}_\lambda$, as proved in Lemma B.1, where

$$\sum_\lambda D_\lambda = \mathbb{1}, \tag{3.13}$$

and

$$\mathsf{tr}[D_{\lambda'} F_\lambda] = \delta_{\lambda\lambda'}. \tag{3.14}$$

In this representation, a completely-positive trace-preserving map [212, 257] $\mathcal{E}$ is represented by a quasistochastic map defined by

$$\xi_{\mathcal{E}}(\lambda'|\lambda) = \mathsf{tr}[D_{\lambda'}\mathcal{E}(F_\lambda)]. \tag{3.15}$$

[6] Quasiprobability representations can also be defined for generalized probabilistic theories [23, 136] (GPTs) beyond quantum theory [260, 305], but we are here only interested in the case of quantum theory and its subtheories.

[7] That is, in which the states span the quantum state space and the effects span the quantum effect space. Note that the stabilizer subtheory is such a theory, which can be seen by noting that the Weyl operators span the space of Hermitian operators, and hence, so do their eigenstates.

As special cases, the representation of a state ρ is given by

$$\xi_\rho(\lambda) = \mathsf{tr}[D_\lambda \rho] \tag{3.16}$$

and the representation of an effect E is given by

$$\xi_E(\lambda) = \mathsf{tr}[F_\lambda E]. \tag{3.17}$$

Note that for a set of effects that sum to the identity, Eq. (3.12) ensures that $\sum_E \xi_E(\lambda) = 1$.

The quantum probabilities are recovered as

$$\mathsf{tr}[E\mathcal{E}(\rho)] = \sum_{\lambda',\lambda} \xi_E(\lambda')\xi_{\mathcal{E}}(\lambda'|\lambda)\xi_\rho(\lambda). \tag{3.18}$$

Note that this is an instance of diagram preservation, wherein one decomposes the probability into the composition of the representations of the state, channel, and effect.

A quasiprobability representation is said to be *nonnegative* if for every process $\mathcal{E}$, $0 \leq \xi_{\mathcal{E}}(\lambda'|\lambda) \leq 1$ for every λ, λ'. In this case, the representation is in one-to-one correspondence with a noncontextual ontological representation [261, 281].

3.3.1 Gross's representation

The particular quasiprobability representation introduced by Gross [133] is for odd dimensional quantum systems and takes the sample space to be a discrete classical phase space V, and so its elements will be labelled by position and momentum, i.e $a := (p, q) \in V$, rather than λ. Hence, the basis operators in Gross's representation are indexed by $a \in V$, and we will denote them by A_a. Note that the Weyl operators form an orthonormal basis for the complex vector space of linear operators on the Hilbert space, where orthonormality is with respect to a rescaled Hilbert-Schmidt inner product:

$$\frac{1}{d}\,\mathsf{tr}[W_{p,q}W^\dagger_{p',q'}] = \delta_{p,p'}\delta_{q,q'}. \tag{3.19}$$

The basis operators in Gross's representation can be decomposed in terms of this orthonormal basis as follows:

$$\{A_a\}_a := \left\{\frac{1}{d}\sum_b \exp([a,b])W_b^{G\dagger}\right\}_a, \tag{3.20}$$

where Gross's Weyl operators $W^G_{p,q}$ are related to ours via

$$W^G_{p,q} := \omega^{2^{-1}pq} W_{p,q}. \tag{3.21}$$

We will sometimes refer to these A_a as phase space point operators. These operators form an orthonormal basis for $\mathsf{Herm}(\mathcal{H})$, and so the basis is essentially self-dual, so that both $\{F_\lambda\}$ and $\{D_\lambda\}$ are proportional to $\{A_a\}$, with $D_\lambda = \frac{1}{d}F_\lambda$. They moreover satisfy a number of useful properties (see, e.g., Lemma 29 of Ref. [133]) including

$$A_a = W_a A_0 W_a^\dagger. \tag{3.22}$$

Eq. (3.22) is a special case of a key feature of Gross's representation, namely *Clifford covariance* [133]:

$$C_{a,S} A_b C_{a,S}^\dagger = W_a M_S A_b M_S^\dagger W_a^\dagger = A_{Sb+a} \tag{3.23}$$

for any Clifford unitary $C_{a,S}$. This property implies, for example, that when one transforms a quantum state under a given Clifford unitary, the representation of the state transforms under the associated symplectic affine map, i.e.

$$\xi_\rho(b) = \xi_{C_{a,S}\rho C_{a,S}^\dagger}(Sb+a). \tag{3.24}$$

To see this, note that the representation of the state is given by Eq. (3.16), which in this instance means that $\xi_\rho(b) := \mathsf{tr}[A_b\rho]$. Using this definition and the properties that we have so far introduced, we obtain:

$$\begin{aligned}
\xi_{C_{a,S}\rho C_{a,S}^\dagger}(Sb+a) &:= \mathsf{tr}[A_{Sb+a} C_{a,S}\rho C_{a,S}^\dagger] && (3.25)\\
&= \mathsf{tr}[C_{a,S}^\dagger A_{Sb+a} C_{a,S}\rho] && (3.26)\\
&\overset{(6.269)}{=} \mathsf{tr}[C_{-S^{-1}a,S^{-1}} A_{Sb+a} C^\dagger_{-S^{-1}a,S^{-1}}\rho] && (3.27)\\
&\overset{(3.23)}{=} \mathsf{tr}[A_{S^{-1}(Sb+a)-S^{-1}a}\rho)] && (3.28)\\
&= \mathsf{tr}[A_b\rho] && (3.29)\\
&=: \xi_\rho(b). && (3.30)
\end{aligned}$$

3.4 Main result

Our main result is a complete characterization of the (non)classicality of the stabilizer subtheory in every finite dimension.

Theorem 3.4.1.

(a) For any stabilizer subtheory (single- or multi-particle) in **odd** *dimensions, the* unique *nonnegative and diagram-preserving quasiprobability representation for it is Gross's representation.*

(b) For any stabilizer subtheory (single- or multi-particle) in **even** *dimensions, there is no nonnegative and diagram-preserving quasiprobability representation.*

The proof is given in Appendix B.3. The proof for the single-particle case proceeds as follows. First, we apply the structure theorem from Ref. [260] to show that any nonnegative and diagram-preserving representation of the stabilizer subtheory must be an exact frame representation. Next, we leverage the fact that noncontextuality implies outcome determinism to find a privileged labeling of the ontic states as points in a phase space. We show that this implies Clifford covariance for all Weyl operators and for the Hadamard. Using this and the fact that Weyl operators form a basis of the linear operators, we then show that the representation is fixed by the outcomes of measurements of Weyl operators on the $\lambda = (0, 0)$ ontic state.[8] We then show that the representation is fixed by a specification of which outcome occurs for measurements of Weyl operators when the ontic state is $\lambda = (0, 0)$. We then show that this specification is constrained by the representation of the Hadamard, by Hermiticity of the phase point operators, and by constraints on outcomes assigned to commuting pairs of Weyl operators. In odd dimensions, we show that the unique solution to these conditions is that which gives Gross's phase point operators. In even dimensions, we show that there is no solution. The generalization to multi-particle stabilizer subtheories is then shown to follow immediately.

As shown in Ref. [52, 284], Gross's representation is identical to Spekkens' epistemically restricted toy theory [282] for odd dimensions [284]. Furthermore, it is shown in Ref. [260] that noncontextual ontological models of an operational theory are in one-to-one correspondence with ontological models of the GPT defined by the operational theory, and also in one-to-one correspondence with diagram-preserving and nonnegative quasiprobability

[8] We believe, but have not shown, that distinct GHW representations differ by exactly these choices of outcomes.

representations of the GPT defined by the operational theory. Through these equivalences, our result can be stated in a number of equivalent ways (e.g., depending on whether one views the stabilizer subtheory as an operational theory or as a GPT). Perhaps the most natural equivalent statement of Theorem 18 is the following: For odd dimensions, the unique noncontextual representation of the stabilizer subtheory is Spekkens' epistemically restricted toy theory. For even dimensions, the stabilizer subtheory is contextual.

There are several senses in which Theorem 18(a) is stronger than that proven by Gross [133]. Most importantly, the principle of generalized noncontextuality is a well-established notion of classicality, while the notion of covariance is not. Additionally, our result starts from the very weak assumption of classical realism [261]—that is, the ontological models framework—while Gross's result requires two additional assumptions beyond this, namely that the representation is on a $d \times d$ phase space and gives the correct marginal probabilities. In our approach, both of these are derived. Finally, our uniqueness result holds in all odd dimensions, while Gross's uniqueness result was proven only for odd prime dimensions.

Theorem 18(b) establishes that every stabilizer subtheory of even dimension exhibits contextuality. While this result has previously been claimed to be true, it had not in fact been proven (to our knowledge). For $d = 2$, there are well-known proofs of contextuality, e.g. in Ref. [186]. It follows that every subtheory that contains all the processes in the qubit stabilizer subtheory is also contextual. However, it is not known whether every even-dimensional stabilizer subtheory contains the qubit stabilizer as a subtheory (see Ref. [133] for details), and so the claim of Theorem 18(b) does not trivially follow in this manner.

3.5 Generalized contextuality as a resource for quantum computation

The stabilizer subtheory is efficiently simulable [131]. However, if one supplements it with appropriate nonstabilizer states, one can achieve universal quantum computation through magic state distillation [43].

Any state that promotes the stabilizer subtheory to universal quantum computation must have negativity in its Gross's representation [309]. Ref. [154] further showed that every such state can be used to generate state-dependent proofs of Kochen-Specker contextuality using stabilizer measurements [154], and hence that contextuality is necessary for universality in this model of quantum computation.

The key argument of Ref. [154] was a graph-theoretic proof that if a state is negative in Gross's representation, then it admits a (state-dependent) proof of Kochen-Specker contextuality using only stabilizer measurements. Our main theorem, Theorem 18, is analogous, establishing that if a state is negative in Gross's representation, then it admits a proof of *generalized* contextuality.

Hence, we arrive at a result akin to that of Ref. [154]: generalized contextuality is necessary for universality in the state injection model of quantum computation.

Theorem 3.5.1. *Consider any state ρ that promotes the stabilizer subtheory to universal quantum computation. There is no generalized noncontextual model for the stabilizer subtheory together with ρ.*

This follows immediately from the fact that negativity in a state's Gross's representation is necessary for it to promote the stabilizer subtheory to universal quantum computation [309], together with our result that negativity in Gross's representation implies generalized contextuality.[9]

One might expect that this result follows immediately from the fact that there is no nonnegative quasiprobability representation of full quantum theory, and that such a proof would hold in every model of quantum computation. However, the mere fact that a universal quantum computer can *simulate* every quantum circuit does not necessarily imply that one can *implement* every quantum circuit. (The loophole here follows from the distinction between computational universality and strict universality [10]. For example, the Toffoli and Hadamard gate together form a computationally universal gate set, and yet composition of these two gates cannot generate arbitrary unitary gates—only those with real matrix elements.) Hence, one cannot without further arguments conclude that a universal quantum computer is capable of implementing circuits with negativity (or contextuality)—one can only conclude that it can simulate such circuits.

[9]The result could also presumably be proven less directly, using the necessity result for Kochen-Specker contextuality [154] together with the fact that Kochen-Specker contextuality implies generalized contextuality [167, 170]. This is not entirely trivial, as the latter implication requires bringing auxiliary operational processes into the argument, and one must establish that all of these additional processes are within the stabilizer subtheory. This seems to be the case. First, one can establish outcome determinism for ontic states in the support of the maximally mixed state following the logic of Ref. [281], but using only stabilizer preparations. One can then establish that every ontic state in the support of the given nonstabilizer state (from the state-dependent proof of Ref. [154]) is also in the support of the maximally mixed state, using the fact that there always exists a decomposition of the maximally mixed state into the given nonstabilizer state together with *only* stabilizer states.

3.5.1 On the sufficiency of generalized contextuality for universal quantum computation

Thus far we have focused on the necessity of contextuality for quantum computation. However, the fact that Gross's representation provides the unique generalized noncontextual representation of the stabilizer subtheory will likely also be useful for discovering in what sense (if any) generalized contextuality is *sufficient* for quantum computation.

Without any caveats, generalized contextuality is clearly not sufficient for universal quantum computation. This can be seen by the example of the stabilizer theory in dimension 2, which admits proofs of contextuality [186] and yet is efficiently simulable [131].

Still, it is conceivable that there is a more nuanced sufficiency result relating contextuality and computation, e.g. by leveraging quantitative measures of generalized contextuality [196] or by focusing on particular dimensions and models of quantum computation.

We now prove a related result (which does not explicitly rely on our main theorem).

From Ref. [15, 309], we know that access to enough copies of any nonstabilizer pure state promotes the stabilizer subtheory to universal quantum computation. Similarly, access to enough copies of any nonstabilizer unitary promotes the stabilizer subtheory to universal quantum computation, since the Clifford unitaries together with any other unitary gate forms a universal gate set [51, 211].

It is well known that every pure nonstabilizer state is negatively represented in Gross's representation [133]. Additionally, it is not hard to see that every nonstabilizer unitary gate is negatively represented in Gross's representation. By the universal gate set property [51, 211], combining the positively represented Clifford gates with any given nonstabilizer unitary allows the approximation of any other unitary—including one that maps some pure stabilizer state to some pure nonstabilizer state. Since the stabilizer state is represented positively and the nonstabilizer state must be represented negatively in Gross's representation, the unitary mapping between them must have negativity in its Gross's representation, and hence so must the given nonstabilizer unitary used to construct it. Hence we obtain the following theorem:

Theorem 3.5.2. *A sufficient condition for any unitary or pure state to promote the stabilizer subtheory to universal quantum computation is that it be negatively represented in Gross's representation.*

For the case of pure states, this result was pointed out in Refs. [15, 309].

Perhaps the most important open question that remains is whether an analogous sufficiency result holds for mixed quantum states and generic quantum channels.

Part II.

Nonclassicality of common-cause processes

Chapter 4

Quantifying Bell: the Resource Theory of Nonclassicality of Common-Cause Boxes (excerpt)

4.1 Motivating our approach and contrasting it with alternatives

4.1.1 Three views on Bell's theorem

The traditional commentary on Bell's theorem [96, 274] takes a particular view on how to articulate the assumptions that are necessary to derive Bell inequalities. Among these assumptions, two are typically highlighted as deserving of the most scrutiny, namely, the assumptions that are usually termed *realism* and *locality*[1]. Abandoning one or the other of these two assumptions is the starting point of most commentaries on what to do in the face of violations of Bell inequalities.[2] Furthermore, a schism seems to have developed between the camps that advocate for each of these two views [319].

Among the researchers who take Bell's theorem to demonstrate the need to abandon realism, there is a contingent that adopts a purely operational attitude towards quantum theory, that is, an attitude wherein the scientist's job is merely to predict the statistical distribution of outcomes of measurements performed on specific preparations in a specified

[1] Note, however, that different authors will formalize these assumptions in different ways.

[2] See, however, the discussion of superdeterminism in footnote 6.

experimental scenario. We shall refer to the members of this camp as *operationalists* [316]. For such researchers, a violation of a Bell inequality is simply a litmus test for the inadequacy of a classical realist account of the experiment. One particular type of operationalist attitude, which we shall term the **strictly operational paradigm**, advocates that physical concepts ought to be defined in terms of operational concepts, and consequently that any properties of a Bell-type experiment, such as whether it is signalling or not and what sorts of causal connections might hold between the wings, must be expressed in the language of the classical input-output functionality of that experiment. In other words, they advocate that the only concepts that are meaningful for such an experiment are those that supervene[3] upon its input-output functionality.[4] Most prior work on quantifying the resource in Bell experiments has been done within this paradigm, and the characteristic of experimental correlations that is usually taken to quantify the resource is simply some notion of distance from the set of correlations that satisfy all the Bell inequalities.

Consider, on the other hand, the researchers who take realism as sacrosanct, and in particular those who take Bell's theorem to demonstrate the failure of locality—that is, the existence of superluminal causal influences [199, 213].[5] Researchers in this camp, whom we shall refer to as advocates of the **superluminal causation paradigm**, would presumably find it natural to quantify the resource of Bell inequality violations in terms of the strength of the superluminal causal influences required to account for them (within the framework of a classical causal model). An approach along these lines is described in Refs. [64, 65]. Earlier work on the communication cost of simulating Bell-inequality violations [198, 299] is also naturally understood in this way.[6]

[3] A-properties are said to supervene on B-properties if every A-difference implies a B-difference.

[4] Some might describe what we have here called the strictly operational paradigm as the "device-independent" paradigm [252], however, we avoid using the latter term here because its usage is not restricted to describing a particular type of empiricist philosophy of science: it also has a more technical meaning in the context of quantum information theory, wherein it indicates whether or not a given information-theoretic protocol depends on a prior characterization of the devices used therein. Indeed, Bell-inequality-violating correlations have been shown to be a key resource in cryptography because they allow for device-independent implementations of cryptographic tasks[7, 8, 27, 87, 97, 157, 222, 253, 308].

[5] Although such influences do not imply the possibility of superluminal signalling, they do imply a certain tension with relativity theory if one believes that the latter does not merely concern anthropocentric concepts such as signalling, but also physical concepts such as causation.

[6] A less common view on how to maintain realism in the face of Bell inequality violations is to hold fast to locality but give up on a different assumption that goes into the derivation of Bell inequalities, namely, that the hidden variables are statistically independent of the setting variables. This is known as the "superdeterministic" response to Bell's theorem [147]. Advocates of this approach would presumably find it natural to quantify the resource of Bell inequality violations in terms of the deviation from such statistical independence that is required to explain a given violation. In particular, the results of Refs. [135] and [26] seeking to quantify the nonindependence needed to explain a given Bell inequality violation might

In recent years, a third attitude toward Bell's theorem—inspired by the framework of causal inference [216]—has been gaining in popularity. In this approach, the assumptions that go into the derivation of Bell inequalities are [324]: Reichenbach's principle (that correlations need to be explained causally), the framework of classical causal modelling, and the principle of no fine-tuning (that statistical independences should not be explained by fine-tuning of the values of parameters in the causal model). Here, a violation of a Bell inequality does not lead to the traditional dilemma between realism and locality, but rather attests to the impossibility of providing a non-fine-tuned explanation of the experiment within the framework of classical causal models. This attitude implies the possibility of a new option for what assumption to give up in the face of such a violation. Specifically, the new possibility being contemplated is that one can hold fast to Reichenbach's principle and the principle of no fine-tuning—and hence to the possibility of achieving satisfactory causal explanations of correlations—by replacing the framework of classical causal models with an intrinsically nonclassical generalization thereof.

As is shown in Ref. [324], because the correlations in a Bell experiment do not provide a means of sending superluminal signals between the wings, the only causal structure that is a candidate for explaining these correlations without fine-tuning is one wherein there is a purely common-cause relation between the wings, that is, one that admits no causal influences between the wings. Therefore, the new approach to achieving a causal explanation of Bell inequality violations is one that posits a common cause mechanism but replaces the usual formalism for causal models with one that allows for more general possibilities on how to represent its components [12][7]. We refer to this attitude as the **causal modelling paradigm**.

The causal modelling paradigm implies not only a novel attitude towards Bell's theorem, but also a change in how one conceives of the resource that powers the information-theoretic applications of Bell-inequality violations. The resource is not taken to be some abstract notion of distance from the set of Bell-inequality-satisfying correlations within the space of all nonsignalling correlations, as advocates of the strictly operational paradigm seem to favour, nor to consist of the strength of superluminal causal influences, as advocates of the superluminal causation paradigm would presumably have it. Rather, we take the resource to be the *nonclassicality* required by any generalized causal model that can explain the Bell

be framed within a resource-theoretic framework. However, given that the setting variables can no longer be considered as freely specifiable within such an approach, it would be inappropriate to conceptualize a Bell experiment as a box-type process as we have done here.

[7] Specifically, in the proposal of Ref. [12], reversible deterministic causal dependences are represented by unitaries rather than bijective functions, and lack of knowledge is represented by density operators rather than by classical probability distributions.

inequality violations without fine-tuning.

We shall show that in the resource theory that emerges by adopting this attitude, the nonclassicality of common-cause processes in Bell experiments cannot be captured solely by the degree of violation of facet-defining Bell inequalities. That is, there are distinctions among such common-cause processes—different ways for these to be nonclassical—which do not correspond to distinctions in the degree of violation of any facet-defining Bell inequality.

4.1.2 Generalized causal models

We will work with the notion of a generalized (i.e., not necessarily classical) causal model that has been developed in Refs. [107, 142] using the framework of generalized probabilistic theories (GPTs) [23, 136]), and refer to it as a **GPT causal model**. Since we are interested in the distinction between classical and nonclassical, without specifically distinguishing quantum and supra-quantum types of nonclassicality, we will not be making use of any of the recent work [12, 88] on devising an intrinsically *quantum* notion of a causal model.[8]

One can then approach the study of nonclassicality in arbitrary causal structures from within the scope of these GPT causal models, and pursue the development of a resource theory of such nonclassical features. One must simply specify the nature of the scenario being considered: the number of wings of the experiment (commonly conceptualized as the laboratories of different parties when discussing information-theoretic tasks), and the causal structure presumed to hold among these wings.[9] The set of all resources one might contemplate are then the set of processes that can be described with a GPT causal model having the appropriate causal structure. In this chapter, we focus on the causal structure wherein there is a common cause that acts on all of the wings, but no causal influences between any of them, which we term a **Bell scenario**.

We conceptualize any experimental configuration as a process from its inputs to its outputs. In the GPT framework for causal models, one has the capacity to consider processes that have GPT systems as inputs and outputs at the various wings. However, we will restrict our attention to processes that have only *classical* inputs and outputs.

[8]However, we will consider the question of when certain correlations that arise in a GPT causal model can be quantumly-realized.

[9]It is perhaps inappropriate to call the relation between the parties in a general communication protocol a "causal structure", insofar as the latter term usually refers to a directed acyclic graph (DAG) in the causal inference literature, and a communication network can have cycles, such as when there exists a communication channel in both directions between two parties. Nonetheless, we will here focus on communication networks that *do* correspond to DAGs.

Such processes can be conceptualized as black-box processes, to which one inputs classical variables and from which classical variables are output. They are therefore precisely the sorts of processes considered in the device-independent paradigm. We further restrict our attention to processes with a classical input and classical output at each wing, where the input temporally precedes the output.[10] In the device-independent paradigm, the term "box" is generally used as jargon for such processes (for instance, as it is used in the term "PR box" [224]). We therefore refer to such processes as **box-type processes** or simply **boxes**. A box-type process is completely characterized by specifying the conditional probability distribution over its outcome variables given its input variables.

We use the term **common-cause box** to refer to box-type processes that can be realized using a causal structure consisting of a common cause acting on all of the wings. In GPT causal models, all common-cause boxes can be decomposed into the preparation of a GPT state on a multipartite system, followed by the distribution of the component subsystems among the wings, followed by each subsystem being subjected to a GPT measurement, chosen from a fixed set according to the classical input variable at that wing (the local setting variable), and the result of which is the classical output variable at that wing (the local outcome variable). In short, such processes can be decomposed in the same manner in which a multipartite Bell experiment is decomposed into a preparation of a correlated resource and local measurements. The distinction between classical and nonclassical common-cause boxes is simply the distinction between whether there is a *classical* causal model underlying the process, or whether one must resort to a causal model that invokes a nonclassical GPT.

4.1.3 Resourcefulness in the causal modelling paradigm

In order to quantify the nonclassicality of common-cause boxes, we will use the approach to resource theories described in Ref. [80]. In this approach, resource theories are defined via *partitioned process theories*. An **enveloping theory of processes** must be specified, together with a subtheory of processes that can be implemented at no cost, called the **free subtheory of processes**. This partitions the set of all processes in the enveloping theory into free and costly (i.e., nonfree) processes. One can then ask of any pair of processes in the enveloping theory whether the first can be converted to the second by embedding it in a circuit composed of processes that are drawn entirely from the free subtheory. The set of higher-order processes that are realized in this way—i.e., by embedding in a circuit composed of processes drawn from the free subtheory—is termed the **set of free**

[10]Thus, we do not consider processes that involve a sequence over time of classical input variables and classical output variables; that is, in the language of Refs [67, 68], we do not consider general n-combs.

operations. Pairwise convertibility relations under the set of free operations define a pre-order on the set of all resources, and a partial order over the equivalence classes of such resources. One can then quantify the relative worth of different resources by their relative positions in this partial order. Functions over resources that preserve ordering relations, termed monotones, provide a particularly simple means of quantifying the worth of resources.

The resource theory considered in this chatper is defined as follows. We take the enveloping theory of processes to consist of the common-cause boxes that can be realized in a GPT causal model, which we term **GPT-realizable**. We take the free subtheory of processes to consist of the common-cause boxes that can be realized in a classical causal model, which we term **classically realizable**.

It follows that the free common-cause boxes are precisely those that satisfy all the Bell inequalities, while the costly common-cause boxes are those that violate some Bell inequality. To determine the ordering relations that hold among these common-cause boxes, one must determine the convertibility relations among them. Given the definition of our resource theory, whether one common-cause box can be converted to another is determined by whether this can be achieved by composing it with classical common-cause boxes. This subsumes correlated local processings of the inputs and outputs of the box.

A note about nomenclature

In this chapter, we avoid describing the resource behind Bell inequality violations as *nonlocality*. This is because we believe that it is *only* for those who take the lesson of Bell's theorem to be the existence of superluminal causal influences that it is appropriate to describe violations of Bell inequalities by this term. Researchers in the operationalist camp have not, generally speaking, avoided using the term *nonlocality*, but seem instead to use it as a synonym for "violation of a Bell inequality" rather than to imply a commitment to superluminal causal influences. However, we believe that such a usage invites confusion and so we opt instead to avoid the term altogether. Nevertheless, our project is very much in line with earlier projects that describe themselves as developing a *resource theory of nonlocality*, such as Refs. [94, 116, 117, 120].

4.1.4 Contrast to the strictly operational paradigm

As noted in the introduction and as will be demonstrated, in the special case of Bell scenarios—the focus of this chapter—the natural set of free operations within our causal modelling paradigm is equivalent to one of the proposals for the set of free operations made in earlier works within the strictly operational paradigm, namely, *local operations and shared randomness* (LOSR), as the latter is defined in Refs. [94, 120]. Additionally, the natural enveloping theory adopted in the strictly operational approach, namely, the set of no-signalling boxes, also coincides with that of our enveloping theory for the case of Bell scenarios, namely, the set of GPT-realizable common-cause boxes (where the equivalence of these two sets can be inferred from the results of Ref. [23]). Therefore, in spite of the difference in the attitude we take towards Bell's theorem, the resource theory that we define for Bell scenarios is the same as the one studied in Refs. [94, 120].

Nonetheless, the difference in our attitude towards Bell's theorem is not inconsequential. We presently outline its significance for the project of this chapter as well as for potential future generalizations of this project.

Most importantly, the causal modelling approach diverges sharply from any strictly operational approach once one considers causal structures beyond Bell scenarios. In a resource theory of nonclassicality for more general causal structures, both the free subtheory and the enveloping theory proposed by the causal modelling approach are radically different from those suggested by the strictly operational approach. In particular, the free subtheory need not be LOSR in a general causal structure and the enveloping theory need not be the set of all nonsignalling operations. Our approach allows us to define a resource theory that is specific to a scenario in which only strict subsets of the wings are connected by common causes [40, 107] (such as the triangle-with-settings scenario) and this provides a concrete example of a case where the free subtheory is not LOSR and the enveloping theory is not all nonsignalling operations. In these cases, the free operations are "local operations and causally admissable shared randomness", wherein only those subsets of wings that are connected by a common cause have shared randomness. This is distinct from the LOSR operations, which assume that randomness is shared between all the wings. It seems unlikely that the resource theory we propose in these cases can be motivated (or even fully characterized) in the strictly operational paradigm.

Even for Bell scenarios, however, the causal modelling approach offers advantages over its competitors. In particular, it singles out a unique set of free operations, while the strictly operational approach does not. From our perspective, the resource underlying Bell inequality violations is the nonclassicality of the causal model required to explain them with a common cause, so *clearly* the free operations should involve only classical common causes acting between the wings. In the strictly operational paradigm, by contrast, any operation that preserves no-signalling and takes local boxes to local boxes might constitute a legitimate candidate for a *free* operation. This ambiguity is reflected in the existence of distinct proposals for the set of free operations in strictly operational resource theories. Aside from LOSR, there is also a proposal called *wirings and prior-to-input classical communication* (WPICC) [116] that allows for classical causal influences among the wings *prior* to when the parties receive their inputs. If one believes that there is a singular concept that underlies the violation of Bell inequalities, then at most *one* of these proposals (LOSR or WPICC) can be taken as the relevant set of free operations.[11] Although WPICC operations meet all desired operational criteria, they are immediately ruled out as candidates for the free operations within the causal modelling paradigm, on the grounds that they involve nontrivial cause-effect influences between the wings.

Another advantage of our approach for the Bell scenario is that it highlights the fact that LOSR is *by construction* a convex set, a fact that is critical for the algorithmic method that we derive for determining the ordering relation between any two resources. In highlighting this fact, our approach led us to notice an oversight in some previous attempts to formalize LOSR.

Finally, we note that prior work of Ref. [120] departs from the strictly operational paradigm through their use of the *unified operator formalism* [6, 11], which is analogous to the quantum formalism, but where nonpositive Hermitian operators are allowed to represent states. They do not characterize boxes primarily by their input-output functionality, but rather as a composition of a bipartite source with local measurements. Indeed in their Fig. 4, they explicitly depict the internal structure of the box. It is in this sense that their approach does not quite fit the mould of a strictly operational approach but is rather somewhat more in the flavour of the causal modelling approach we have described here.

Nonetheless, the unified operator formalism differs significantly from the GPT formalism of Refs. [107, 142] with respect to the *independence* of the nonclassical common cause from the measurements employed in realizing nonclassical boxes. In the unified operator formalism, the Hermitian operator describing the shared state cannot be chosen freely for a given set of quantum measurements, because some choices would yield negative numbers rather than valid probabilities. By contrast, in the GPT formalism that we adopt here, the set of GPT states is contained within the dual of the set of GPT product measurements, and hence any measurement scheme can be paired with any shared state while yielding valid probabilities. The causal modelling paradigm must reject any dependence of the shared state on the choice of measurements, while such dependence is unavoidable within the unified operator formalism. As defined in Ref. [216], a causal model is a directed acyclic graph, or equivalently, a circuit of causal processes, wherein the distinct processes in the circuit are required to be *autonomous* (i.e., independently variable). We therefore

[11] Competing sets of free operations may be interesting for studying phenomena *other* than the resource powering violations of Bell inequalities, but this is not the issue at stake in this chapter.

classify Ref. [120] as neither within the causal modelling paradigm nor within the strictly operational paradigm, while still exhibiting some features of each of these approaches.

4.1.5 Contrast to the superluminal causation paradigm

To our knowledge, advocates of the superluminal causation paradigm have not attempted to develop a resource theory for Bell inequality violations (although Refs. [64, 65] are related in spirit). If it *were* attempted (within the framework of Ref. [80]), then the commitments of the approach suggest that it would also be done differently from the way we have done so here. Those who endorse the superluminal causation paradigm do not shy away from the notion of causation, and hence a resource theory developed within their paradigm could be presented using the same framework that we use here — that of causal models. However, such an approach would likely be framed entirely in terms of *classical* causal models, rather than introducing the notion of GPT causal models.

Advocates of the superluminal causation paradigm would naturally define the free boxes to be those that involve only subluminal causes. Hence, in scenarios wherein the inputs and the outputs at one wing are space-like separated from those at the other wings, so that subluminal causal influences cannot act between the wings, a box is free if and only if it can be realized by a classical common cause. Thus, the natural choice of the free subtheory in the superluminal causation paradigm coincides with the free subtheory in the causal modelling paradigm. On the other hand, the natural choice of the enveloping theory in the superluminal causation paradigm consists of the set of boxes that are classically realizable given superluminal causal influences between the wings. This differs from the enveloping theory in the causal modelling paradigm because it includes boxes that are signalling. In the superluminal causation paradigm, therefore, it is natural to try and quantify the resource in terms of the strength of the superluminal causal influence between the wings that is required to explain it in a classical causal model.[12]

Because the enveloping theory within this paradigm includes not only non-signalling boxes that violate Bell inequalities but signalling boxes as well, the resource theory is rich enough to describe communication between the wings. Therefore, defining the resource theory in this way would not distinguish classical and nonclassical common-cause resources (as we propose to do here), but would instead draw a line between classical common-cause

[12]It should be noted that no *finite* speed of superluminal causal influences can satisfactorily account for the predictions of quantum theory, per Ref. [16], so such influences would need to be assumed to be of infinite speed.

resources and everything else — including classical signalling resources.[13] If one were to go this route, then all of classical Shannon theory would be subsumed in the resource theory. A potential response to this expansion in the scope of the project might be to try to eliminate such signalling resources *by hand*, by demanding that the enveloping theory was constrained to those boxes that are non-signalling among the wings. Such a response, however, seems to compromise the ideals of the superluminal causation paradigm, because no-signalling is an operational notion rather than a realist one.[14]

[13]Note, therefore, that if one seeks to partition resources of a given type into classical and nonclassical varieties, then defining the enveloping theory correctly is just as important as defining the free subtheory correctly.

[14]John Bell famously argued against the idea that no-signalling could embody an assumption of locality in a fundamental physical theory on the grounds that it was too anthropocentric [35]:

> ...the "no signaling" notion rests on concepts which are desperately vague, or vaguely applicable. The assertion that *we cannot signal faster than light* immediately provokes the question: Who do we think *we* are? *We* who can make *measurements*, *we* who can manipulate *external fields*, *we* who can *signal* at all, even if not faster than light? Do *we* include chemists, or only physicists, plants, or only animals, pocket calculators, or only mainframe computers?

Chapter 5

The type-independent resource theory of local operations and shared randomness

Abstract: In space-like separated experiments and other scenarios where multiple parties share a classical common cause but no cause-effect relations, quantum theory allows a variety of nonsignaling resources which are useful for distributed quantum information processing. These include quantum states, nonlocal boxes, steering assemblages, teleportages, channel steering assemblages, and so on. Such resources are often studied using nonlocal games, semiquantum games, entanglement-witnesses, teleportation experiments, and similar tasks. We introduce a unifying framework that subsumes the full range of nonsignaling resources, as well as the games and experiments that probe them, into a common resource theory: that of local operations and shared randomness (LOSR). Crucially, we allow these LOSR operations to locally change the type of a resource, so that players can convert resources of *any* type into resources of any other type, and in particular into strategies for the specific type of game they are playing. We then prove several theorems relating resources and games of different types. These theorems generalize a number of seminal results from the literature, and can be applied to lessen the assumptions needed to characterize the nonclassicality of resources. As just one example, we prove that semiquantum games are able to perfectly characterize the LOSR nonclassicality of every resource of *any* type (not just quantum states, as was previously shown). As a consequence, we show that any resource can be characterized in a measurement-device-independent manner.

5.0.1 Introduction

A key focus in quantum foundations is the study of nonclassicality. Starting from the Einstein-Podolsky-Rosen paradox [100], special focus has been given to experiments involving space-like separated subsystems. In the modern language of causality [12, 29, 88, 324], the key feature of these scenarios is that the subsystems that are being probed share a classical common cause, but do not share any cause-effect channels between them. In such scenarios, quantum theory allows for distributed quantum channels that act as valuable nonclassical resources for accomplishing tasks that would otherwise be impossible.

The most common examples of such resources are entangled quantum states [151] and boxes producing nonlocal correlations [44]; but there are many other types of useful resources. We develop a resource-theoretic [80] framework that unifies a wide variety of these, including quantum states [314], boxes [44], steering assemblages [55, 320], channel steering assemblages [220], teleportages [57, 145], distributed measurements [37], measurement-device-independent steering channels [58], Bob-with-input steering channels [32], and generic no-signaling quantum channels [314]. Free (or classical) resources are those that can be generated freely by local operations and shared randomness (LOSR), encompassing the specific cases of separable quantum states, local boxes, unsteerable assemblages, and so on. Any resource that cannot be simulated by LOSR operations is said to be nonfree, or nonclassical. A resource is said to be at least as nonclassical as another resource if it can be transformed to the second using LOSR transformations. Crucially, such comparisons can be made for resources of arbitrary and potentially differing types.

Some works in the past have focused on LOSR as a resource theory *in specific scenarios*, such as for quantum states [48, 256], for nonlocal correlations [94, 116, 256, 322], and for steering assemblages [58] (albeit under a different name). These previous works focused on one or two types of resources, and most commonly on quantum states. Our framework is more general, but subsumes each of these as a special case.

In addition to introducing this encompassing framework, our second primary goal herein is to study how the type of a resource impacts the methods by which one can characterize its nonclassicality in practice. For example, nonlocal boxes have classical inputs and outputs, and so only weak assumptions [36, 233] about one's laboratory instruments are required for their characterization. However, when a resource has a quantum output, one requires a well-characterized quantum measurement to probe that output and consequently the resource [244]. In such a case, the test of nonclassicality is said to be *device-dependent*, while in adversarial scenarios such as cryptography, the terminology of *trust* is also used [223]. The same idea applies to a quantum input, which must be probed using a well-characterized quantum state preparation device. Thus, only nonlocal boxes can be probed in a *device-independent* manner; *a priori*, quantum states require well-characterized quantum measurement devices; while other objects, such as steering assemblages, require a mixture of both [56]. Consequently, it is important to determine under what circumstances devices of one type may be converted into devices of a second type *in a manner that does not degrade their usefulness as a resource.* If such a conversion is possible, then one may be able to lessen the assumptions and technological requirements needed to characterize one's devices.

In some particular cases, previous work has studied this question of whether the nonclassicality of a quantum state can be characterized by first applying free operations that convert it to another type of resource. For example, we know that some Werner states [22, 315] have a local model for all measurements; such nonclassical states can only be transformed into classical boxes, and so all information about their nonclassicality is lost in the conversion. In contrast, the main result of Ref. [48] proves that every entangled state can have its nonclassicality encoded in a semiquantum channel. Additionally, in Ref. [57], it is shown that every entangled state can generate a type of no-signaling channel (recently termed a teleportage [145]) that could not be generated by any separable state and which is useful for some task related to quantum teleportation [188].

It is useful to distinguish between qualitative versus quantitative characterizations of nonclassicality. To highlight the distinction, it is instructive to examine one particular line of research. Ref. [48] is often advertised as proving that the nonclassicality of every entangled state can be revealed in a generalization of nonlocal games termed *semiquantum games* (which were later used to construct *measurement-device-independent entanglement witnesses* [41]). However, this claim is actually a (qualitative) corollary of the (quantitative) main theorem, which showed that the performance of states in semiquantum games *exactly* reproduces the classification of entangled states under LOSR transformations. Subsequent works [41, 242] focused on the qualitative distinction between classical and nonclassical resources, but still later works reinterpreted the payoffs of semiquantum games as measures of entanglement [245, 273], thus reconnecting with the quantitative nature of Buscemi's original work. Note also that the quantitative study of entanglement is historically linked to entanglement monotones [311]. However, the study of nonclassicality cannot be reduced to a single such measure, as there are many inequivalent species of nonclassicality even in the simplest cases [322]. Informed by the recent formalization of resource theories [80], we study the fundamental mathematical object—the preorder of resources under LOSR transformations. One can then derive specific nonclassicality witnesses and monotones [246], each of which provides an incomplete characterization of the preorder.

As implied just above, the mathematical structure that best allows for comparison between objects that need not be strictly ordered is a *preorder*. Formally, a preorder is an ordering relation that is reflexive ($a \succeq a$) and transitive ($a \succeq b$ and $b \succeq c$ implies $a \succeq c$)[1]. Our work focuses on three distinct preorders, which the reader should be careful to distinguish. First, there is the preorder $R \succeq_{\text{LOSR}} R'$ (sometimes denoted $R \xmapsto{\text{LOSR}} R'$) that indicates if a resource R can be converted into another resource R' by LOSR transformations (Definition 1). Second, there is the preorder $\succeq_{\text{type}}$ over resource types that orders those types according to their ability to encode nonclassicality (Definition 2). Finally, there is the preorder $\succeq_{\mathcal{G}_T}$ that ranks resources according to their performance with respect to the set $\mathcal{G}_T$ of all games of a particular type T (Definition 5).

This chapter is best read alongside Ref. [246]. In the current chapter, we present a general framework to study quantum resources of arbitrary types, and we quantify the nonclassicality of these resources within a type-independent resource theory of local operations and shared randomness. Here, our main results center on showing how resources of one type can be more easily characterized by first converting them to resources of a second type. In Ref. [246], our aim is practical and computational, focusing on how data can be used to characterize one's resources using off-the-shelf software. There, we include type-independent techniques for computing witnesses that can certify the nonclassicality of a resource, as well as techniques for computing the value of type-independent monotones (which we introduce therein).

5.0.2 Organization of the chapter

In Section 5.1, we discuss various types of resources. We inventory the 9 possible types of a single party's partition of a resource, where that party's input and output may each be trivial, classical, or quantum. Focusing on the 81 bipartite resource types for simplicity, we recognize 10 types that have been studied in the literature and identify 5 new nontrivial resource types. All other bipartite resource types are either trivial or equivalent up to a symmetry. We then define LOSR transformations between resources of arbitrary types, as well as the ordering over resources that this induces.

In Section 5.2, we define a precise sense in which some types can express the LOSR nonclassicality of other types. In many cases, conversions from a resource of one type to another type necessarily degrade the nonclassicality of the resource, as in Werner's example. In other cases, one can perfectly encode the nonclassicality of any given resource into some resource of the target type, as in Buscemi's example. For every single-party type, we ask which can perfectly encode the nonclassicality of which others, and we answer this question

[1] A preorder is distinguished from a partial order by the fact that $a \succeq b$ and $b \succeq a$ need not imply $a = b$. In a partial order, $a \succeq b$ and $b \succeq a$ implies $a = b$.

for almost every pair, with the exception of one open question. From these considerations of single party types, one can deduce encodings of more complicated resource types which involve multiple parties. Most strikingly, we show that semiquantum channels (with quantum inputs and classical outputs) are universal, in the sense that the nonclassicality of all resources can be encoded into them.

In Section 5.3, we give an abstract framework for probing the nonclassicality of resources, subsuming as special cases the notions of nonlocal games [44], semiquantum games [48], steering [58, 320] and teleportation [57] experiments, and entanglement witnessing [73]. In our framework, every type of resource has a corresponding type of game, where a game of some type maps every resource of that type to a real number. (E.g., in nonlocal and semiquantum games, this number is the usual average game payoff). We then show how resources of any type can be used to play a game designed for one specific type. In some cases, games of one type can *completely* characterize the nonclassicality of every resource of another type. For example, Ref. [48] showed that the LOSR nonclassicality of every quantum state is perfectly characterized by the set of semiquantum games. We generalize these ideas by proving that if one type can encode another, then games of the first type can perfectly characterize the LOSR nonclassicality of all resources of the second type. Together with our results on which types can encode which others, this expands the known methods for quantifying LOSR nonclassicality in practice and in theory. For example, our result on the universality of the semiquantum type implies that any resource of any type can be characterized by some semiquantum game, and hence can be characterized in a measurement-device-independent manner.

In Section 5.4, we relate our work to existing results. First, we note how our results generalize the main result of Ref. [48], showing that semiquantum games can completely characterize the LOSR nonclassicality of arbitrary resource, not just of quantum states. Next, we show that the results of Ref. [58] are a special case of two of our theorems when one applies steering experiments to quantify the nonclassicality of quantum states; further, our theorems provide a generalization of these arguments to more general experiments and types of resources. Finally, we show that the LOSR nonclassicality of every quantum state is *completely* characterized by the set of teleportation games, and thus that the results of Ref. [57] can be extended to be quantitative as well as qualitative.

5.1 Resource types and LOSR transformations between them

We are interested in scenarios where the relevant parties share a classical common cause but do not share any cause-effect channels. For example, parties who perform experiments at space-like separation cannot access classical communication. For simplicity, we henceforth focus on bipartite scenarios; however, all of our results generalize immediately to arbitrarily many parties. We will consider only nonsignaling resources [28, 224] throughout this chapter.[2] We will not specifically consider post-quantum channels in this chapter, although one might naturally extend our work to include these as resources. Hence, in this chapter a resource is a completely positive [212], trace-preserving, nonsignaling quantum channel. The parties may share various types of resources, that we now classify by type.

5.1.1 Partition-types and global types

In this chapter, we use the term **type** (of a resource) to refer exclusively to whether the various input and output systems are trivial (I), classical (C), or quantum (Q). A system is said to be trivial if it has dimension one, is said to be classical if all operators on its Hilbert space are diagonal, and is otherwise said to be quantum. (See Ref. [246] for more details.) Additionally, if a resource has more than one input (output), which may be of different types, we imagine grouping them together, yielding an effective input (output) whose type is the least expressive type which embeds all those in the grouping, where quantum systems embed classical systems, which embed trivial systems.

We will denote the type of a single party's share of a resource by $T_i := X_i \!\to\! Y_i$, where i labels the party and $X, Y \in \{\mathsf{I}, \mathsf{C}, \mathsf{Q}\}$, with X labeling whether the input to that party is trivial (I), classical (C), or quantum (Q) and Y labeling the output similarly. We will refer to T_i as the **partition-type** of party i.

We can then denote the **global type** of an n-party resource as $T := T_1 T_2 ... T_n \simeq X_1 X_2 ... X_n \!\to\! Y_1 Y_2 ... Y_n$. Note that while the specification of the global type of a resource fixes the number of parties and the types of their partitions of the resource, the specification of a partition-type does not constrain either the number of other parties who share the

[2] In fact, if one wishes to interpret resourcefulness as *nonclassicality*, then one must further restrict the enveloping theory to those resources that can be generated by local operations and quantum common causes. For non-signaling resources that *cannot* be realized in this manner [31], resourcefulness may originate in the nonclassicality of a common-cause process *or* in *classical* communication channels (which are fine-tuned so as to not exhibit signaling).

resource, nor the types of those other partitions. One could also consider partition-types for partitions of a resource that involve more than one party, but this chapter makes use only of partition-types that involve a single party.

We now describe the ten examples of resource types from Fig. 5.1, setting up some explicit terminology and conventions as we go. We graphically depict trivial, classical, and quantum systems by the lack of a wire, a single wire, and a double wire, respectively.

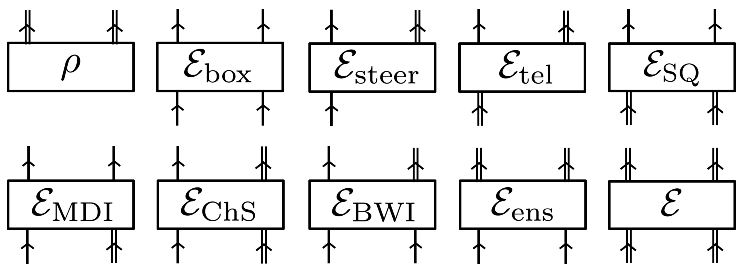

Figure 5.1: Common types of no-signaling resources, where classical systems are represented by single wires and quantum systems are represented by double wires. (a) A quantum state ρ has type II→QQ. (b) A box $\mathcal{E}_{\text{box}}$ has type CC→CC. (c) A steering assemblage $\mathcal{E}_{\text{steer}}$ has type CI→CQ. (d) A teleportage $\mathcal{E}_{\text{tel}}$ has type QI→CQ. (e) A semiquantum channel $\mathcal{E}_{\text{SQ}}$ has type QQ→CC. (f) A measurement-device-independent steering channel $\mathcal{E}_{\text{MDI}}$ has type CQ→CC. (g) A channel steering assemblage $\mathcal{E}_{\text{ChS}}$ has type CQ→CQ. (h) A Bob-with-input steering channel $\mathcal{E}_{\text{BWI}}$ has type CC→CQ. (i) An ensemble-preparing channel $\mathcal{E}_{\text{ens}}$ has type CC→QQ. (j) A quantum channel $\mathcal{E}$ has type QQ→QQ.

Fig. 5.1(a) depicts a **quantum state**, the canonical quantum resource. Bipartite quantum states have type II→QQ; that is, they have no inputs and both outputs are quantum. The nonclassicality of quantum states is often quantified using the resource theory of local operations and classical communication (LOCC). While this is appropriate in some contexts, allowing classical communication for free is not appropriate in the context of space-like separated experiments, nor in any other scenario where distributed systems are unable to causally influence one another. In such cases, LOSR operations are the relevant ones for quantifying nonclassicality of any resource, including quantum states, and it is *LOSR-entanglement*, not LOCC-entanglement, that is relevant, as argued extensively in Ref. [256].

Fig. 5.1(b) depicts another canonical type of resource [28, 44], often termed a correlation or a box-type resource, or **box** for short. Bipartite boxes have type CC→CC; that is, both parties have a classical input and a classical output. Extensive research has been done on boxes, e.g. to characterize the set of local boxes [44] and the possible LOSR conversions between them [94, 150, 322]. The fact that we wish to subsume boxes in our framework provides another reason to focus on LOSR as opposed to LOCC, since LOSR has been argued to be the appropriate set of free operations in this context [322] Furthermore, under unbounded LOCC *all* boxes would be deemed free, even nonlocal or signaling boxes.

Fig. 5.1(c) depicts the type of resource that arises naturally in a steering scenario [55, 100, 115, 221, 265, 277, 301, 320], often termed an **assemblage** [227]. Such resources have type CI→CQ; that is, the first party has a classical input and classical output, while the second party has no input and a quantum output.

Fig. 5.1(d) depicts a type of resource that arises naturally in a teleportation scenario [57, 294], termed **teleportages** [145]. Such resources have type QI→CQ. Intuitively, given a teleportage, one would complete the standard teleportation protocol by applying one of a set of unitaries on the quantum output, conditioned on the classical output. The precise operational sense in which these teleportages relate to the possibility of implementing an effective quantum channel is still being investigated [188][3].

Fig. 5.1(e) depicts the type of resource that arises naturally in semiquantum games, namely type QQ→CC. We will term these **distributed measurements** or **semiquantum channels**, since they arise in multiple contexts where one term [37] or the other [48] is more natural.

Fig. 5.1(f) depicts the type of resource that arises naturally in measurement-device-independent (MDI) steering scenarios [58], namely type CQ→CC. We will term these **MDI-steering channels**.

Fig. 5.1(g) depicts the type of resource that arises naturally in channel steering scenarios [220], often termed a **channel assemblage**. Such resources have type CQ→CQ.

Fig. 5.1(h) depicts the type of resource that arises when one generalizes a steering scenario to have a classical input on the steered party [32], termed a **Bob-with-input steering channel**. Such resources have type CC→CQ.

Fig. 5.1(i) depicts a distributed classical-to-quantum channel, of type CC→QQ. We will term these **ensemble-preparing channels**. An interesting example of such a channel can be found in Ref. [31] (see Eq. 82).

Fig. 5.1(j) depicts a generic bipartite **quantum channel**, of type QQ→QQ.

[3]While LOSR is clearly the correct set of free operations for studying resources in Bell scenarios and other common cause scenarios, the same is not true for teleportation experiments, which might be better described by another resource theory (such as LOCC). The surprising insight that follows from Ref. [57] is that a great deal can nonetheless be learned about teleportation scenarios by studying LOSR.

This list is not exhaustive. Even in the bipartite case, one might wonder how many nontrivial resource types there are, and whether all of these have been studied. First, note that the partition-type $\mathsf{I}{\to}\mathsf{I}$ corresponds to a trivial party. As there are no nonclassical resources involving only one party, all bipartite types involving partition-type $\mathsf{I}{\to}\mathsf{I}$ for either party are trivial. Two other partition-types, $\mathsf{C}{\to}\mathsf{I}$, and $\mathsf{Q}{\to}\mathsf{I}$, are also trivial, since the no-signaling principle guarantees that their input cannot affect the operation of the remaining parties [246]. Moreover, some global types are equivalent up to exchange of parties, in which case we will consider only a single representative. This leads us to our first open question.

Open Question 1. *Even in the bipartite case, there are five nontrivial global types of resources that have not (to our knowledge) been previously studied, namely* $\mathsf{QC}{\to}\mathsf{CQ}$*,* $\mathsf{CQ}{\to}\mathsf{QQ}$*,* $\mathsf{IQ}{\to}\mathsf{QQ}$*,* $\mathsf{QQ}{\to}\mathsf{CQ}$*, and* $\mathsf{CI}{\to}\mathsf{QQ}$*. Do any of these correspond to scenarios that are interesting in their own right?*

At the very least, each new type implies a novel form of 'nonlocality'. What remains to be seen is whether these will be directly relevant for quantum information processing tasks.

5.1.2 Free versus nonfree resources

A nonsignaling resource (of any type) is **free** with respect to LOSR, or **classical**[4], if the parties can generate it freely using local operations and shared randomness. This notion of being free with respect to LOSR subsumes the established notions of classicality for every type of resource in Fig. 5.1; e.g. for states it coincides with separability [151], for boxes, it coincides with admitting of a local hidden variable model [44], for assemblages it coincides with unsteerability [55, 301], for teleportages it coincides with the inability to outperform classical teleportation [57], and so on, as pictured in Fig. (5.2).

Any resource that cannot be simulated by local operations and shared randomness is *non-free* and constitutes a resource of LOSR nonclassicality. The purpose of our type-independent resource theory of LOSR is to quantitatively characterize nonfree resources of arbitrary types, as we now do.

[4]In reference to the fact such resources can be generated by classical common causes. Classicality of a *resource* is not to be confused with classicality of input and output systems.

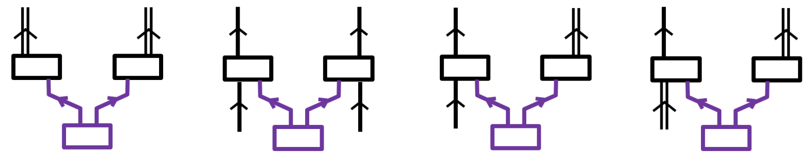

Figure 5.2: Free LOSR resources are those that can be simulated by local operations (in black) and shared randomness (in purple). We depict four canonical types of free resources here: separable states, local boxes, unsteerable assemblages, and classical teleportages.

5.1.3 Type-changing LOSR operations

Two parties in an LOSR scenario transform resources using free LOSR operations. Most previous works that studied LOSR focused on conversions between specific types of resources; for example, Refs. [94, 116, 322] considered LOSR conversions from boxes to boxes, Ref. [48] considered LOSR conversions from quantum states to quantum states, and Ref. [58] considered LOSR conversions[5] from quantum states to assemblages. In keeping with our aim to unify a range of scenarios in one framework, and because local operations can freely change the type of a resource, we do *not* restrict attention to conversions among resources of fixed type, but rather allow conversions among resources of all types.

We denote the set of all operations that can be generated by local operations and shared randomness by LOSR. As depicted in Fig. 5.3(a), the most general local operation on a given party is given by a comb [68], and the different parties may correlate their choice of comb using their shared randomness. Note that this shared randomness can be transmitted down the side channel of each local comb, which implies that this depiction of LOSR is completely general and is convex [322] for conversions from one fixed type to another. We will denote an element of this set by $\tau \in$ LOSR and a generic resource of arbitrary type by R.

As in any resource theory [80], the set of free operations induces a preorder over the set of all resources. Here, we write $R \xmapsto{\text{LOSR}} R'$ whenever there exists some $\tau \in$ LOSR such that $R' = \tau \circ R$, and we say that R is **at least as nonclassical** (as resourceful) as R'. We denote the ordering relation for the preorder defined by LOSR conversions as $\succeq_{\text{LOSR}}$:

[5]In this last case, the authors introduced the term local operations with steering and shared randomness (LOSSR); however, the operations they consider involve all and only the subset of LOSR operations from quantum states to assemblages, so there is no need for the new term LOSSR.

Definition 1. *For resources R and R' of different and arbitrary type, we say that $R \succeq_{\mathrm{LOSR}} R'$ iff $R \xmapsto{\mathrm{LOSR}} R'$.*

This definition allows us to make rigorous, quantitative comparisons of LOSR nonclassicality among resources of arbitrary types. The relation $\succeq_{\mathrm{LOSR}}$ is a preorder, as there exists an identity LOSR transformation (reflexivity), and LOSR transformations compose (transitivity).

Two resources R and R' are equally nonclassical if they are interconvertible under LOSR; that is, if $R \xmapsto{\mathrm{LOSR}} R'$ and $R' \xmapsto{\mathrm{LOSR}} R$. We denote this $R \xleftrightarrow{\mathrm{LOSR}} R'$, and we say that R and R' are in the same LOSR equivalence class.

We give several examples of conversions among resource types in Fig. 5.3, depicting wires of unspecified (and arbitrary) type by dashed double lines.

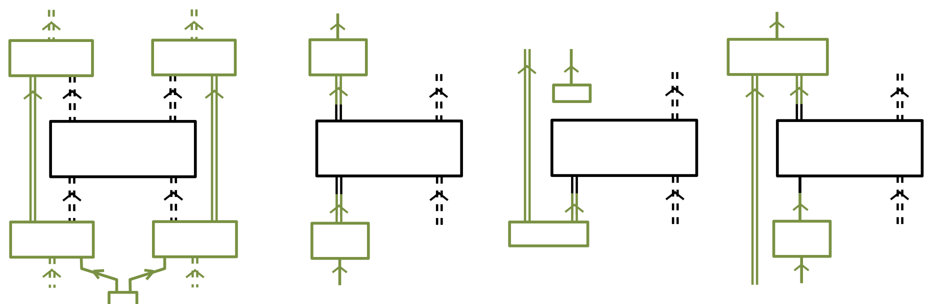

Figure 5.3: Some type-changing operations (in green), as described in the main text. Dashed wires denote systems of arbitrary and unspecified type. (a) A generic bipartite type-changing LOSR transformation. (b) A transformation taking partition-type $\mathsf{Q}\to\mathsf{Q}$ to $\mathsf{C}\to\mathsf{C}$. (c) A transformation taking partition-type $\mathsf{Q}\to\mathsf{I}$ to $\mathsf{I}\to\mathsf{Q}$. (d) A transformation taking partition-type $\mathsf{C}\to\mathsf{Q}$ to $\mathsf{Q}\to\mathsf{C}$.

Fig. 5.3(a) depicts a generic bipartite type-changing LOSR operation. Fig. 5.3(b) depicts an example of a specific transformation that takes the left partition of the resource from $\mathsf{Q}\to\mathsf{Q}$ to $\mathsf{C}\to\mathsf{C}$. It is generated by composition with a local ensemble-preparing channel and a local measurement channel, respectively. Fig. 5.3(c) depicts an example of a specific transformation that takes the left partition of the resource from $\mathsf{Q}\to\mathsf{I}$ to $\mathsf{I}\to\mathsf{Q}$. The transformation is generated by (sequential) composition with half of an entangled state and parallel composition with a classical system in some fixed state. In this example, the output system type is quantum, since it is comprised of a classical and quantum system. Fig. 5.3(d) depicts an example of a specific transformation that takes the left partition of the resource from $\mathsf{C}{\to}\mathsf{Q}$ to $\mathsf{Q}{\to}\mathsf{C}$, generated by a stochastic transformation on the classical input to the resource and performing a joint quantum measurement channel on the quantum output of the resource together with some new quantum input.

5.2 Encoding nonclassicality of one type of resource in another type

We now consider a preorder over *types of resources* (rather than over the resources themselves). This allows us to formally compare the different manifestations of nonclassicality. For example, this preorder provides a *formal* sense in which entanglement and nonlocality are incomparable types of nonclassicality. Surprisingly, we will also show that not all types of nonclassicality are incomparable.

Definition 1. *Global type T* **encodes the nonclassicality of** *global type T', denoted $T \succeq_{\text{type}} T'$, if for every resource R' of type T', there exists at least one resource R of type T such that $R' \overset{\text{LOSR}}{\longleftrightarrow} R$.*

In other words, there exists some resource of the higher type in every equivalence class of resources of the lower type. Several well-known examples of such encodings will be given shortly.

To study the preorder over global types, it is also useful to consider a preorder over partition-types; that is, over the nine possible types $T_i := X_i{\to}Y_i$ of a single party's share of a resource. Considering without loss of generality the first party, denoted by subscript 1, we say that type T_1 is higher in the preorder than type T_1' if for every resource of type $T_1'T_2...T_n$, there exists a resource of type $T_1T_2...T_n$ that is in the same LOSR equivalence class (for all numbers of parties n). Equivalently, this means that the LOSR equivalence class of any resource with partition-type T_1' on the first party always contains at least one resource of partition-type T_1 (on the first party). We denote this second ordering relation $\succeq_{\text{type}}$:

Definition 2. *We say that $T_1 \succeq_{\text{type}} T_1'$ iff for all R' of type $T_1'T_2...T_n$ (as one ranges over all $T_2, ..., T_n$ and all n), there exists R of type $T_1T_2...T_n$ in the LOSR equivalence class of R', that is, satisfying $R' \overset{\text{LOSR}}{\longleftrightarrow} R$.*

In such cases, we say that partition-type T_1 encodes (the nonclassicality of) all resources of partition-type T_1', or more simply that type T_1 encodes type T_1'.

If every partition-type of some given global type is higher than the corresponding partition-type of a second global type on every partition, then the first type is necessarily higher in the preorder over global types. Hence, orderings over global types can often be deduced from orderings over partition-types.

As a trivial example, it is clear that the global type $\mathsf{QQ}{\to}\mathsf{QQ}$ (that of bipartite quantum channels) is above every other bipartite type. For example, it is above the global type $\mathsf{II}{\to}\mathsf{QQ}$ (that of bipartite quantum states) in the preorder, so that $\mathsf{QQ}{\to}\mathsf{QQ} \succeq_{\text{type}} \mathsf{II}{\to}\mathsf{QQ}$, since the former is an instance of the latter where the inputs to the channel are trivial. In other words: given any bipartite quantum state, there is a bipartite quantum channel that is in the same LOSR equivalence class—namely, the quantum state itself, viewed as a channel from the trivial system to a quantum system on each partition. We will refer to such trivial instances of ordering among types as **embeddings** of one type into the other.

Two resource types are in the same equivalence class over types if any resource of either type can be converted into a resource of the other type which is in the same LOSR equivalence class. For example, the three partition-types $\mathsf{I}{\to}\mathsf{I}$, $\mathsf{C}{\to}\mathsf{I}$, and $\mathsf{Q}{\to}\mathsf{I}$ are all in the lowest equivalence class over partition-types, since (as discussed above) they never play any role in the nonclassicality of any nonsignaling resource.

Understanding the scope of nonclassicality-preserving conversions between resources of different global types is particularly useful for devising experimental measures and witnesses of nonclassicality, as we discuss in Section 5.3.3 (and in Ref. [246]). Abstractly, this is because one type is above another type if there exists an embedding of the partial order over equivalence classes of resources of the lower type into the partial order of the higher type. When this is the case, techniques for characterizing the preorder of the higher type give direct information about the preorder of the lower type.

5.2.1 Determining which types encode the nonclassicality of which others

In this section, we derive all but two of the ordering relations that hold between the possible pairings of partition-types by leveraging various results from the literature. These results are summarized in Table 5.1. As discussed above, orderings over global types can be deduced from these.

As discussed above, there are no nonfree resources that nontrivially involve the types $\mathsf{I}{\to}\mathsf{I}$, $\mathsf{C}{\to}\mathsf{I}$, or $\mathsf{Q}{\to}\mathsf{I}$, so we need not discuss them further. There remain 6 nontrivial types, and hence 36 ordering relations to check. These are all shown in the table. If the column

Does type T encode the nonclassicality of type T'?

T' \ T	$I \to C$	$I \to Q$	$C \to C$	$C \to Q$	$Q \to C$	$Q \to Q$
$I \to C$	✓ embed	✓ embed	✓ embed	✓ embed	✓ embed	✓ embed
$I \to Q$	✗ trans.	✓ embed	✗ Werner states	✓ embed	✓ semi–quantum games	✓ embed
$C \to C$	✗ trans.	✗ LOSR cannot entangle	✓ embed	✓ embed	✓ embed	✓ embed
$C \to Q$	✗ trans.	✗ trans.	✗ trans.	✓ embed	✓ Thm 3	✓ embed
$Q \to C$	✗ trans.	✗ trans.	✗ trans.	?	✓ embed	✓ embed
$Q \to Q$	✗ trans.	✗ trans.	✗ trans.	?	✓ Thm 3	✓ embed

Table 5.1: A green check mark in a given cell indicates that the column type T is higher in the order over partition-types than the row type T' (denoted $T \succeq_{\text{type}} T'$), while a red cross indicates that it is not higher (denoted $T \not\succeq_{\text{type}} T'$. The text in each cell alludes to the proof (given in the main text) of that ordering relation. Two relations are unknown, as indicated by blue question marks.

resource type T is higher in the order than the row type T', so that $T \succeq_{\text{type}} T'$, then we indicate this with a green check mark in the corresponding cell in the table. If instead $T \not\succeq_{\text{type}} T'$, we indicate this with a red cross. In each case, we briefly allude to the logic behind the proofs for that particular ordering—proofs which we now give.

As stated in Section 5.2, a type is higher in the order than all types that it embeds, where quantum systems embed classical systems, which embed trivial systems. In the table, we indicate these trivial ordering relations by the word 'embed'.

Next, recall that Werner proved the existence of entangled states that cannot violate any Bell inequality involving projective measurements [315]. It was subsequently proved that this holds true even for arbitrary local measurements [22], a result that holds even if the choice of local measurements are made in a correlated fashion using shared randomness. This constitutes the most general LOSR conversion scheme from quantum states to boxes. In other words, an entangled Werner state cannot be converted into *any* nonfree box, much less into a box that is in its LOSR equivalence class (as would be required for encoding its nonclassicality into a box-type resource). It follows that global type $\mathsf{CC} \to \mathsf{CC}$ is not above global type $\mathsf{II}{\to}\mathsf{QQ}$, which in turn implies that partition-type $\mathsf{C}{\to}\mathsf{C}$ is not above partition-type $\mathsf{I}{\to}\mathsf{Q}$. That is, $\mathsf{C}{\to}\mathsf{C} \not\succeq_{\text{type}} \mathsf{I}{\to}\mathsf{Q}$, as is indicated in the table by the phrase 'Werner states'.

In addition, it is well known that LOCC can generate arbitrary boxes and yet cannot generate any entangled state. Since LOSR operations form a subset of LOCC operations, this implies that LOSR operations applied to any box (of type $\mathsf{CC}{\to}\mathsf{CC}$) cannot generate *any* nonfree state (of type $\mathsf{II}{\to}\mathsf{QQ}$), much less a state in its LOSR equivalence class. Hence, global type $\mathsf{II}{\to}\mathsf{QQ}$ is not above global type $\mathsf{CC}{\to}\mathsf{CC}$, which in turn implies that partition-type $\mathsf{I}{\to}\mathsf{Q}$ is not above partition-type $\mathsf{C}{\to}\mathsf{C}$. That is, $\mathsf{I}{\to}\mathsf{Q} \not\succeq_{\text{type}} \mathsf{C}{\to}\mathsf{C}$, as is indicated in the table by the phrase 'LOSR cannot entangle'.

We can use transitivity of the ordering relation to prove that $\mathsf{I}{\to}\mathsf{C}$ is not above $\mathsf{I}{\to}\mathsf{Q}$ and is not above $\mathsf{C}{\to}\mathsf{C}$, and that none of $\mathsf{I}{\to}\mathsf{C}$, $\mathsf{I}{\to}\mathsf{Q}$, or $\mathsf{C}{\to}\mathsf{C}$ are above any of $\mathsf{C}{\to}\mathsf{Q}$, $\mathsf{Q}{\to}\mathsf{C}$, and $\mathsf{Q}{\to}\mathsf{Q}$. For example, from the fact that $\mathsf{C}{\to}\mathsf{C}$ is above $\mathsf{I}{\to}\mathsf{C}$ and the fact that $\mathsf{C}{\to}\mathsf{C}$ is not above $\mathsf{I}{\to}\mathsf{Q}$, it must be that $\mathsf{I}{\to}\mathsf{C}$ is not above $\mathsf{I}{\to}\mathsf{Q}$. If it were otherwise, one would have $\mathsf{C}{\to}\mathsf{C}$ above $\mathsf{I}{\to}\mathsf{C}$ above $\mathsf{I}{\to}\mathsf{Q} \implies \mathsf{C}{\to}\mathsf{C}$ above $\mathsf{I}{\to}\mathsf{Q}$, which is false. The other transitivity arguments run analogously. In the table, we indicate all such ordering relations by the abbreviation 'trans.'.

One of the authors proved in Ref. [48] that there exists some semiquantum channel (of type $\mathsf{QQ}{\to}\mathsf{CC}$) in the same equivalence class as any given quantum state (of type $\mathsf{II}{\to}\mathsf{QQ}$). A slight reframing of this result implies that the semiquantum partition-type $\mathsf{Q}{\to}\mathsf{C}$ is higher in the order than $\mathsf{I}{\to}\mathsf{Q}$, as we show below. That is, $\mathsf{Q}{\to}\mathsf{C} \succeq_{\text{type}} \mathsf{I}{\to}\mathsf{Q}$, as is indicated in the table by the phrase 'semiquantum games'.

Finally, as we prove in Theorem 1, the semiquantum partition-type $\mathsf{Q}{\to}\mathsf{C}$ is higher in the order than all other partition-types. The ordering relations that follow from our proof but not from previous work, namely $\mathsf{Q}{\to}\mathsf{C} \succeq_{\text{type}} \mathsf{C}{\to}\mathsf{Q}$ and $\mathsf{Q}{\to}\mathsf{C} \succeq_{\text{type}} \mathsf{Q}{\to}\mathsf{Q}$, are indicated in the table by the phrase 'Thm 3'.

This proves all the results shown in the table. There remain two unknown ordering relations, indicated in the table by question marks; namely whether $\mathsf{C}{\to}\mathsf{Q}$ is higher in the order than either $\mathsf{Q}{\to}\mathsf{C}$ or $\mathsf{Q}{\to}\mathsf{Q}$. Because $\mathsf{Q}{\to}\mathsf{C}$ and $\mathsf{Q}{\to}\mathsf{Q}$ are in the same equivalence class (at the top of the order), the answer to both of these questions must be the same; that is, either $\mathsf{C}{\to}\mathsf{Q}$ encodes them both, or it encodes neither. Such an encoding could have dramatic practical consequences. For example, if the encoding can be done with a fixed transformation (which is not a function of the resource to be converted), then this would enable the possibility of *preparation-device-independent* quantification of nonclassicality.

Open Question 2. *Can the LOSR nonclassicality of any resource be perfectly characterized in a* preparation-device-independent *manner?*

5.2.2 Semiquantum channels are universal encoders of nonclassicality

To complete the arguments of the last section, we prove that the semiquantum partition-type can encode any other partition-type. The consequences of this fact are fleshed out further in Section 5.3.3.

Theorem 1. *The semiquantum partition-type* $\mathsf{Q}{\to}\mathsf{C}$ *is in the unique equivalence class at the top of the order over partition-types. That is, it can encode the nonclassicality of all other partition-types.*

Proof. Consider a bipartite channel $\mathcal{E}$ that has a quantum output of dimension d, together with arbitrary other outputs and inputs (denoted by dashed double lines), as shown in black in Fig. 5.4(a). One can transform $\mathcal{E}$ into a resource with a quantum input of dimension d and a classical output of dimension d^2 by composing $\mathcal{E}$ with a Bell measurement as shown in green in Fig. 5.4(a); that is, by performing a measurement in a maximally entangled basis on the quantum output of $\mathcal{E}$ and a new quantum input of the same dimension d. To

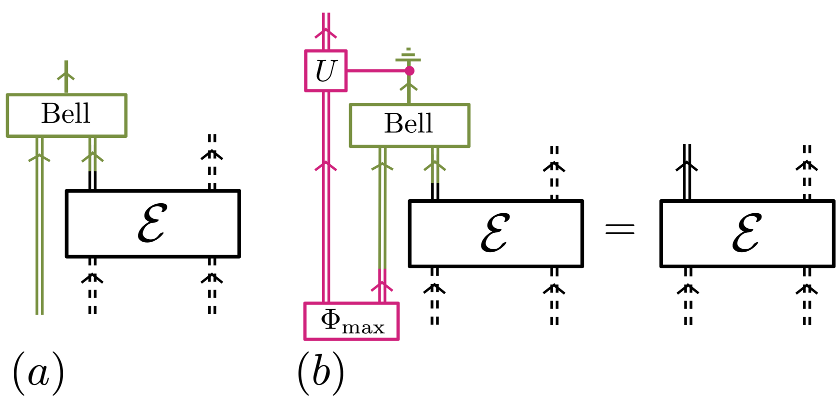

Figure 5.4: (a) A free transformation (in green) that converts a quantum output to a classical output together with a new quantum input. (b) This transformation does not change the LOSR equivalence class, since it has a left inverse (shown in pink) which is a free transformation.

see that this transformation preserves LOSR equivalence class, it suffices to note that there exists a local (and hence free) operation, shown in pink on the left-hand side of Fig. 5.4(b), which takes the transformed channel back to the original channel $\mathcal{E}$. In particular, this local operation feeds one half of a maximally entangled state Φ_{max} into the Bell measurement, and then performs a correcting unitary operation U on the other half of the entangled state, conditioned on the classical outcome of the Bell measurement. For the correct choice of correction operations, the overall transformation on $\mathcal{E}$ is just the well-known teleportation protocol [38], and so the equality shown in Fig. 5.4(b) holds. Hence, the channel in Fig. 5.4(a) is in the same LOSR equivalence class as $\mathcal{E}$, which implies that every partition of a resource can be transformed to a resource of type $\mathsf{Q}\to\mathsf{C}$ in the same equivalence class. □

Note that $\mathsf{Q}\to\mathsf{Q}$ is trivially also at the top of the order, since every other type embeds into it. It is thus in the same equivalence class as $\mathsf{Q}\to\mathsf{C}$.

5.3 A unified framework for distributed games of all types

A variety of 'games' have been studied for the purposes of quantifying nonclassicality of various types of resources. For instance, the nonclassicality of quantum states has been studied from the point of view of nonlocal games and semiquantum games, as well as teleportation, steering, and entanglement witnessing experiments. Nonlocal games have also been used to study the nonclassicality of boxes.

In fact, there is a natural class of distributed tasks for every type of resource, including one for each of the common types in Section 5.1.

Definition 3. *For a given global type T, we define a distributed* **T-game** *as a linear map from resources of type T to the real numbers.*

The set $\mathcal{G}_T$ of all such maps for fixed T is the set of T-games, and a resource of type T is said to be a **strategy** for a T-game. This last terminology is motivated by the fact that no matter how complicated the players' tactics, their score for a given T-game only depends on the resource of type T that they ultimately share with the referee. We will refer to any game of any type as a distributed game.

In Fig. 5.5, we depict four distributed games together with the type of resource that acts as a strategy for that game. We represent a game diagrammatically as a monolithic comb with appropriate input and output structure such that composition of the comb corresponding to a game G_T with a strategy $\mathcal{E}_T$ of type T yields a circuit with no open inputs or outputs, representing the real number $G_T(\mathcal{E}_T)$.

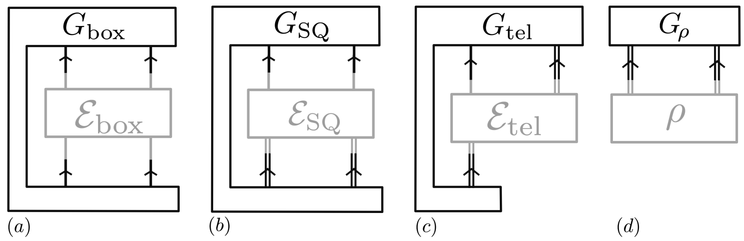

Figure 5.5: Some games and their strategies. (a) Boxes are strategies for nonlocal games. (b) Semiquantum channels are strategies for semiquantum games. (c) Teleportages are strategies for teleportation games. (d) Entangled states are strategies for entanglement witnesses.

5.3.1 Implementations of a game

We have noted that a variety of games and experiments can be viewed abstractly under the umbrella of T-games. The *practical* meaning of such games is made more clear by considering the following two-step procedure, by which a referee can implement any game (of any type T). This procedure is depicted on the right-hand side of Fig. 5.6.

First, the referee performs a tomographically complete measurement on the composite system defined by the collection of output systems of the given strategy $\mathcal{E}_T$, and implements a preparation drawn at random from a tomographically complete set of preparations on the composite system defined by the collection of all the systems that are inputs of $\mathcal{E}_T$. In fact, it suffices for the referee to perform tomographically complete measurements and preparations *independently* on every input and output, as depicted in the dashed box in Fig. 5.6. We will refer to this process as the application of an **analyzer** Z to the given strategy. That is, an analyzer Z is a linear and tomographically complete map from strategies to correlations of the form $P_{Z\circ\mathcal{E}_T}(ab|xy) := Z \circ \mathcal{E}_T$, with a, b labeling the values of the classical outputs of Z and x, y the values of the classical inputs of Z. Second, the referee uses a fixed payoff function $F_{\text{payoff}}(abxy)$ to assign a real number $G_T(\mathcal{E}_T) = \sum_{abxy} F_{\text{payoff}}(abxy) P_{Z\circ\mathcal{E}_T}(ab|xy)$ to strategy $\mathcal{E}_T$.

This point of view on games is useful for the proof of Theorem 2, and it is also useful for establishing a physical picture of games of each type. For example, in a Bell experiment, one applies LOSR operations (or often just LO operations) in order to convert one's quantum state to a conditional probability distribution, and the payoff function in the game constitutes the Bell inequality that one tests. As a second example, see Ref. [188] for a study of various teleportation games. As noted therein, there are interesting teleportation tasks (which admit of a simple operational interpretation) beyond merely attempting to establish an identity channel between two parties using shared entanglement. However, in the rest of this chapter it will be simpler to view a game in the abstract (simply as a linear map from resources of a given type to the real numbers), and we will leave the further investigation of such games (beyond the cases which have already been studied) to future work.

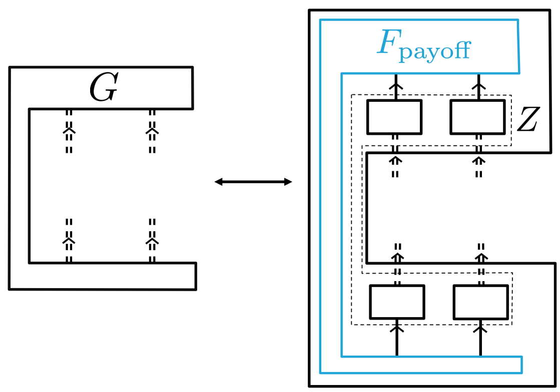

Figure 5.6: A depiction of the concrete two-step process by which a referee can implement a game (of any type). The referee first applies a tomographically complete analyzer Z, and then assigns a real number to the resulting statistics using a payoff function F_{payoff}.

5.3.2 Performance of resources of arbitrary type with respect to a game

By definition, every T-game assigns a real number to every resource of type T. At this stage, the number need not be related in any way to the nonclassicality of resources; e.g., the score need not behave monotonically under LOSR. Nonetheless, one can use any T-game to learn about the LOSR ordering of resources of type T; indeed, the full set of T-games perfectly characterizes this preorder. (In case this is not completely obvious, it will follow as a corollary of our Theorem 3.) Furthermore, one can use a T-game to (partially) quantify the nonclassicality of a resource *of arbitrary type*, not only of type T. For example, nonlocal games and semiquantum games have been used to probe the nonclassicality of quantum states [48, 245, 273, 293].

This is because—although a T-game does not *directly* assign a score to resources of any type other than T—it can quantify the performance of a resource of any type by a maximization over all $\tau \in$ LOSR that map the given resource to one of type T. That is:

Definition 4. *The (optimal)* performance *of a resource R of arbitrary type with respect to a game G_T of arbitrary type T is given by*

$$\omega_{G_T}(R) = \max_{\tau:\mathbf{Type}[R]\to T} G_T(\tau \circ R). \tag{5.1}$$

Clearly, $\omega_{G_T}(R)$ is a measure of how well an arbitrary resource R can perform at LOSR-game G_T. Because of the maximization over LOSR operations, $\omega_{G_T}(R)$ is by construction a monotone with respect to LOSR. Constructions of this sort are often termed *yield monotones* [129]. We discuss monotones further in Ref. [246], as monotones are useful tools for obtaining partial information about the preorder over resources and for relating the preorder to practical tasks.

The set $\mathcal{G}_T$ of all games of a given type T defines a preorder over all resources of all types, where resource R is above R' in the order if for every T-game, R can achieve a value at least as high as R' can. We denote this third ordering relation $\succeq_{\mathcal{G}_T}$:

Definition 5. *For resources R and R' of different and arbitrary type, we say that $R \succeq_{\mathcal{G}_T} R'$ iff $\omega_{G_T}(R) \geq \omega_{G_T}(R')$ for every $G_T \in \mathcal{G}_T$.*

Next, we prove that if one resource outperforms a second at all possible games of a given type, then it can also generate *any specific strategy* of that type that the second resource can generate. This is a nontrivial result, since it need not be the case that the first resource is higher in the LOSR order.

Theorem 2. *For resources R and R' of different and arbitrary type and a resource $\mathcal{E}_T$ of arbitrary type T, $R \succeq_{\mathcal{G}_T} R'$ iff $R' \xmapsto{\text{LOSR}} \mathcal{E}_T \implies R \xmapsto{\text{LOSR}} \mathcal{E}_T$. That is, any strategy $\mathcal{E}_T$ for games of type T that can be freely generated from R' can* also *be freely generated from R.*

Proof. If $R' \xmapsto{\text{LOSR}} \mathcal{E}_T \implies R \xmapsto{\text{LOSR}} \mathcal{E}_T$, then R can generate any strategy for any given game G_T that R' can, and so always performs at least as well as R' at T-games, and so $R \succeq_{\mathcal{G}_T} R'$.

To prove the converse, consider a set of games of type T defined by ranging over all possible payoff functions $F_{\text{payoff}}(abxy)$ for some fixed analyzer Z—that is, a specific tomographically complete measurement for each output system of the resource and a specific tomographically complete set of states for each input system of the resource. Assume that $R' \xmapsto{\text{LOSR}} \mathcal{E}_T$ for some strategy $\mathcal{E}_T$, and define $P_{Z\circ\mathcal{E}_T}(ab|xy) = Z \circ \mathcal{E}_T$. For $R \succeq_{\mathcal{G}_T} R'$, it must be that $R \xmapsto{\text{LOSR}} \mathcal{E}'_T$ for at least one strategy $\mathcal{E}'_T$ satisfying $P_{Z\circ\mathcal{E}'_T}(ab|xy) = Z \circ \mathcal{E}'_T$. If this were *not* the case, then the convex set $S(R)$ of all correlations that R can generate in this scenario, $S(R) := \{P_{Z\circ\tau\circ R}(ab|xy) = Z\circ\tau\circ R\}_{\tau\in\text{LOSR}}$, would not contain $P_{Z\circ\mathcal{E}_T}(ab|xy)$, and the hyperplane that separated $P_{Z\circ\mathcal{E}_T}(ab|xy)$ from S would constitute a payoff function F_{payoff} for which R' outperformed R, which would be in contradiction with the claim that $R \succeq_{\mathcal{G}_T} R'$. By tomographic completeness, the preimage of every correlation under Z contains at most one strategy. Hence, if two strategies map to the same correlation, then they must be the same strategy, and so it must be that $\mathcal{E}_T = \mathcal{E}'_T$ in argument above. That is, we have shown that if $R \succeq_{\text{SQ}} R'$ and $R' \xmapsto{\text{LOSR}} \mathcal{E}_T$, then $R \xmapsto{\text{LOSR}} \mathcal{E}_T$. □

5.3.3 Implications from the type of a resource to its performance at games

We now prove that games of a higher type perfectly characterize the LOSR nonclassicality of resources of a lower type.

Theorem 3. *If $T \succeq_{\text{type}} T'$, then for resources R_1, R_2 of type T', $R_1 \succeq_{\text{LOSR}} R_2$ iff $R_1 \succeq_{\mathcal{G}_T} R_2$. Equivalently: if type T is above type T', then for resources of type T', the orders defined by $\succeq_{\text{LOSR}}$ and $\succeq_{\mathcal{G}_T}$ are identical.*

Proof. Consider the set $\mathcal{G}_T$ of all games of type T and two resources R_1 and R_2, both of type T', where $T \succeq_{\text{type}} T'$. Clearly $R_1 \succeq_{\text{LOSR}} R_2$ implies $R_1 \succeq_{\mathcal{G}_T} R_2$, since $R_1 \succeq_{\text{LOSR}} R_2$ implies that R_1 can be used to freely generate R_2 and hence to generate any strategy that can be generated using R_2. Next, we prove that $R_1 \succeq_{\mathcal{G}_T} R_2$ implies $R_1 \succeq_{\text{LOSR}} R_2$. By assumption, $T \succeq_{\text{type}} T'$, and so for R_2 of type T', there exists a strategy $\mathcal{E}_T$ for games of type T such that $R_2 \xleftrightarrow{\text{LOSR}} \mathcal{E}_T$. Since $R_1 \succeq_{\mathcal{G}_T} R_2$, Theorem 2 tells us that $R_2 \xmapsto{\text{LOSR}} \mathcal{E}_T$ implies $R_1 \xmapsto{\text{LOSR}} \mathcal{E}_T$, and hence $R_1 \xmapsto{\text{LOSR}} R_2$ by transitivity. Hence we have proven that the two orderings are the same; that is, $R_1 \succeq_{\text{LOSR}} R_2$ if and only if $R_1 \succeq_{\mathcal{G}_T} R_2$. □

A consequence of this result is that if $T \succeq_{\text{type}} T'$, then every nonfree resource of type T' is useful for some T-game. Two special cases of this fact that were previously proved are that all entangled states are useful for semiquantum games and that all entangled states are useful for teleportation.

If one views the encoding of one type into another type as an embedding of the partial order over equivalence classes of resources of the lower type into the partial order of the higher type, then this result can be seen as a consequence of the fact that games of type T are sufficient for characterizing the partial order over resources of type T.

A corollary of Theorem 1 and Theorem 3 is that semiquantum games fully characterize the LOSR ordering among all resources of arbitrary type.

Corollary 1. *For any resources R and R' (which may be of arbitrary and different types), $R \succeq_{\text{LOSR}} R'$ if and only if $R \succeq_{\mathcal{G}_{\text{SQ}}} R'$.*

This generalizes the main result of Ref. [48] from quantum states to resources *of arbitrary type*. Since semiquantum games characterize the LOSR nonclassicality of arbitrary resources, and since referees in semiquantum games do not require any well-characterized quantum measurement devices [58], it follows that the nonclassicality of any resource of any type can be characterized in a measurement-device-independent manner.

Note that for such tests to be practically useful, it must be possible to convert an *unknown* resource into a semiquantum channel in the same LOSR equivalence class. This is indeed possible, because for all resources of a given type, there is a *single* transformation that implements the conversion, namely, the Bell measurement in Fig. 5.4(a). Critically, this transformation is not a function of the resource to be converted.

5.4 Extending results from the literature

We now give further applications of our results, in particular showing how our framework extends a number of seminal results from the literature.

5.4.1 Applying semiquantum games to perfectly characterize arbitrary quantum channels

Buscemi proved in Ref. [48] that the order over quantum states with respect to LOSR is equivalent to the order over quantum states defined by their performance with respect to semiquantum games. This result is an instance of our Corollary (1) where R and R' are both quantum states.

For concreteness, we now briefly reiterate the argument in this specific context. The existence of the invertible transformation in Fig. 5.4 implies that II→QQ is below QQ→CC in the order on global types, and hence that

$\forall \sigma,\ \exists\, \mathcal{E}_{\text{SQ}}\ s.t.$

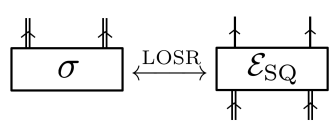

For this σ and $\mathcal{E}_{\text{SQ}}$ such that $\sigma \xmapsto{\text{LOSR}} \mathcal{E}_{\text{SQ}}$, Theorem 2 states that if $\rho \succeq_{\mathcal{G}_{\text{SQ}}} \sigma$, then $\rho \xmapsto{\text{LOSR}} \mathcal{E}_{SQ}$. Since $\mathcal{E}_{\text{SQ}} \xmapsto{\text{LOSR}} \sigma$, transitivity gives that $\rho \succeq_{\mathcal{G}_{\text{SQ}}} \sigma \implies \rho \xmapsto{\text{LOSR}} \sigma$. Since the converse implication is self-evident, one sees that the LOSR order over quantum states is equivalent to the order over quantum states defined by their performance with respect to semiquantum games.

This proof is inspired by the original argument in Ref. [48], but our framework makes the proof shorter and more intuitive. As we saw in Corollary (1), it also allowed us to generalize the result from quantum states to arbitrary resources. As stated above, this implies that the LOSR nonclassicality of *any* resource can be witnessed and quantified in a measurement-device-independent [41, 58] manner.

5.4.2 Applying measurement-device-independent steering games to perfectly characterize assemblages

Cavalcanti, Hall, and Wiseman proved in Ref. [58] that the LOSR order over quantum states defined by subset inclusion over the assemblages that each can generate via LOSR is equivalent to the order over quantum states defined by their performance with respect to steering games. This result is a special case of our Theorem 2, where R and R' are quantum states and $\mathcal{E}_T$ is a steering assemblage:

Corollary 2. $\rho \succeq_{\mathcal{G}_{\text{steer}}} \sigma$ *iff* $\sigma \xmapsto{\text{LOSR}} \mathcal{E}_{\text{steer}} \implies \rho \xmapsto{\text{LOSR}} \mathcal{E}_{\text{steer}}$.

Our Theorem 2 extends this result to arbitrary resource types and games.

Additionally, the existence of the invertible transformation in Fig. 5.4 immediately implies that

$\forall \mathcal{E}_{\text{steer}},\ \exists\, \mathcal{E}_{\text{MDI}}$ *s.t.*

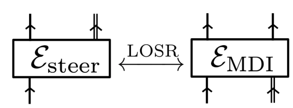

In other words, CI→CQ is below CQ→CC in the order on global types. Our Theorem 3 then gives a new result, which is the direct analogue of the result in Ref. [48] in this new context: the LOSR order over assemblages is equivalent to the order over assemblages defined by performance relative to all measurement-device-independent steering games. Explicitly: the fact (proven in Section 5.2.1) that $T_{\text{MDI}} \succeq_{\text{type}} T_{\text{steer}}$ implies that

Corollary 3. *For two assemblages* $\mathcal{E}_{\text{steer}}$ *and* $\mathcal{E}'_{\text{steer}}$*, one has* $\mathcal{E}_{\text{steer}} \succeq_{\text{LOSR}} \mathcal{E}'_{\text{steer}}$ *iff* $\mathcal{E}_{\text{steer}} \succeq_{\mathcal{G}_{\text{MDI}}} \mathcal{E}'_{\text{steer}}$.

Indeed, this theorem holds not just for assemblages, but for any resource type that is lower in the global order than measurement-device-independent steering channels, including channel steering assemblages and Bob-with-input assemblages.

5.4.3 Applying teleportation games to perfectly characterize quantum states

Cavalcanti, Skrzypczyk, and Šupić proved in Ref. [57] that the nonclassicality of every entangled state can be witnessed by some teleportation experiment. We apply arguments analogous to those of the last two subsections to strengthen their results, most notably in Corollary 4, which provides the quantitative analogue of their (qualitative) main result.

First, the existence of the invertible transformation in Fig. 5.4 again implies that

$$\forall \sigma,\ \exists\, \mathcal{E}_{\text{tel}}\ s.t.$$

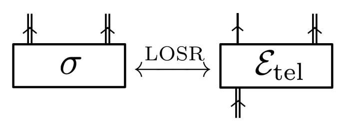

In other words, II→QQ is below QI→CQ in the order on global types. Our Theorem 3 again yields a result analogous to that in Ref. [48], namely, that the LOSR order over entangled states is equivalent to the order over entangled states with respect to performance at teleportation games[6]. Explicitly: denoting the type of quantum states by T_ρ, the fact that $T_{\text{tel}} \succeq_{\text{type}} T_\rho$ implies that

Corollary 4. $\rho \succeq_{\text{LOSR}} \sigma$ *iff* $\rho \succeq_{\mathcal{G}_{\text{tel}}} \sigma$.

Indeed, this theorem holds not just for quantum states, but for any resource type that is lower in the global order than teleportages, including, for example, steering assemblages.

Our Theorem 2 can also be applied to teleportation games, yielding a result analogous to that in Ref. [58]. That is, any resource that outperforms a second resource at all teleportation games can generate any specific strategy that the second can generate:

Corollary 5. $R \succeq_{\mathcal{G}_{\text{tel}}} R'$ *iff* $R' \xmapsto{\text{LOSR}} \mathcal{E}_{\text{tel}} \implies R \xmapsto{\text{LOSR}} \mathcal{E}_{\text{tel}}$.

[6] It is worth noting that there are subtleties in the relationship between teleportation games (as defined here, and see also Ref. [188]) and the usual conception of teleportation experiments (as attempts to establish an identity channel between two parties using shared entanglement). For example, note that any nonfree assemblage constitutes a special instance of a teleportage that is useless for generating a coherent quantum channel between two parties, and yet that *is* useful for some teleportation game.

5.5 Open questions

Our framework suggests a great deal of open questions for future study, two important examples of which were highlighted above.

Ideally, one would have type-independent methods for characterizing nonclassicality in practice. We begin developing such a toolbox in Ref. [246].

For each of the fifteen bipartite global types mentioned above, it is interesting to study the basic features of the (type-specific) LOSR resource theory. While this has been done for boxes, little attention has been given to this problem in other cases, even for quantum states.

Such features include the geometry of the free set of resources, the LOSR preorder, useful monotones and witnesses, and so on. Ultimately, we advocate not just for these type-specific investigations, but for research in the type-independent context.

Part of the project of characterizing the preorder will be to characterize the sense in which there exist inequivalent kinds of nonclassicality. At the top of the preorder, the situation for bipartite LOCC-entanglement is quite simple: there is a single maximally entangled state of a given dimension, from which all other states can be obtained by LOCC transformations. This is no longer the case for multipartite LOCC-entanglement [98], nor for LOSR-entanglement even in the bipartite case [256]. For resources beyond quantum states and for more parties, the situation gets even more complex. As an example, our work implies that there exist semiquantum channels in the equivalence class of Werner states, and semiquantum channels in the equivalence class of nonlocal boxes, and that these semiquantum channels exhibit inequivalent forms of nonclassicality.

Open Question 3. *What are the key features of the type-independent preorder over LOSR resources? What inequivalent forms of nonclassicality do these resources exhibit?*

If one were interested only in *witnessing* nonclassicality as opposed to *quantifying* it, one could consider a preorder over types defined by a less restrictive condition, where type T is above type T' if every nonfree resource of type T' could freely generate at least one nonfree resource of type T. All the known results in Table 5.1 hold for this definition as well; however, the two definitions might yield different answers for the open questions that remain.

One could also consider modifying our Definition 2 such that local operations were taken to be free rather than local operations and shared randomness. Note that the operations required in the proof of Theorem 5.4 do not make use of any shared randomness, and so the theorem would still hold. In fact, one can readily verify that all the orderings in Figure 5.1 would continue to hold. However, Theorem 2 requires convexity (through its use of the separating hyperplane theorem), as do Theorem 3 and Corollary 1 (since they rely on Theorem 2).

If one were to modify Definition 2 so that local operations and classical communication were free, the situation is less clear, as one would presumably need to widen the scope of applicability to signaling resources.

Open Question 4. *What can be learned by considering a type-independent framework of LOCC nonclassicality?*

This would be the relevant resource theory, for example, for distributed parties who share quantum memories and the ability to communicate classically.

Our framework has focused on the divide between classical and quantum resources. One can also study the divide between quantum and post-quantum resources, as we do in Ref. [254].

A final open question regards the relationship between our work and self-testing [201, 206, 291]. In self-testing, correlations (e.g. of type $\mathsf{CC}\to\mathsf{CC}$) certify the existence of an underlying valuable quantum resource (say $\mathsf{II}\to\mathsf{QQ}$). For example, the quantum correlations violating the CHSH inequality [75] maximally are a signature of an underlying quantum state that is at least as nonclassical as a singlet state (see [291] for a pedagogical derivation). Recently, the self-testing line of research has expanded beyond self-testing of states, and now has also been applied to steering assemblages [122, 292], entangled measurements [17, 237], prepare-and-measure scenarios [297], and quantum gates [266]. However, the correlations that are a signature of the given resource cannot be converted back to that quantum state, and so are not in the same LOSR equivalence class. Rather, they merely allow one to *infer the prior existence* of the self-tested resource. As such, the precise relationship with our work is left for exploration.

In the present work, we did not consider the Hilbert space dimensions as part of the resource type. One could consider a more fine-grained study of conversions between resources of different sizes. For example, the notion of nonclassical dimension for bipartite quantum states is encoded by the Schmidt rank [289]. We leave as an open question the generalization of this notion to other resource types; note that Ref. [246] includes a discussion of Hilbert space dimensions solely for the purposes of implementing numerical algorithms.

As a final remark, we recall that the semiquantum games introduced in [48] to test bipartite states in a measurement-device independent fashion [41], can be transformed into guessing games suitable for testing, always in a measurement-device independent fashion, quantum channels and quantum memories [49, 243]. More generally, such single-party guessing games have found application in the context of measurement resources [49, 276, 278] and general convex resource theories [295, 296, 302, 303]. We leave further investigations about relations between these works and ours for future research.

5.6 Conclusions

We have presented a resource-theoretic framework that unifies various types of resources of nonclassicality that arise when multiple parties have access to classical common causes but no cause-effect relations. This type-independent resource theory allows us to compare the LOSR nonclassicality of resources of arbitrary types and to quantify them using games of arbitrary types. We then derived several theorems that ultimately can be used to simplify the methods by which one characterizes the nonclassicality of resources. Our theorems additionally generalize, unify, and simplify the seminal results of Refs. [48, 57, 58], and our framework leads to a number of exciting questions for future work.

Part III.

Causal-inferential theories

Chapter 6

Unscrambling the omelette of causation and inference: The framework of causal-inferential theories

Abstract: Using a process-theoretic formalism, we introduce the notion of a causal-inferential theory: a triple consisting of a theory of causal influences, a theory of inferences, and a specification of how these interact. Recasting the notions of operational and realist theories in this mold clarifies what a realist account of an experiment offers beyond an operational account. It also yields a novel characterization of the assumptions and implications of standard no-go theorems for realist representations of operational quantum theory, namely, those based on Bell's notion of locality and those based on generalized noncontextuality. Moreover, our process-theoretic characterization of generalized noncontextuality is shown to be implied by an even more natural principle which we term *Leibnizianity*. Most strikingly, our framework offers a way forward in a research program that seeks to circumvent these no-go results. Specifically, we argue that if one can identify axioms for a realist causal-inferential theory such that the notions of causation and inference can differ from their conventional (classical) interpretations, then one has the means of defining an intrinsically quantum notion of realism, and thereby a realist representation of operational quantum theory that salvages the spirit of locality and of noncontextuality.

6.1 Introduction

One of the key disagreements among quantum researchers is the question of which elements of the quantum formalism refer to ontological concepts and which refer to epistemological concepts. The importance of settling this issue was famously noted by E.T. Jaynes [1]:

> [O]ur present [quantum mechanical] formalism is not purely epistemological; it is a peculiar mixture describing in part realities of Nature, in part incomplete human information about Nature — all scrambled up by Heisenberg and Bohr into an omelette that nobody has seen how to unscramble. Yet we think that the unscrambling is a prerequisite for any further advance in basic physical theory. For, if we cannot separate the subjective and objective aspects of the formalism, we cannot know what we are talking about; it is just that simple.

In our view, the most constructive way of defining 'realities of Nature' is as causal mechanisms acting on causal relata. That is, we here take an account of an operational phenomenon to be *realist* if it secures a *causal explanation* of that phenomenon. Hence, the particular omelette of ontology and epistemology that we will be endeavouring to unscramble in this work is the one that results from the mixing up of the concepts of *epistemic inference* on the one hand, and of *causal influence* on the other.

Scrambling of this sort is not unique to the quantum formalism—it arises also in the standard formalism for classical statistics. In that context, the difference can be characterized as follows: Bayesian inference stipulates how *learning* the value of one variable allows an agent to update their information about the value of another, while causal influence stipulates how the value of one variable *determines* the value of another (with a consequence being that an agent who *controls* the first variable can come to have some control over the second). Despite the apparent clarity of the distinction, it is often challenging to disentangle the two concepts. The statistical phenomena known as 'Simpson's paradox' [275] and 'Berkson's paradox' [39], for example, have the appearance of paradoxes *precisely because of* our tendency to inappropriately slide from statements about conditional probabilities (which merely support inferences) to statements about cause-effect relations. A satisfactory understanding of these phenomena was only found after the development of the mathematical framework of causal modeling [216, 290] that incorporated certain formal distinctions between inference and influence which are absent in the standard framework for statistical reasoning.

The conceptual difficulty of disentangling influence and inference is only compounded in the quantum realm, where the interpretation of the elements of the mathematical formalism is even less clear than it is in classical statistics.

The current chapter takes up this challenge more broadly, by pursuing the unscrambling project for two mathematical frameworks that have been used extensively in attempts to understand the conceptual content of quantum theory. The first is the framework of *operational theories*, which aims to clarify what is distinctive about quantum theory by situating it in a landscape of other possible theories, all characterized in a minimalist way in terms of their operational predictions. The second is the framework of *realist theories* (including ontological models), which has been used to constrain the possibilities for causal explanations of the operational predictions of quantum theory (and other operational theories).

We aim to recast both types of theory within a new mathematical framework that incorporates a formal distinction between inference and influence— a distinction that is lacking in previous frameworks.

A theory in our framework is termed a *causal-inferential theory*, and is constructed out of two components:

- a causal theory, which describes physical systems in the world and the causal mechanisms that relate them, and
- an inferential theory, which describes an agent's beliefs about these systems and about the causal mechanisms that relate them, as well as how such beliefs are updated under the acquisition of new information.

The full causal-inferential theory is defined by the interplay between these two components, and allows one to describe a physical scenario in a manner that cleanly distinguishes causal and inferential aspects.

Different causal-inferential theories can be obtained by varying the causal theory, varying the inferential theory, or varying both simultaneously. Note, however, that these two components are required to interact in a particular manner, so that the choice of one may be limited by the choice of the other. In this chapter, we explore, in detail, two particular choices of the causal theory and a single choice of inferential theory, as we now outline.

We take the inferential theory to be Bayesian probability theory combined with Boolean logic. Although we do not explicitly construct any alternatives to this choice in this chapter, we will have much to say about the possiblity of *nonclassical* alternatives to this inferential theory. Given that such a putative nonclassical inferential theory is the primary contrast class for us, we will refer to the inferential theory consisting of Bayesian probability theory and Boolean logic as the *classical* theory of inference.

The two types of causal theory that we consider correspond to operational and realist theories, respectively. In the first type, systems are conceptualized as the causal inputs and causal outputs of experimental procedures, and the causal mechanisms holding between such systems are simply descriptions of these experimental procedures. The causal-inferential theories that one can construct from this causal theory together with the classical inferential theory are called *operational causal-inferential theories*, and can be viewed as a refinement of the notion of operational theory described in Ref. [281] and as a competitor to the framework of Operational Probabilistic Theories [69].

In the second type of causal theory we consider, systems are classical variables and the causal mechanisms holding between these are functions. The (unique) causal-inferential theory that we construct from this is termed a *classical realist causal-inferential theory* and is a refinement of the notion of a *structural equation model* in the field of causal inference [216].[1]

In order to make a connection to other standard notions of operational and realist theories, it is useful to introduce a notion of *inferential equivalence*. Two processes are said to be inferentially equivalent if they lead one to make the same inferences whatever scenario they might be embedded within. If one quotients a causal-inferential theory with respect to the congruence relation associated to inferential equivalence, one obtains a novel type of theory, which we term a *quotiented* causal-inferential theory. Importantly, the latter sort of theory blurs the distinction between causation and inference, and hence constitutes a partial rescrambling of the causal-inferential omelette. *Generalized probabilistic theories* (GPTs) [23, 69, 136, 138], we argue, are best understood as subtheories of quotiented operational causal-inference theories[2] and consequently, unlike *unquotiented* operational causal-inferential theories, they necessarily involve some scrambling of causal and inferential concepts. We also show that *ontological models* [281] (or, more precisely, the ontological theories that are the codomain of ontological modelling maps) are best understood as subtheories of quotiented classical realist causal-inferential theories, and hence are *also* guilty of such scrambling.

Our framework leverages the mathematics of process theories [82, 124, 269], which allows it to be manifestly compositional, and consequently to apply to operational or realist scenarios with arbitrarily complex causal and inferential structure. Many previous frameworks for operational theories [69, 138] have also availed themselves of the mathematics

[1] Although such frameworks achieved significantly more unscrambling of the causal-inferential omelette than the statistical frameworks that preceded them—as noted above in our discussion of statistical 'paradoxes'—our novel framework achieves some further unscrambling.

[2] This view of GPTs as quotiented operational causal-inferential theories is closely related to the quotienting of operational probabilistic theories of [69].

of process theories to allow the representation of arbitrarily complex structures. These did not, however, explicitly distinguish the structures that are causal and those that are inferential, as we do here. Our use of the mathematics of process theories represents more of an innovation on the realist side, since previous frameworks for realist theories focused almost exclusively on the simple structures that arise when describing prepare-measure scenarios, sometimes with an intervening transformation or sequence of transformations [181, 186, 192, 228, 231]. (These frameworks also did not distinguish structures that are causal from those that are inferential.)

One of the motivations for the standard framework for ontological models was to answer the question of whether the predictions of a given operational theory admit of an explanation in terms of an underlying ontology. The counterpart of this question in our new framework is whether the predictions of a given *operational causal-inferential theory* admit of an explanation in terms of an underlying *classical realist causal-inferential* theory. Such an explanation is deemed possible if the former can be *represented* in terms of the latter. We refer to this as a *classical realist representation* of an operational causal-inferential theory.

The key constraint we impose on such representations is that they preserve the causal and inferential structures encoded in the diagrams, a property that is formalized by demanding that the map between the two process theories (operational and classical realist) is *diagram-preserving* [260]. We show that this constraint involves no loss of generality in terms of the sorts of realist models of experimental phenomena one can describe in the framework. Moreover, in concert with standard hypotheses about the causal and inferential structure, it acts as an umbrella principle which subsumes many principles that have previously been used to derive no-go theorems for classical realist representations of operational quantum theory [260].

In particular, in the case of a Bell scenario, the assumption of diagram preservation subsumes the causal and inferential assumptions that go into deriving Bell inequalities (when this derivation is conceptualized in terms of causal modeling [324]). However, it is much more general than this, and subsumes the causal and inferential assumptions that go into deriving Bell-like inequalities (also known as *causal compatibility inequalities*) for scenarios that have a causal structure distinct from the Bell scenario [63, 104, 106, 216, 323, 324]. Our framework therefore also constitutes a refinement of (or alternative to) recently proposed frameworks [216, 323] for identifying causal compatibility constraints in such scenarios.

We also demonstrate how a principle proposed by Leibniz and used extensively by Einstein [285] can be generalized in a natural way to apply to theories incorporating epistemological claims in addition to ontological claims and that this principle implies a formal constraint on realist representations of an operational causal-inferential theory that we term *Leibnizianity*: the representation must *preserve inferential equivalences*. We also demonstrate that the principle of Leibnizianity implies a rehabilitated version of the principle of generalized noncontextuality [281], such that no-go theorems for generalized-noncontextual classical realist representations of operational quantum theory imply no-go theorems for Leibnizian classical realist representations. The question of whether the reverse implication holds remains open. We also discuss the connections between this principle and the old version of generalized noncontextuality.

Bell's no-go theorem is understood in our framework as follows. If quantum theory is cast as an operational causal-inferential theory, then it predicts distributions for certain causal-inferential structures that cannot be realized by a classical realist causal-inferential theory with the same causal-inferential structure. Meanwhile, noncontextuality no-go theorems are understood in our framework similarly, but where one demands that inferential equivalences *as well as* the causal-inferential structure are preserved.

The conventional realist responses to the standard no-go theorems are unsatisfactory in various ways, such as requiring superluminal causal influences, requiring fine-tuning, and running afoul of Leibniz's principle. In light of this, it has been suggested that a more satisfactory way out of these no-go theorems may be achieved by modifying the notion of a realist representation (see, e.g., Sec. 1.C of Ref. [284]). This has been described in past work as 'going beyond the standard ontological models framework', but here is understood as seeking a nonclassical generalization of the notion of a classical realist representation. Our process-theoretic framework provides the formal means of achieving this because it allows the interpretation of causal and inferential concepts to be determined by the axioms of the process theories that describe them and hence to differ from the conventional, classical interpretations of these concepts. This is analogous to how, in nonEuclidean geometries, the concepts of point and line acquire novel meanings distinct from their conventional ones. Success in such a research program consists in finding an intrinsically quantum notion of a realist causal-inferential theory which can provide a Leibnizian representation of operational quantum theory. We propose natural constraints on the axioms describing a theory of causal influences, a theory of epistemic inferences, and their interactions. We also point to pre-existing work that offer clues for how to proceed.

Thus, the work we present here provides a significant step forward in this research program. On the one hand, it provides, for the first time, a concrete proposal for the mathematical form of the sought-after theory, and, on the other hand, it provides a set of ideas for the form of its axioms, thereby providing a road map for future research.

6.1.1 Process theories and diagram-preserving maps

We formalize the ideas discussed in the introduction using the mathematical language of process theories and diagram-preserving maps. This section serves as a brief introduction to these concepts. Process theories provide a mathematical framework for describing an extremely broad class of theories, finding utility both in physics [127, 269] and beyond [18, 83, 84]. They can ultimately be given a category-theoretic foundation, but in this chapter we will require only the diagrammatic approach. We do provide some "category-theoretic remarks" when a given definition or result can be expressed concisely in this language, but these remarks can be skipped without impacting the comprehensibility of the rest of the chapter.

As the name would suggest, the basic building blocks of a process theory are processes. These could correspond to physical processes in the world, but equally well could apply to computational processes, mathematical processes, etc. In our work, we will focus on causal and inferential processes.

Definition 2 (Processes)**.** *A* process *is defined as a labeled box with labeled input and output* systems*, e.g.:*

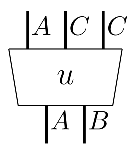

. (6.1)

The label of a system, e.g., A, is known as the type *of the system while the label on the box, e.g., u, is simply the* name *of the process.*

Note that it is allowed for a process to have no input systems and/or no output systems. Processes with no inputs are called states*, those with no outputs are called* effects*, and those with neither inputs nor outputs are simply called* closed diagrams *(or sometimes* scalars*). That is,*

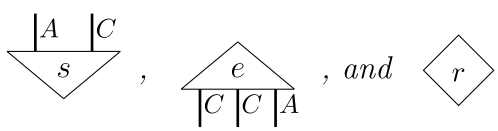

(6.2)

are examples of a state, an effect and a closed diagram respectively.

Before we can define a process theory we must introduce the idea of a diagram of processes.

Definition 3 (Diagrams)**.** *A* diagram *is defined as a 'wiring together' of a finite set of processes—that is, an output of one process is connected to the input of another, such that the system types match and no cycles are created. For example,*

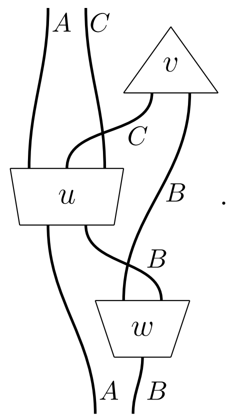

(6.3)

Only the connectivity *of the diagram— which systems are wired together and the ordering of the input and output systems—is relevant. That is, two diagrams are the same if one can be deformed into the other while preserving this connectivity. For example,*

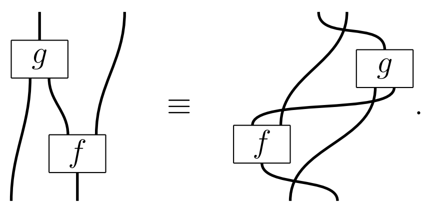

(6.4)

We will now formally define what we mean by a process theory.

Definition 4 (Process theories)**.** *A process theory is defined as a collection of processes,* T*, which is closed under forming diagrams. For example, we can draw a box around the* *above diagram* (6.3) *and view it as another process in the theory. That is,*

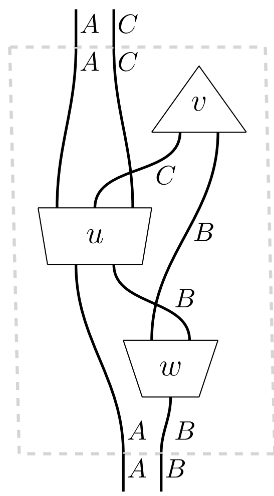

$\in$ T (6.5)

for $u, v, w \in \mathrm{T}$.

This completes the definition of a process theory. However, it is sometimes useful to introduce elementary notions from which any diagram can be built up. To start, we highlight certain elements of the diagrams by picking them out with dashed boxes below:

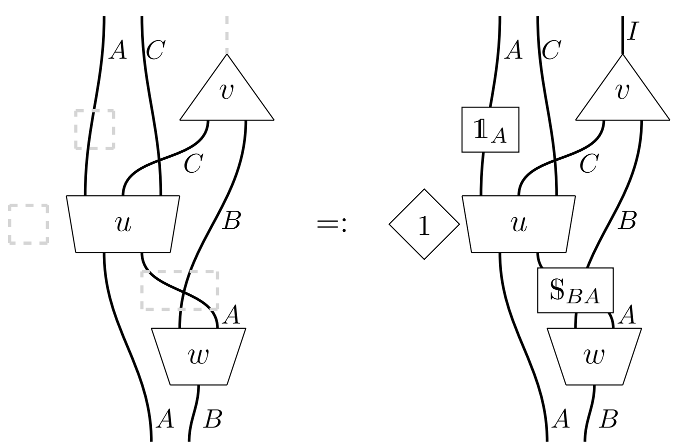

(6.6)

That is: i) one can view the empty box on the left as a special closed diagram (as it has no input and no output) which we refer to as the scalar 1; ii) one can view the box with the A wire running through it as an identity process $\mathbb{1}_A$; iii) one can view the box with the crossed A and B wires as a swap process $\mathbb{S}_{BA}$; and iv) one can view the output system of v as the trivial system I. Clearly our diagrammatic notation implies constraints on these extra elements. In particular, wiring identity processes onto any other process leaves that process invariant; the composite of a trivial system with any other system is just that system; swapping twice is the identity on the two systems; and finally, composing a process with the scalar 1 leaves that process invariant. Together with the elements just introduced, one can introduce two elementary notions of composition from which any diagram can be built up: sequential composition of processes, denoted

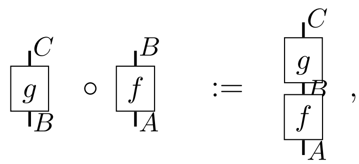

(6.7)

and parallel composition of processes, denoted

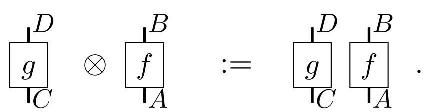

(6.8)

Note that, given any diagram, there are a (generally infinite) number of ways in which it can be expressed in terms of these primitive notions of composition, and yet these are all the same diagram. Hence, we view the diagrammatic representation as being the fundamental description, and we view the elementary notions from which they can be built as an (at times) convenient mathematical representation of them.

Remark 4. *Having defined these extra structures implicitly in the diagrammatic notation, it should be clear how to identify the structure of a process theory with that of a* strict symmetric monoidal category *(SMC). In short, we take processes to be morphisms and systems to be objects, with sequential and parallel composition providing morphism composition and the monoidal product, respectively. For a more formal treatment, see Ref. [214].*

We will often consider higher-order processes such as

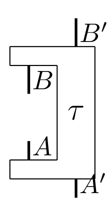

, (6.9)

which we will call *clamps*. These can be thought of as objects that map a process from A to B to a process from A' to B' via

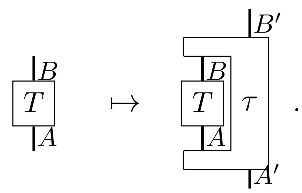

(6.10)

These are not primitive notions within the framework of process theories, but rather are constructed out of two processes x_τ and y_τ connected together with an auxiliary system W_τ, as

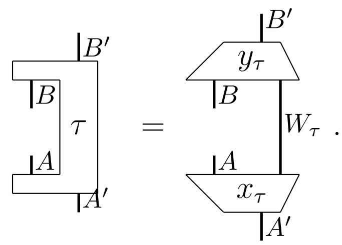

(6.11)

We will consider a number of different process theories, some of which are sub-process-theories of others.

Definition 5. *A sub-process-theory $T \subseteq T'$ is a process theory where the processes are a subset of the processes in T' and composition of processes in T is given by composition in T'. Note that since a sub-process-theory is a process theory, T must be closed under composition.*

Remark 5. *In terms of the associated SMCs, this is simply defining T as a subSMC of T'.*

As well as these process theories and sub-process-theories, we will also consider structure-preserving maps between these. The relevant structure which we demand be preserved is the composition of processes as described by diagrams.

Definition 6 (Diagram-preserving maps)**.** *A diagram-preserving map, $\mathbf{m} : \mathrm{T} \to \mathrm{T}'$, is a map from processes in T to processes in T' such that wiring together processes in T to form a diagram and then applying the map $\mathbf{m}$ is the same as applying $\mathbf{m}$ to each of the component processes and then wiring them together in T'. We depict these maps as shaded* *boxes, e.g.*

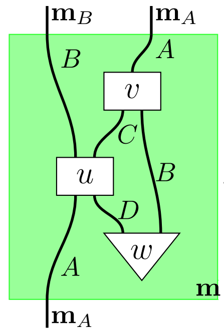

, (6.12)

where the diagram in the green box is a diagram in T *(with input* A *and outputs* B *and* A*) which is mapped by the green box* $\mathbf{m}$ *to a process in* T' *with input* $\mathbf{m}_A$ *and outputs* $\mathbf{m}_A$ *and* $\mathbf{m}_B$*. In this example, the constraint of diagram preservation is simply that*

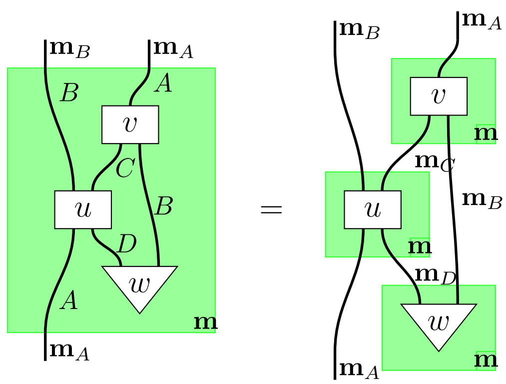

. (6.13)

Remark 6. *If one views each process theory as an SMC, then such diagram-preserving maps are simply* strict symmetric monoidal functors *between the SMCs. The diagrammatic notation as shaded regions was introduced to us by Ref. [110], which was itself was based on Ref. [204].*

It will also be useful to consider partial diagram preserving (DP) maps where the domain is limited in scope.

Definition 7 (Partial diagram-preserving maps)**.** *A partial diagram-preserving map* $\mathbf{m} : \mathrm{T}' \to \mathrm{T}''$ *is a diagram-preserving map from some sub-process-theory* $T \subseteq \mathrm{T}'$ *to* T''*.*

Remark 7. *Categorically, such a map is a partial strict symmetric monoidal functor between the relevant SMCs.*

Remark 8 (Category of process theories, ProcessTheory)**.** *The category of process theories is defined as follows: The objects of* ProcessTheory *are process theories and the morphisms are diagram-preserving maps. It is simple to see that this is indeed a category as one can easily define identity morphisms and morphism composition satisfying the relevant conditions.*

6.2 Causal primitives

We denote a generic process theory of causal relations by Caus. The primitive elements of such a theory are systems and the causal mechanisms that relate them. Systems correspond to physical degrees of freedom in the conventional sense of being the loci of causal relations, i.e., the causal relata. Causal mechanisms are autonomous physical relationships between these systems, relationships governed by the laws of nature and by the arrangement of relevant physical systems and apparatuses.

In classical theories, systems are often represented by sets, and causal mechanisms by functions between these (e.g., in structural equation models [216]). In quantum theory, systems are typically represented by Hilbert spaces, and causal mechanisms by unitaries between these [12][3]. In operational causal-inferential (CI) theories, one typically does not have a direct description of the causal mechanisms, but rather only a very coarse-grained description of them in terms of laboratory procedures that are implemented on the relevant systems; these systems are represented only as an abstract label, and typically represent the physical systems one imagines are input and output from the apparatuses.

In the next two sections, we will consider two distinct classes of causal primitives in more detail, first those relevant for operational CI theories, and then those relevant for classical realist CI theories. In Section 6.9.1, we return to the question of what properties any process theory must satisfy for it to be considered a good theory of causal relations.

6.2.1 Process theories of laboratory procedures

We now define the sort of process theory that will ultimately constitute the causal component of an operational CI theory. We denote it Proc (as shorthand for 'procedure' not 'process'). The systems in Proc label the primitive causal relata, while the processes,

[3] Although how to decompose a given unitary with multiple outputs and a given internal causal structure remains an open problem [12, 190].

which describe the potential causal relations between them, are labeled by laboratory procedures, conceptualized as a list of instructions of what to do in the lab, and presumed to be individuated by the system they act on (the input system) and the system they prepare (the output system).

We will label general systems by A, B, etc. Some systems (which could, for instance, represent setting or outcome variables) will be deemed *classical*. A classical system X will be associated with a set X, which represents the set of distinct states of the classical system. Diagrammatically, a general laboratory procedure t with input system A and output system B will be drawn as

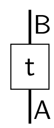

. (6.14)

We define a measurement m on system A as a process with a generic input system A and a classical output system X:

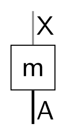

. (6.15)

Classical systems in Proc will be drawn with a light grey wire, as was done in this diagram. We denote the set of operations with input system A and output system B as $\overline{\mathsf{A}\to\mathsf{B}}$. The set of measurements on a system A having outcome space X is therefore denoted $\overline{\mathsf{A}\to\mathsf{X}}$. (Classical systems allow one to describe more than just measurement outcomes; for instance, they can also represent classical control systems.)

One can compose these operations to describe experiments. For example, a preparation procedure on system A followed by a measurement on A with outcome space X is described by the diagram

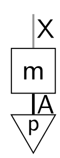

, (6.16)

while controlling a transformation from B to C on the output X of a measurement on A would be described by the diagram

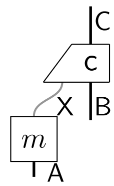

. (6.17)

An example of a more general diagram is

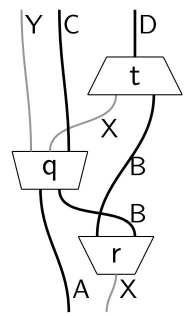

. (6.18)

In our formalism, such a diagram represents a hypothesis about the fundamental causal structure. In Appendix C.2 we discuss the consequences of this choice and how it differs from the choice typically made in operational frameworks. Here, two wires in parallel should be interpreted as independent subsystems (e.g., independent degrees of freedom), where one can talk independently about either subsystem as a potential causal influence on other systems.

Note that we have here defined a *class* of process theories, insofar as we have not specified the *particular* set of systems and procedures that define Proc. Perhaps the most common operational theory to consider is that containing all known physical systems and laboratory procedures on them. One might also consider a restriction of this set, for example, the set of two-level systems and laboratory procedures on them. Finally, one might consider a foil operational theory [71], with a set of hypothetical systems and procedures on them. Each possible choice for Proc defines a different operational CI theory.

Remark 9. *Unlike the other process theories that we deal with in this chapter,* Proc *is a* free *process theory. This means that there are no equalities other than those defined by the framework of process theories—two diagrams are equal if and only if they can be transformed into one another by sliding the processes around on the page while preserving the wiring.*

6.2.2 Process theory of classical functional dynamics

We now define the process theory that will ultimately act as the causal component of our notion of a classical realist causal-inferential theory. We denote it Func.

The systems in Func again label the primitive causal relata. However, what distinguishes them from the systems in Proc is that we assume that these relata are described by ontic state spaces, that is, some finite[4] sets $\Lambda, \Lambda', \ldots$. The processes in Func are functions $f : \Lambda \to \Lambda'$ describing dynamics on these ontic state spaces.[5]

Diagrammatically, a function f with input Λ and output Λ' will be drawn as

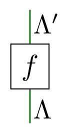

. (6.19)

We denote the set of functions with input Λ and output Λ' as $\overline{\Lambda \to \Lambda'}$. We take the trivial system to be the singleton set $\star = \{*\}$. Hence, states correspond to functions $s : \star \to \Lambda$ that are in one-to-one correspondence with the elements of Λ; there is a unique effect $u : \Lambda \to \star$ for each system, defined by $u(\lambda) = *$ for all $\lambda \in \Lambda$; and, there is a unique scalar $1 : \star \to \star$ corresponding to the identity function on the singleton set.

We can compose these functions to describe ontological scenarios. A function that prepares some ontic state of system Λ followed by a function describing the functional dynamics of the system is described by the diagram

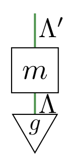

, (6.20)

---

[4]This assumption of finiteness is made for simplicity of presentation, but could be removed in future work.

[5]Note that we will allow arbitrary functions in this chapter, although in some cases one might wish to restrict the dynamics, e.g., to reversible functions or to symplectomorphisms.

while some more general ontological scenario could be described by the diagram

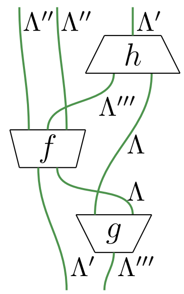

. (6.21)

The key formal distinction between Func and Proc is that Func is *not* a free process theory. Indeed, there are many nontrivial equalities provided by composition of functions. An example is provided by Eq. (6.20), the diagram, which is made up of two functions $g : \star \to \Lambda$ and $m : \Lambda \to \Lambda'$, is *strictly equal* to the diagram

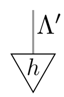

, (6.22)

where h is the sequential composition of m and g, i.e., $h : \star \to \Lambda' :: * \to m(g(*))$.

Composite systems are given by the Cartesian product of the associated sets:

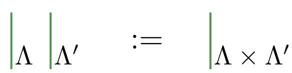

. (6.23)

Parallel composition of functions is given by their Cartesian product,

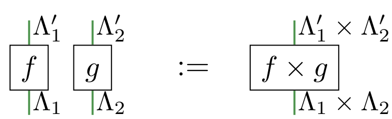

, (6.24)

and sequential composition is given by composition of functions,

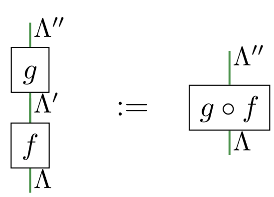

. (6.25)

It follows that any diagram is equal to the function from its inputs to its outputs, and that one can compute this effective function directly from the specific functions which comprise the diagram.

Remark 10. *Categorically we are simply defining the symmetric monoidal category* FinSet *whose objects are finite sets, morphisms are all functions between them, and the monoidal structure is given by the cartesian product and the singleton set.*

6.3 Inferential primitives

We denote a generic process theory of inference by Inf. The primitive inferential notions in our framework are systems and the processes which specify what an agent knows about them and how the agent reasons—for example, how they update what they know about one system given new information about another. When the inferential theory is defined intrinsically, the systems within the theory are simply understood as the entities about which one has states of knowledge and about which one asks questions, regardless of what these entities are: one could be making inferences about physical systems, or about mathematical truths, and so on. When an inferential theory is considered as a part of a causal-inferential theory, however, the entities about which one makes inferences are causal mechanisms and the causal relata that these act on.

The processes in Inf are particular inferences (i.e., updates of the knowledge one has about one system given new information about another), while the rewrite rules in Inf encode the laws of inference that an agent should follow if they are to be rational. The standard laws of inference are those of Bayesian probability theory and Boolean logic. We formalize the laws of inference that will be relevant in this chapter within a single process theory, namely SubStoch, the process theory of substochastic maps. SubStoch is the only explicit example of an inferential process theory that we will consider in this chapter[6]. In Section 6.9.1, we return to the question of what properties any process theory must satisfy for it to be considered a good process theory of inference.

To diagrammatically distinguish the causal structure encoded in a given diagram of Caus from the inferential structure encoded in a given diagram of Inf, we draw diagrams in the former vertically, and diagrams in the latter horizontally. We will term the systems in the former *causal systems*, and systems in the latter *inferential systems*.

[6] However, another good example that one could consider is the category Rel of finite sets and relations. This would describe possibilistic reasoning as opposed to the probabilistic reasoning of SubStoch.

6.3.1 Bayesian probability theory

The first component of SubStoch is (classical) Bayesian probability theory, describing an agent's states of knowledge and the updating thereof. We denote this process theory Bayes.

Systems in Bayes are represented by finite sets $X, Y, \ldots$. A process from X to Y in this theory is interpreted as the propagation of an agent's knowledge about X to her knowledge about Y, and is represented by a stochastic map. Such a process will be depicted diagrammatically as

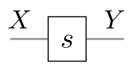

. (6.26)

The trivial system is the singleton set, $\star := \{*\}$, and a map σ from the trivial system to system X, depicted as

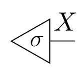

, (6.27)

corresponds to a probability distribution over X. We will denote the point distribution $\delta_{X,x}$ on some element $x \in X$ as $[x]$. There is a unique effect for each system that corresponds to marginalisation over the variable. This is drawn as

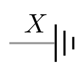

. (6.28)

As this is the unique effect, it is clear that any closed diagram is associated with the number 1, which is the unique scalar in Bayes. It follows that given a state of knowledge σ about a pair of variables, X and Y, one can define the marginal distribution on the variable X as

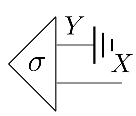

. (6.29)

Note that it is possible to take convex combinations of stochastic processes provided that they have matching system types. We denote a convex combination of stochastic processes $\{s_i : X \to Y\}$ with weights $\{p_i\}$ as

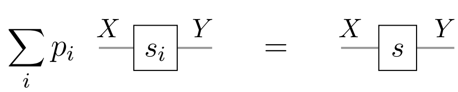

, (6.30)

where $s : X \to Y$ is another stochastic process in the theory, namely, $s = \sum_i p_i s_i$.

Remark 11. *Categorically,* BAYES *is simply the symmetric monoidal category* FINSTOCH *where objects are finite sets and morphisms are stochastic maps between them.*

6.3.2 Boolean propositional logic

The second component of SUBSTOCH is the description of propositions, as governed by (classical) Boolean propositional logic. We denote the process theory describing this as BOOLE. The systems in BOOLE are finite sets labeled by $X, Y, ...$, just as in BAYES. However, the processes in BOOLE are partial functions; that is, functions that may only be defined on a subset of their domain.

Many of the key processes in BOOLE are simply functions; we introduce these first, and only later discuss the more general processes in BOOLE for which partial functions are required.

First, we consider states in BOOLE: functions from the trivial system $\star$ to a generic system X. These are in one-to-one correspondence with the elements of X, and so we can simply label each function by the element x that is the image of $*$ under it. Hence, we can write $x : \star \to X :: * \mapsto x$, depicted diagrammatically as

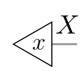

. (6.31)

We will refer to states in BOOLE as **value assignments**, because they can naturally be viewed as assigning a value x to the variable ranging over the set X.

Yes-no questions about a system X can be represented as functions from X to the answer set $\mathrm{B} := \{\mathrm{Y}, \mathrm{N}\}$, we refer to this answer set B as the **Boolean system**. Diagrammatically, these yes-no questions are denoted by

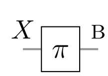

. (6.32)

A value assignment $x \in X$ assigns the answer 'yes' or 'no' to such questions via composition:

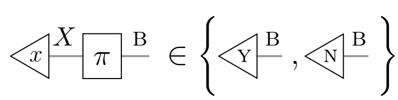

. (6.33)

Now, each of these yes-no questions can be uniquely characterized by the subset of X for which the answer is 'yes', that is, $\{x \in X \text{ s.t. } \pi :: x \mapsto \text{Y}\}$. This means that they are in one-to-one correspondence with the elements of the powerset of X which we will view as a Boolean algebra, and, hence which we will denote by $\mathcal{B}(X)$. Due to this correspondence, we will use the symbol π both to denote the function $\pi : X \to \text{B}$ and the element of the Boolean algebra $\pi \in \mathcal{B}(X)$. We therefore refer to such functions as **propositional questions**[7].

We now show how the structure of the Boolean algebra $\mathcal{B}(X)$ can be diagrammatically represented using such propositional questions. To begin, there are two distinguished propositions in the Boolean algebra, the tautological proposition $\top$ (corresponding to the subset $X \subseteq X$) and the contradictory proposition $\perp$ (corresponding to the empty subset $\emptyset \subseteq X$). As propositional questions these can be defined diagrammatically via

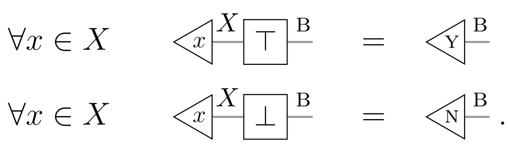

(6.34)

(6.35)

Moreover, our representation of propositions as functions in BOOLE allows us to diagrammatically represent unary and binary logical operations on the propositions, by defining them in terms of their action on Boolean systems. For example, the NOT operation on an arbitrary proposition can be represented as

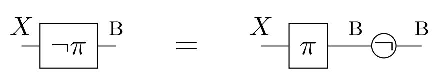

(6.36)

where the dot decorated by the $\neg$ symbol is the function $\neg : \text{B} \to \text{B}$ whose action on the Boolean system reflects the truth table of the logical NOT, namely $\neg(\text{Y}) = \text{N}$ and $\neg(\text{N}) = \text{Y}$ (implying that $\neg$ is self-inverse). Similarly, one can represent the logical OR operation (disjunction), denoted $\vee$, as

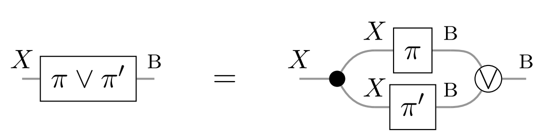

(6.37)

where the black dot is the copy function $\bullet : X \to X \times X$ defined by $\bullet(x) = (x, x)$, and the dot decorated by the $\vee$ symbol is the function $\vee : \text{B} \times \text{B} \to \text{B}$ whose action on the Boolean

[7] These could also have been termed 'predicates'.

system reflects the truth table of the logical OR, namely, $\mathrm{Y} \vee \mathrm{Y} = \mathrm{Y} \vee \mathrm{N} = \mathrm{N} \vee \mathrm{Y} = \mathrm{Y}$ and $\mathrm{N} \vee \mathrm{N} = \mathrm{N}$. In a similar manner, one can construct representations of the logical AND, logical implication, exclusive OR, etc.

We have therefore diagrammatically represented the basic operations required to define a Boolean algebra $\mathcal{B}(X)$. Moreover, the basic properties of a Boolean algebra (associativity, absorption, commutativity, identity, annihilation, idempotence, complements, and distributivity) can also be shown to hold. These are defined and proven in Appendix C.3.

One can also express propositional questions about composite systems, for example

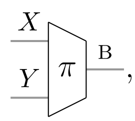

(6.38)

where $\pi \in \mathcal{B}(X \times Y)$. We discuss in Appendix C.3 how these can be constructed out of single-system propositions.

Arbitrary functions are also valid processes in Boole, since (as we now show) each is a valid *Boolean algebra homomorphism*, that is, a map between Boolean algebras that preserves the logical connectives $\wedge$ and $\vee$ as well as the top and bottom elements $\top$ and $\bot$. Consider a general function

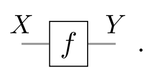

(6.39)

In the process theory Bayes, this process would be viewed as a stochastic map acting on the right to take states of knowledge about X to states of knowledge about Y. Within Boole, however, this process is viewed as a map acting on the left, taking a propositional question about Y to a propositional question about X via

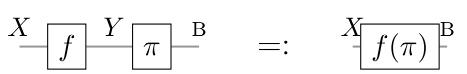

(6.40)

We now show that $f(\_)$ defines a Boolean Algebra homomorphism from $\mathcal{B}(Y)$ to $\mathcal{B}(X)$, where each subset of Y is mapped to the subset of X which is the preimage under $f(\_)$ of the subset of Y. We will sometimes refer to such generic functions simply as **propositional maps**. It is easy to see that propositional maps do indeed preserve $\wedge$, $\vee$, $\top$, and $\bot$. The fact that the propositional map $\top : Y \to \mathrm{B} :: y \mapsto \mathrm{Y}$ is preserved follows immediately from the fact that $f(\_)$ maps every x to *some* y, which $\top$ then necessarily maps to Y. Preservation of $\bot$ is analogous. To see that $\wedge$ is preserved—namely, that $f(\pi \wedge \pi') = f(\pi) \wedge f(\pi')$—note that

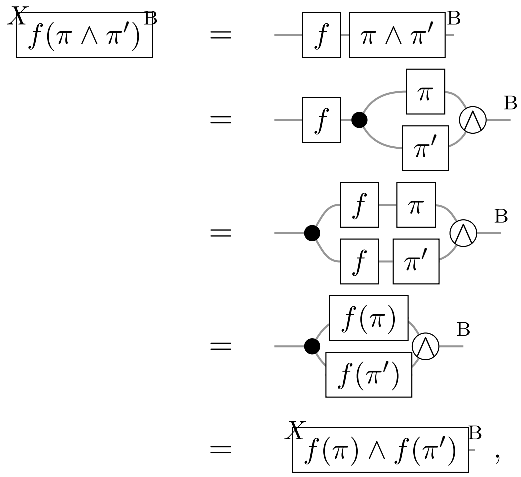

(6.41)

(6.42)

(6.43)

(6.44)

(6.45)

where the equality between Eq. (6.42) and Eq. (6.43) simply states that copying the output of a function is the same as copying the input and applying the function to the two copies of the input. The proof of preservation of $\vee$ is analogous.

Remark 12. *Categorically, this dual picture of $f(\_)$ as a function from X to Y and as a Boolean Algebra homomorphism from $\mathcal{B}(Y)$ to $\mathcal{B}(X)$ is the duality between the categories* FinSet *and* FinBoolAlg *where $B = \{\text{Y}, \text{N}\}$ is the dualising object.*

To express the truth value assigned to a proposition, we must introduce scalars and effects, which moreover requires us to go beyond functions and consider partial functions[8]. A partial function $\hat{f} : X \to Y$ is a function from some (possibly empty) subset $\chi_{\hat{f}} \subseteq X$ to Y; the partial function is simply undefined on the elements of X outside of this subset.

There are two scalars in Boole, which we identify with true and false; namely, the function $\mathsf{True} : \star \to \star :: * \mapsto *$ and the partial function $\mathsf{False} : \star \to \star$, which is defined only on the empty set $\emptyset$, that is, $\chi_{\mathsf{False}} = \emptyset$. The scalar True is depicted by the empty diagram, since composing it with any other process leaves that process invariant:

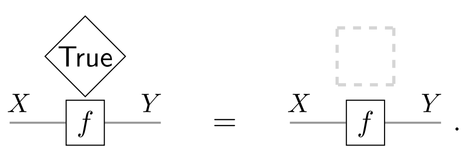

(6.46)

[8]This is because in the process theory of functions, there is a unique scalar and a unique effect. The unique scalar is the function taking the singleton set to itself, while the unique effect is the function from X to the singleton set which maps every element $x \in X$ to $*$.

The scalar False behaves as a 'zero scalar', in the following sense. Defining the 'zero process' for a pair of systems (X, Y) as the unique partial function $\mathbf{0} : X \to Y$ such that $\chi_{\mathbf{0}} = \emptyset$, it follows that composing any other partial function $\hat{f} : X \to Y$ with False will give this zero process:

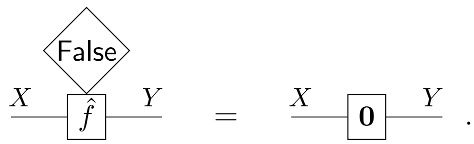

(6.47)

Now, effects within BOOLE are partial functions taking $X \to \star$, diagrammatically denoted by

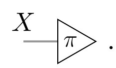

(6.48)

We will see that these are also in one-to-one correspondence with the elements of the Boolean algebra $\mathcal{B}(X)$, which justifies our labelling them by propositions π. To see this, first note that value assignments $x \in X$ assign a truth-value to effects within BOOLE when the two are composed together:

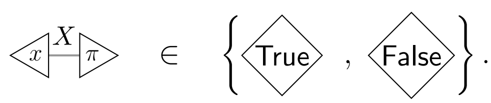

(6.49)

Hence, one can uniquely associate a partial function with the subset of X for which we obtain True; indeed, this subset is the domain χ_π of the effect, viewed as a partial function. We will call such partial functions **propositional effects**.

At this point, we have three uses of the symbol π: we have $\pi \in \mathcal{B}(X)$ as a subset (an element of a Boolean Algebra), $\pi : X \to \text{B}$ as a propositional question, and now $\pi : X \to \star$ as a propositional effect. The distinction between these should be clear from context.

To more explicitly see the connections between propositional effects and propositional questions, let us consider the special case of propositional effects for the Boolean system B. There are four of these, corresponding to the four subsets of B on which the partial function from B to $\star$ can be defined, namely $\{\text{Y}\}$, $\{\text{N}\}$, $\{\text{Y}, \text{N}\}$, and $\emptyset$.

We denote these effects, respectively, as

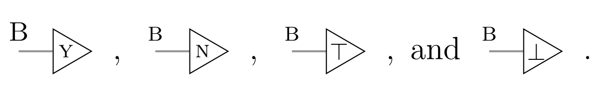

(6.50)

Then, we can write a given propositional effect in terms of the associated propositional question via

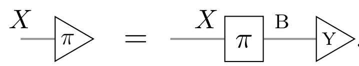

(6.51)

Value assignments to propositional effects are then consistent with value assignment to propositional questions, in the sense that[9]

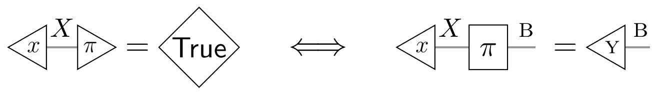

(6.52)

It turns out that *all* partial functions can be generated from the elements we have introduced so far, and hence all of these are in Boole. An arbitrary partial function $\hat{f}$ can be written as

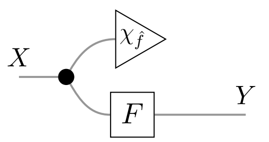

(6.53)

where $\chi_{\hat{f}}$ is an arbitrary propositional effect, $F : X \to Y$ is an arbitrary propositional map, and the black dot is the copy function. Here, the top part of the diagram defines the subset of the domain on which the partial function is defined, and then F defines the action of the partial function on that domain. Hence, we have that

Remark 13. *Categorically, the process theory* Boole *is the symmetric monoidal category* $\text{FinSet}_{\text{part}}$ *where objects are finite sets and morphisms are partial functions.*

We discussed how the functions in Boole correspond to Boolean algebra homomorphisms. In contrast, partial functions in Boole are more general. In general, partial functions do not map propositional questions to propositional questions (via Eq. (6.40)), but rather take them to other partial functions. However, partial functions do map propositional effects to propositional effects via

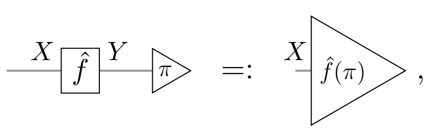

(6.54)

and so we can ask which structures of the Boolean algebra of propositional effects are preserved by such a map. We show in Appendix C.3 that $\perp$, $\vee$ and $\wedge$ are preserved, but

[9] It is worth noting that one could dispense with the notion of propositional questions and express all claims in terms of propositional effects. We include the notion of a propositional question here because it helps to clarify certain conceptual distinctions.

$\top$ and $\neg$ are not. We then show that a partial function from X to Y corresponds to a Boolean Algebra homomorphism from $\mathcal{B}(Y)$ to $\mathcal{B}(\chi_{\hat{f}})$.

6.3.3 The full inferential process theory

We now show how the probabilistic and the propositional parts of the inferential theory interact—for example, allowing one to compute the probability one should assign to a propositional effect on a system X, given an arbitrary state of knowledge about X. This interaction is possible because Bayes and Boole define a collection of processes on the same types of systems. However, the processes in Bayes are stochastic maps, while those in Boole are partial functions, and so it remains to define how these interact with one another.

We proceed by showing that both Bayes and Boole can be faithfully represented within the process theory of substochastic linear maps, SubStoch (see Remark 14). More formally, there is a diagram-preserving inclusion map from Bayes into SubStoch, and there is a diagram-preserving map (given by Eq. (6.57) below) from Boole into SubStoch. Hence, we have

$$\text{Bayes} \rightarrowtail \text{SubStoch} \leftarrowtail \text{Boole}. \tag{6.55}$$

Moreover, we show that any substochastic map can be realised by composing the processes from these two representations.

First, let us consider Boole. Note that any function $f : X \to Y$ can be represented by an associated stochastic map, $(\mathbf{f}_x^y)_{x\in X}^{y\in Y}$, via:

$$\mathbf{f}_x^y = \begin{cases} 1 & \text{if } f(x) = y \\ 0 & \text{otherwise.} \end{cases} \tag{6.56}$$

Any such stochastic map is *deterministic*, meaning that there is precisely a single 1 in each column, with the rest of the elements are 0. It is not difficult to check that these stochastic maps compose in the same way as the underlying functions, and so this gives us a representation of the process theory of functions within the process theory of stochastic maps.

More generally, a partial function $\hat{f}$ is also associated with a stochastic map via

$$\hat{\mathbf{f}}_x^y = \begin{cases} 1 & \text{if } \hat{f}(x) = y \\ 0 & \text{otherwise,} \end{cases} \tag{6.57}$$

with the only difference being that for functions, the 0 case occurs only when $f(x) \neq y$, whereas for partial functions it will also occur when $\hat{f}$ is not defined on x. These are more general than deterministic stochastic maps in that some of the 1s may be replaced by 0s—that is, they are *substochastic* maps. Again, one can check that the representative substochastic maps compose in the same way as the underlying partial functions.

Clearly, the probabilistic part of the theory, which is described by stochastic maps, can also be represented as substochastic maps, as the former are simply a special case of the latter.

Hence, both Bayes and Boole can be represented within SubStoch.

Within this representation, certain processes in Bayes and certain processes in Boole correspond to the same substochastic map, and hence are identified. For example, we have the identification

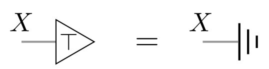

(6.58)

since both are represented by the all-ones column vector. As another simple example, the representation of any function in the propositional theory will coincide with some deterministic stochastic map in Bayes; hence, such processes in SubStoch can be viewed either acting on the left as propositional maps, or acting on the right as stochastic maps. As a final example, the representation of a delta function probability distribution $[x]$ from Bayes coincides with the representation of the value assignment asserting $X = x$ from Boole.

Consider now the representation of the scalars in Bayes and Boole. The unique scalar 1 in Bayes remains the same in this representation, while the pair of scalars in Boole, namely True and False, are represented respectively by the scalars 1 and 0 within SubStoch.

These two representations interact in the obvious way. For example, we expect the diagram

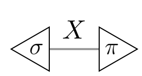

(6.59)

to give the probability $\mathsf{Prob}(\pi : \sigma)$ that the proposition π about X is true, given a state of knowledge σ about X. Indeed, this can be computed within the theory of substochastic maps:

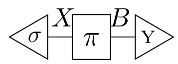

$$= \sum_{x \in X} \sigma(x) \tag{6.60}$$

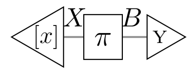

$$= \sum_{x \in X} \sigma(x)\delta_{\pi(x),\mathrm{Y}} \tag{6.61}$$

$$= \sum_{x \in \pi} \sigma(x) \tag{6.62}$$

$$= \mathsf{Prob}(\pi : \sigma). \tag{6.63}$$

It turns out that arbitrary substochastic maps can be realized by the interaction of stochastic maps and partial functions. An arbitrary substochastic map can be represented as

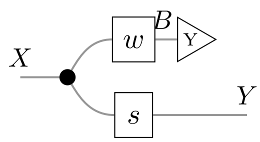

, (6.64)

where s and w are arbitrary stochastic maps. Here, w specifies the normalization of each column of the substochastic map, while s specifies the action of each column (apart from the normalization factor). Hence, we have that

Remark 14. *Categorically,* SubStoch *is the symmetric monoidal category* FinSubStoch *where objects are finite sets and morphisms are substochastic maps between them, together with a selected subobject classifier* $\{\mathrm{Y}, \mathrm{N}\}$*. It is, however, useful to see how this structure arises from the interaction between the probabilistic part (as described by* FinStoch*) and the propositional part (as described by* $\textsc{FinSet}_{\textsc{part}}$*).*

Remark 15. *Strictly speaking, in what follows, we will have certain inferential systems labeled by sets of infinite cardinality. Recalling that* Proc *is a free process theory, its hom-set will typically be of infinite cardinality. (In contrast, note that the hom-sets in* Func *are finite, so the issue will not arise there). To formally deal with this, rather than working with* FinSubStoch*, we should work with* SubStoch*—defined as the Kleisli category of the subdistribution monad on* Set*.*[10] *In this process theory, states are subnormalised probability distributions with finite support, and general processes are substochastic maps that do not generate distributions with infinite support. In future work, it will be important to consider*

[10] Thanks are due to Martti Karvonen for recognizing this issue, giving the resolution to it, and then explaining the resolution to us.

more sophisticated measure-theoretic approaches to infinite sets—in particular, to allow for ontic state spaces of infinite cardinality.

6.4 Causal-inferential theories

Having introduced theories of causal primitives and of inferential primitives, we can now describe how the two interact to define causal-inferential theories. We first (non-exhaustively) describe features that we expect of a generic causal-inferential theory. We then develop two special cases that instantiate these features, namely causal-inferential theories of lab procedures, or *operational CI theories*, and the causal-inferential theory of functional dynamics, namely the *classical realist CI theory.*

A causal-inferential theory consists of a triple of process theories and a triple of DP (partial) maps between them:

$$\text{CAUS} \xrightarrow{\mathbf{e}} \text{C-I} \underset{\mathbf{p}}{\overset{\mathbf{i}}{\rightleftarrows}} \text{INF}\ , \tag{6.65}$$

where the process theory C-I includes all of the causal and inferential systems coming from CAUS and INF, respectively. (That is, **i** and **e** are injective on objects.)

Remark 16. *Note that we will continue to draw the causal systems (i.e., those in the image of* **e***) vertically and the inferential systems (i.e., those in the image of* **i***) horizontally (i.e., just as in the respective domains of* **e** *and* **i***). This choice is merely a convention; we could alternatively have just used a different style of wire or labelling system to keep track of this information. That is, on a formal level,* C-I*, like* CAUS *and* INF*, is simply a symmetric monoidal category.*

We define how the causal and inferential systems interact by introducing three special processes that involve both causal and inferential systems, and that are subject to a collection of rewrite rules. These three processes allow one to (i) specify a state of knowledge about a particular causal dynamics, (ii) gain information about a classical causal system, and (iii) ignore causal systems. These three generators are denoted, respectively, by

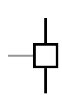 , 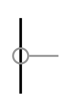 and . (6.66)

In our examples, the theory corresponding to C-I in Eq. (6.65) is defined as the process theory constructed out of arbitrary composition of these three 'generators', together with the processes in the image of **i**.[11]

The rewrite rules that these satisfy will be (to some extent) dependent on the exact causal-inferential theory that one is interested in – indeed, there will be an important distinction between the rewrite rules for our two key examples. An important direction for future research is therefore to determine which aspects of these rewrite rules are in fact generic to *all* CI theories. We return to this question in Section 6.9.1.

Finally, we introduce a partial map **p** that allows one to make inferential predictions, by mapping diagrams in C-I with only inferential inputs and outputs to a particular inferential process within SubStoch. That is, in our examples, the map **p** makes a probabilistic prediction given one's knowledge about a particular causal scenario.

We will for simplicity sometimes refer to C-I as 'the causal-inferential theory' or use the symbol C-I to refer to the full causal-inferential theory, since C-I is the primary process theory of relevance. Strictly speaking, however, a causal-inferential theory is given by a triple as in Eq. (6.65) (including, in particular, the prediction map).

6.4.1 Operational causal-inferential theories

We will use the term 'operational causal-inferential (CI) theory' to refer to a causal-inferential theory of lab procedures; that is, taking Caus = Proc and Inf = SubStoch in Eq. (6.65), i.e.,

$$\textsc{Proc} \xrightarrow{\mathbf{e}} \text{P-S} \underset{\mathbf{p}}{\overset{\mathbf{i}}{\leftrightarrows}} \textsc{SubStoch} \, . \tag{6.67}$$

We have already defined Proc and SubStoch, but it remains to explicitly define P-S and the diagram-preserving maps **e**, **i** and **p** between these three process theories.

SubStoch is a subprocess theory of P-S, explicitly represented by the inclusion of SubStoch into P-S via a DP map $\mathbf{i} : \textsc{SubStoch} \to \text{P-S}$. Diagrammatically, we denote

[11] Note that, unlike the situation with the **i** map, the construction of C-I does not require any of the processes in the image of **e** to be generators, since (as we will see) the set of all such processes can be obtained by combining elements of SubStoch (the image of **i**) with the first generator via Eq. (6.71). Indeed, this is how the map **e** will be *defined*; see Eq. (6.81).

this as a green map, e.g.

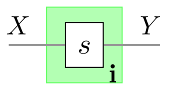

(6.68)

That is, $\mathbf{i}$ denotes that some process in P-S is a member of the subprocess theory SubStoch. For example, in the equation

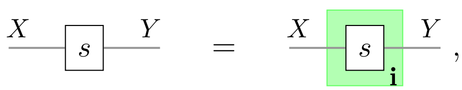

(6.69)

the process s on the RHS is a process in SubStoch, shown being mapped by $\mathbf{i}$ to the process s in P-S, shown on the LHS. In this case, s as a process in P-S (on the LHS) is in the image under $\mathbf{i}$ of a process s in SubStoch (on the RHS).

However, Proc is not a sub-process-theory of P-S; that is, the DP map $\mathbf{e} : \text{Proc} \to \text{P-S}$ is a more complicated embedding. All of the systems from Proc are directly included as systems within P-S. In order to fully define the embedding map $\mathbf{e}$, we must define how the causal and inferential systems in P-S interact.

To proceed, we discuss the interpretation within P-S of the three fundamental generators of interactions between the causal and inferential systems.

The first generator allows us to specify our state of knowledge about which procedure occurs. There is one such generator for each pair (A, B) of systems, depicted as

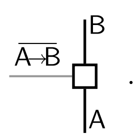

(6.70)

We then interpret

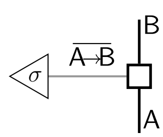

(6.71)

as describing that we have state of knowledge σ about the transformation procedure taking A to B. Indeed, σ is here a probability distribution over the set $\overline{\mathsf{A}\to\mathsf{B}}$ of transformation procedures. We will denote the delta function state of knowledge on $t \in$ Proc by $[t]$, so that

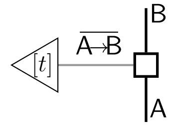

(6.72)

represents certainty that A transforms into B via the procedure t.

Now, suppose that we have states of knowledge about each of two transformation procedures where the output of one is the only input of the other (so that they are purely cause-effect related), then there is a stochastic map that represents how to update knowledge about each of the two individual transformations to a state of knowledge of the composite transformation procedure. We denote this stochastic map as

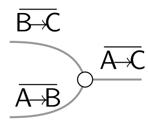

, (6.73)

defined by linearity and its action on delta-function states of knowledge, namely

$\forall t, t'$

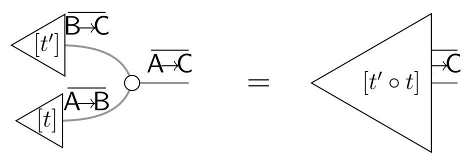

(6.74)

where $\circ$ denotes sequential composition in Proc. To reproduce the intuitive notion of composition, we demand that

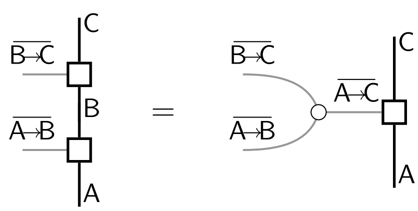

. (6.75)

This rewrite rule can be understood as the equality of two different methods of specifying one's knowledge that the causal structure is a chain $A \to B \to C$. The fact that B is a complete causal mediary between A and C can be encoded in the causal structure of a diagram (as on the LHS), but the RHS encodes it in the inferential structure, as a state of knowledge about the transformation from A to C that is specified in terms of one's state of knowledge about a transformation from A to B and about a transformation from B to C.

Similarly, we can define a stochastic map which represents how one combines states of knowledge about transformations that are causally disconnected. Specifically, suppose that in a transformation from AC to BD, B is influenced only by A and D is influenced only by C. The relevant stochastic map,

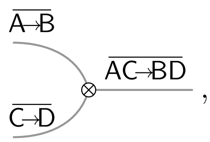

(6.76)

can be defined by linearity and its action on delta-function states of knowledge, namely

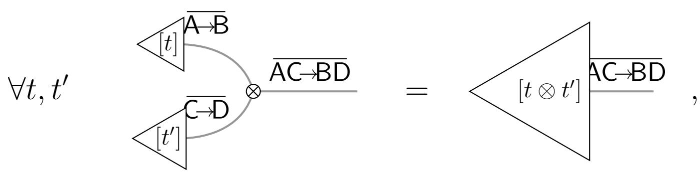

(6.77)

where $\otimes$ denotes the parallel composition of processes within Proc. Then, in analogy to Eq. (6.75), we demand that

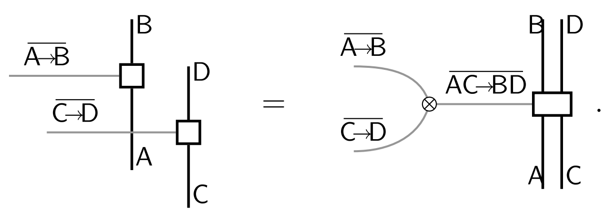

(6.78)

Finally, it will often be useful to be able to interpret some bits of wiring as themselves being processes in P-S, namely, the identity procedure $\mathbb{1}$ and swap procedure $\mathbb{S}$ respectively. We therefore impose that

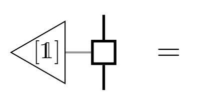

(6.79)

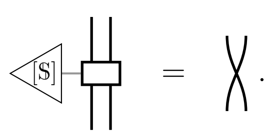

(6.80)

Essentially, these constraints allow us to lift the compositional, that is, causal, structure of Proc *into our theory of states of knowledge about* Proc, that is, into P-S. In particular, given the basic generators and constraints we have just introduced, one can construct a DP map $\mathbf{e} : \text{Proc} \to \text{P-S}$ that embeds procedures into P-S as delta-function states of knowledge:

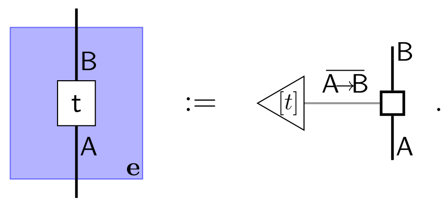

(6.81)

It is simple to check that our constraints on this generator imply that this map is indeed diagram-preserving. For example,

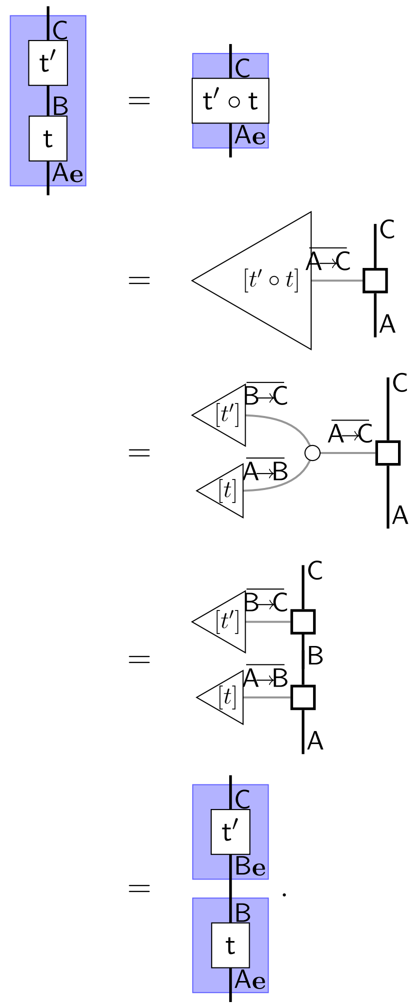

(6.82)

(6.83)

(6.84)

(6.85)

(6.86)

The second generator allows us to directly gain knowledge from a classical causal system. There is one such generator for each classical system X:

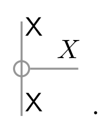

(6.87)

This can equivalently be interpreted as a generator that allows us to ask a question about a classical system by attaching a proposition to it. For example, a proposition π about the outcome of a measurement m is depicted as

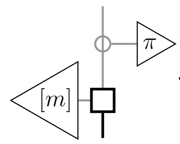

(6.88)

Note that there is no such generator for systems that are not classical, since for these, there is no way to directly gain information about the system; rather, one can only probe them indirectly via their interaction with classical systems.

As with the previous generator, this generator must satisfy certain constraints. First, it must satisfy

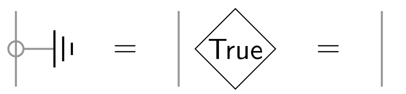

(6.89)

That is, asking about the tautological proposition on a system is the same as not asking anything at all about the system.

Additionally, under sequential composition we demand

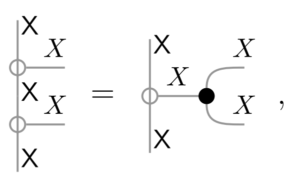

(6.90)

where • is the stochastic broadcasting map that can be defined by linearity and its action on delta-function states of knowledge, namely,

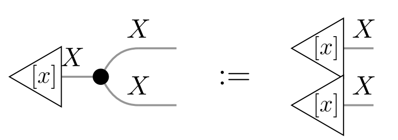

(6.91)

Eq. (6.90) states that directly gaining knowledge about the same system twice is the same as copying the knowledge gained from the system.

We also have a constraint for parallel composition:

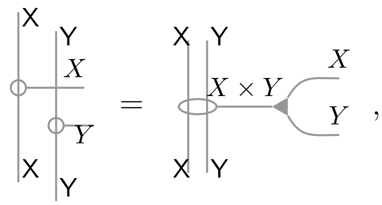

(6.92)

where ◀ is really just the identity stochastic map, but where one is changing from diagrammatically denoting a pair of systems as a single wire to denoting them by a pair of wires. Similarly, we have ▶, which merges a pair of wires into a single wire:

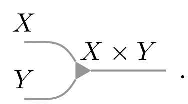

(6.93)

Finally, we introduce our last generator, which represents the ignoring of a causal system as

(6.94)

We depict the process of ignoring a system by the same symbol (albeit smaller) as marginalisation in the inferential theory to make clear that this is not a physical discarding process (such as physically annihilating a system somehow). It merely represents the fact that one is no longer interested in this system. That is, this *ignoring process* is applied whenever an agent decides that they will consider no further propositions about a system or its causal descendents.

The ignoring process satisfies the constraint

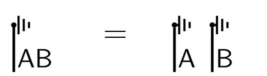

(6.95)

stating that ignoring a composite system is the same as ignoring each of its components. Moreover, it has a nontrivial interaction with the generator of Eq. (6.70), as we demand that

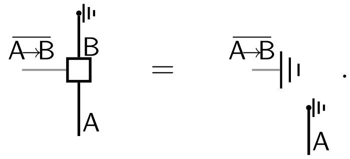

(6.96)

That is, if one is not going to ask any propositional questions about B, then one can ignore the identity of the transformation from A to B, as well as A itself. We term this the constraint of *ignorability*.

We will assume that the ignoring process for the trivial system, I, is simply given by the empty diagram and hence we obtain two special cases of Eq. (6.96):

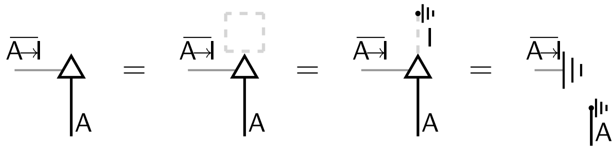

(6.97)

and

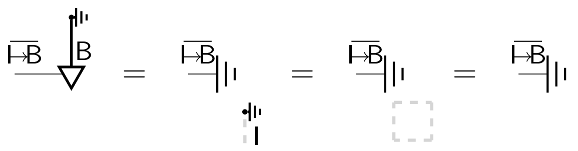

(6.98)

Due to its compositional nature, this framework is clearly able to express scenarios far more general than the well-studied prepare-measure scenario. Even within a simple prepare-measure scenario, our framework allows us to express generality that is typically neglected. In a prepare-measure scenario, the conventional states of knowledge one has and propositions one considers (i.e., the conventional inferential structure) are represented in our framework by

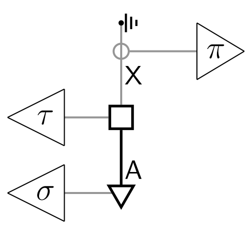

, (6.99)

where σ is a state of knowledge about the preparation of A, τ is a state of knowledge about the measurement on A, and π is a proposition about the measurement outcome X. One can change the inferential structure of the scenario without changing the causal structure of the scenario, e.g., to

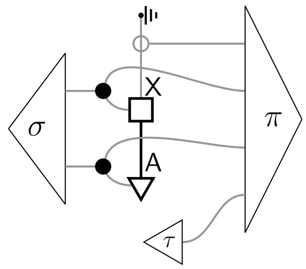

, (6.100)

where σ is a joint state of knowledge about the preparation and measurement procedures, τ is a state of knowledge about some auxiliary inferential system, and π is a joint proposition about which preparation and measurement procedures were performed, the outcome of the measurement, and the auxiliary inferential system. This extra generality is useful as it allows one to model scenarios where the choice of preparation and the choice of measurement are correlated. This occurs, for instance, in two-party cryptography, wherein one party prepares a system and the other measures it, and they correlate their actions based on private randomness that they share. Another example arises for a pair of communicating parties when the preparations and measurements are done relative to local reference frames in the labs of the parties, and where these are correlated with one another, but uncorrelated with a background reference frame.

It will be helpful to introduce a notation that represents a generic diagram in P-S while not displaying all the internal structure—that is, how the diagram is built up out of the generators—but rather only shows its open inputs and outputs. We draw such a generic diagram as

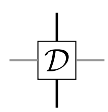

. (6.101)

In a process theory, diagrams without any inputs and outputs are termed *closed diagrams*. Analogously, diagrams whose only inputs and outputs are inferential will be termed *causally closed* diagrams and diagrams whose only inputs and outputs are causal will be termed *inferentially closed* diagrams.

At this point, we can make a useful observation: any diagram in P-S that can be written using generators that do not involve causal systems can be considered as the image of some process in SubStoch under **i**, as in Eq. (6.69).

We have so far introduced the inferential and causal components of P-S and the maps from these into P-S, namely,

$$\textsc{Proc} \xrightarrow{\mathbf{e}} \text{P-S} \xleftarrow{\mathbf{i}} \textsc{SubStoch}. \tag{6.102}$$

We now introduce the prediction map **p**, which describes the inferences one can make in a given scenario.[12] **p** is a partial DP map whose domain is the set of causally closed processes in P-S, and whose co-domain is SubStoch.

An example of the sort of causally closed processes on which **p** is defined is:

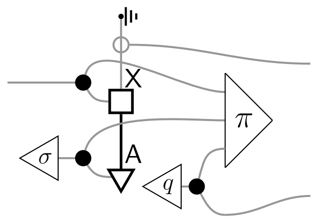

, (6.103)

A diagram that has open causal inputs or outputs is *not* in the domain of the prediction map, because the open causal wires correspond to systems about which either no state of knowledge or no propositional question has been specified. For example, the inferences one should make in a situation described by the diagram

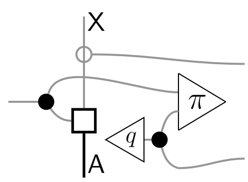

(6.104)

depend on what one knows about previous procedures on A as well as any propositions one considers about X.

Consider first the simple case of closed diagrams in P-S. These are mapped to closed

[12]The idea of separating out the descriptive and the probabilistic components of one's notion of an operational theory can be found in earlier works, notably Ref. [69].

diagrams (scalars) in SubStoch—i.e., elements of $[0,1]$. In the following example,

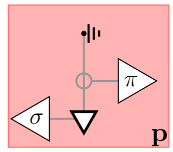

$$= \mathsf{Prob}(\pi : \sigma), \tag{6.105}$$

the prediction map specifies the probability that one should assign to the proposition π being true given that one's state of knowledge is σ. Meanwhile,

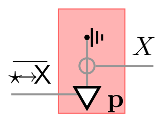

(6.106)

is a stochastic map in SubStoch, which takes a state of knowledge about the preparation X as input and returns a state of knowledge about X.

There is an obvious consistency constraint on processes in P-S that are also processes in the sub-process-theory SubStoch, that is, those that are in the image of $\mathbf{i}$. If one maps a process in SubStoch to P-S by the inclusion map $\mathbf{i}$ and then back to SubStoch via the prediction map $\mathbf{p}$, one should clearly obtain the process itself back again, so that

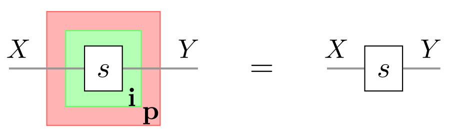

(6.107)

Hence it is a partial left inverse of $\mathbf{i}$, that is, $\mathbf{p} \circ \mathbf{i} = \mathbb{1}_{\text{SubStoch}}$.

Although $\mathbf{p}$ is only a partial map, it is still diagram-preserving on its domain; e.g., one can write

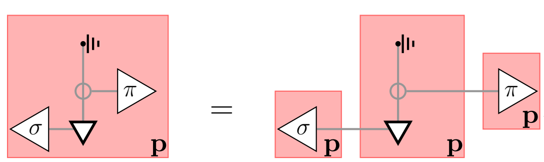

(6.108)

In summary, an operational CI theory is specified by a triple of process theories and a triple of DP maps between them, succinctly drawn as

$$\text{Proc} \xrightarrow{\mathbf{e}} \text{P-S} \underset{\mathbf{p}}{\overset{\mathbf{i}}{\leftrightarrows}} \text{SubStoch}, \tag{6.109}$$

where we use a dashed line to denote the fact that **p** is partial.

Properties of the prediction map

Our constraint of ignorability, Eq. (6.96), implies that the probabilities assigned to propositions about systems are independent of what is known about the future processes applied to the system. For example, the probability

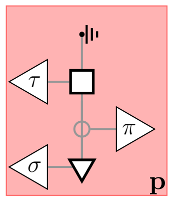

=

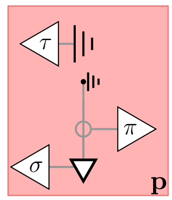

(6.110)

=

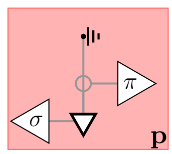

(6.111)

is seen to be independent of the state of knowledge τ. In Ref. [69], this constraint is termed 'causality', and taken to be of central importance. In our framework, however, it does not express any notion of causality. As we discuss further in Appendix C.2, the causal structure in our framework is primitive, and *cannot* be defined in terms of any probabilistic facts such as those expressed by Eq. (6.111). In our framework, the condition of ignorability, Eq. (6.96), does not play a particularly special role; it is simply a fact about the way one makes inferences. In addition to implying Eq. (6.111), it implies many similar independence relations. For example, it implies that a state of knowledge τ about a causal process occurring on one subsystem of a composite whose output is ignored is irrelevant for making inferences about the other subsystem:

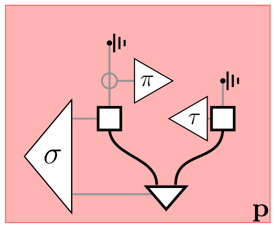 $=$ 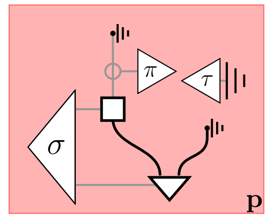 (6.112)

$=$ 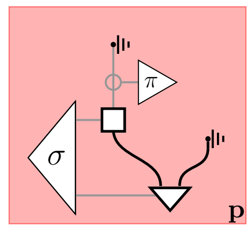 . (6.113)

Conveniently, $\mathbf{p}$ can be fully specified by a relatively simple set of data: the probabilities assigned to point-distributed states of knowledge and atomic propositions. This is exactly the form of data provided in traditional approaches to operational theories.

Theorem 6.4.1. *For every process $\mathcal{D} \in$ P-S in the domain of $\mathbf{p} :$ P-S $\to$ SubStoch, i.e., which is causally closed, the image $\mathbf{p}(\mathcal{D})$ of $\mathcal{D}$ under $\mathbf{p}$ is fully specified by the probabilities assigned to atomic propositions on its inferential output and point distributions on its inferential input.*

Proof. Consider an arbitrary causally closed process $\mathcal{D}$ and imagine mapping it into SubStoch via

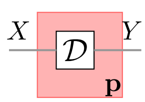 . (6.114)

This is simply a substochastic map and hence is fully characterized by the set of scalars

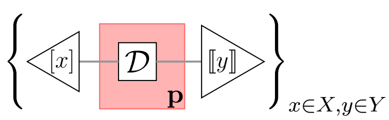 , (6.115)

where $[\![y]\!]$ is an atomic proposition on Y. Such scalars can be rewritten as

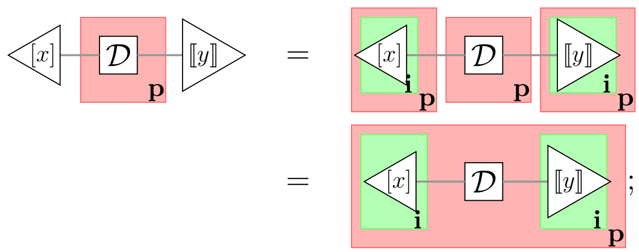

(6.116)

(6.117)

that is, they are the probabilities assigned to atomic propositions $[\![y]\!]$ on Y given point distributions $[x]$ on X, which is what we set out to prove. □

Quantum theory as an operational CI theory

The most straightforward way to cast quantum theory as an operational CI theory is as follows. The causal subtheory for quantum theory, which we denote $\textsc{Proc}_Q$, contains laboratory procedures whose inputs and outputs are classical and quantum systems. The inferential subtheory is the classical one, $\textsc{SubStoch}$. The full theory, $\mathrm{P_QS}$, is constructed as

$$\textsc{Proc}_Q \xrightarrow{\mathbf{e}_Q} \mathrm{P_QS} \underset{\mathbf{p}_Q}{\overset{\mathbf{i}_Q}{\leftrightarrows}} \textsc{SubStoch}\ . \tag{6.118}$$

Here, the specific prediction map $\mathbf{p}_Q$ singles out quantum theory, and is defined as follows. To every quantum system is associated an algebra of operators on a complex Hilbert space of some dimension, and to every classical system is associated an algebra of *commuting* operators on such a Hilbert space; to every diagram is associated a completely-positive [257] trace-nonincreasing map between these; then, the joint probability distributions (on any set of propositions attached to the classical systems) can be computed by composition of these completely-positive trace-preserving maps.

6.4.2 Classical realist causal-inferential theories

Next, we turn our attention to the second class of causal-inferential theories that we will consider, namely classical realist CI theories. These are very similar to the operational CI theories just introduced, but the causal theory is not taken to be a process theory $\textsc{Proc}$ of laboratory procedures, but rather a process theory representing fundamental dynamics of ontic states of systems. In our case, we will take this to be the process theory $\textsc{Func}$

of functional dynamics, introduced in Section 6.2.2. The extra structure in Func relative to Proc accounts for all the differences within our framework between an operational CI theory and a classical realist CI theory, and implies that there is essentially a unique classical realist CI theory, insofar as there is a unique prediction map.

We will use the term 'classical realist CI theory' to refer to the following causal-inferential theory of functional dynamics, namely

$$\textsc{Func} \xrightarrow{\mathbf{e}'} \text{F-S} \underset{\mathbf{p}^*}{\overset{\mathbf{i}'}{\rightleftarrows}} \textsc{SubStoch}. \tag{6.119}$$

We have labeled the diagram-preserving maps here by $\mathbf{e}'$, $\mathbf{i}'$ and $\mathbf{p}^*$ to distinguish them from those in an operational CI theory.

The construction proceeds much like that in the previous section. SubStoch is a subprocess theory of F-S, explicitly represented by the inclusion of SubStoch into F-S via a DP map $\mathbf{i}' : \textsc{SubStoch} \to \text{F-S}$, diagrammatically represented as

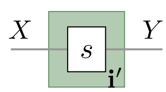

. (6.120)

Func is not a sub-process-theory of F-S, but rather embeds into F-S via a map $\mathbf{e}' : \textsc{Func} \to \text{F-S}$, which we will define after introducing some relevant generators.

The first generator again allows one to specify a state of knowledge about the functional dynamics. There is one such generator for each pair of systems (Λ, Λ'), depicted as

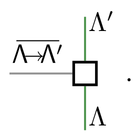

. (6.121)

Then, the diagram

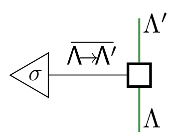

(6.122)

represents the state of knowledge σ about the function from Λ to Λ' describing the dynamics.

Naturally, we demand that constraints analogous to those in Eq. (6.75) and Eq. (6.78) are satisfied, which then implies that we can construct a DP map $\mathbf{e}' : \text{FUNC} \to \text{F-S}$ defined as

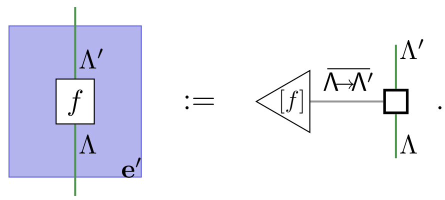

(6.123)

The second generator allows us to directly gain knowledge from an ontological system, or equivalently, to ask a question about a system by attaching a proposition to it. Here, we see the first key distinction between ontological and operational CI theories—for operational CI theories, we could only define such a generator for classical systems; however, because all systems in FUNC are sets Λ, this generator

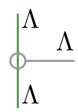

(6.124)

can be defined for any system in F-S. Naturally, we demand that each such generator satisfies the constraints stipulated in Eqs. (6.90) and (6.92).

Finally, we introduce a generator

(6.125)

which represents ignoring the system Λ and that satisfies constraints analogous to Eq. (6.95) and Eq. (6.96).

We now have the tools to describe a wide range of scenarios. For example, the scenario

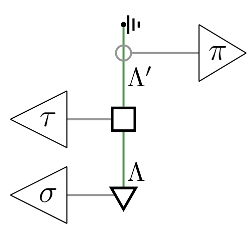

(6.126)

might arise as a classical realist model of the operational scenario in Diagram (6.99). This is analogous to a prepare-measure scenario. Even in this simple causal structure, however, we can also describe more general inferential structures; for example, an analogue of Diagram (6.100), namely

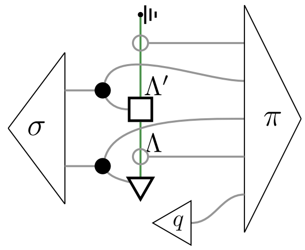

. (6.127)

In fact, this is even more general than Diagram (6.100), since in F-S (unlike in P-S), one can consider propositions about *arbitrary* systems.

Perhaps the central distinction between P-S and F-S is that in F-S, there is a constraint on the interactions between the first two generators we introduced. This constraint is a consequence of the fact that one can attach propositions to any system in F-S, together with our assumption that the causal mechanisms are described by functions. This means that we can, in certain situations, propagate what we know about one physical system to knowledge about another. It is due to this single extra constraint that classical realist CI theories have so much more structure than operational CI theories.

This interaction between the two generators is governed by the equality:

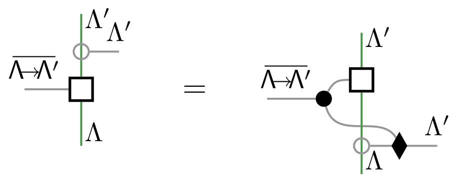

, (6.128)

where the black diamond converts a state of knowledge about Λ and a state of knowledge about $\overline{\Lambda\mapsto\Lambda'}$ into a state of knowledge about Λ', and is defined by linearity and its action on delta-function states of knowledge, namely,

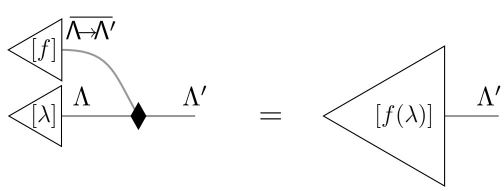

, (6.129)

or equivalently,

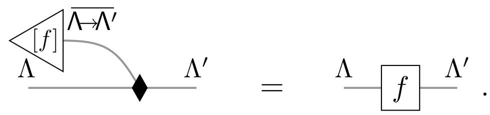

(6.130)

Eq. (6.128) ensures that one can specify what one knows about the output Λ' of some dynamical process in one of two equivalent ways: either by knowing about it directly, or by taking what is jointly known about the dynamics and the state that is input to the dynamics, and then propagating one's beliefs accordingly (i.e., according to the stochastic map ⧫). To give a simple example, suppose we have a delta function state of knowledge that the ingoing system Λ is prepared in state λ, and that the functional dynamics are given by f, then this rewrite rule allows us to reason as follows:

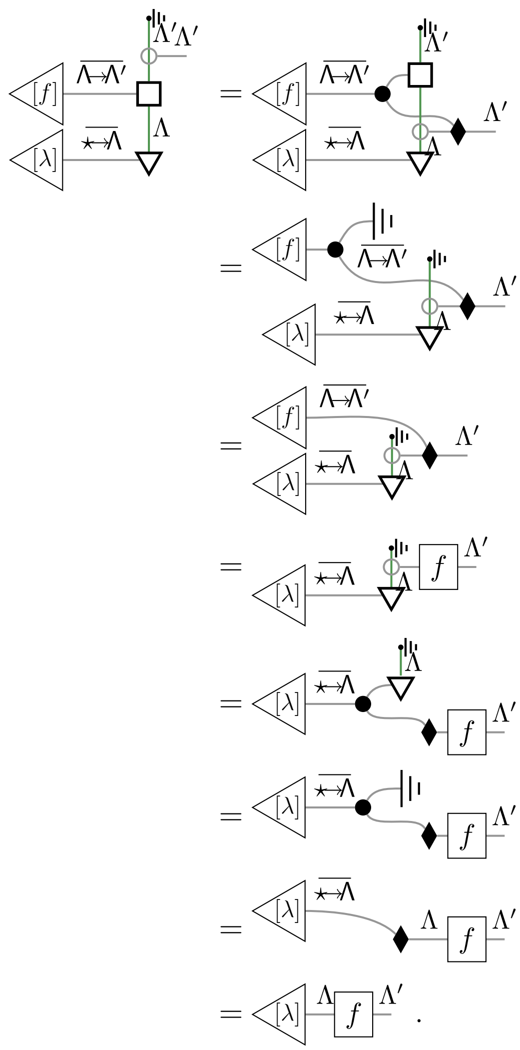

(6.131)
(6.132)
(6.133)
(6.134)
(6.135)
(6.136)
(6.137)
(6.138)

From Eq. (6.133) to Eq. (6.134), the black diamond turns a delta function state of knowledge about functional dynamics from Λ to Λ' into a (functional) propagation of one's state of knowledge about Λ to one's state of knowledge about Λ'. The rewrites from Eq. (6.134)-Eq. (6.138) are analogous to Eq. (6.131)-Eq. (6.134), but slightly more subtle insofar as they involve the special case where an input is trivial.

For this special case with a trivial input system, Eq. (6.128) becomes

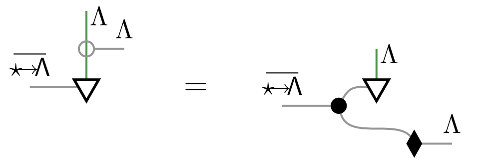

(6.139)

where, in this special case, $\blacklozenge$ is simply the isormorphism between $\overline{\star\!\to\!\Lambda}$ and Λ. For convenience, we will henceforth denote this isomorphism by

(6.140)

so that

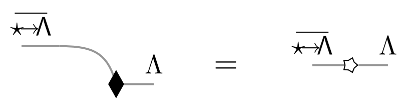

(6.141)

and

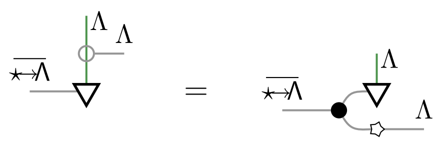

(6.142)

Predictions are made in a classical realist CI theory in a manner analogous to how predictions are made in an operational CI theory. They are represented by a partial diagram-preserving map, $\mathbf{p}^* : \text{F-S} \to \textsc{SubStoch}$, whose domain is given by the set of causally closed processes in F-S. As before, the prediction map is a partial left inverse of $\mathbf{i}'$, so that $\mathbf{p}^* \circ \mathbf{i}' = \mathbb{1}_{\textsc{SubStoch}}$. For example, closed diagrams in F-S are mapped to closed diagrams (scalars) in SubStoch—elements of $[0, 1]$; e.g.,

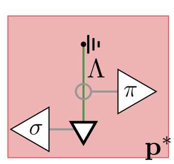

$$= \mathsf{Prob}(\pi : \sigma) \ . \qquad (6.143)$$

Meanwhile,

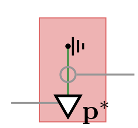

(6.144)

is a stochastic map in SubStoch, as before. Analogous to Theorem 6.4.1, one has that for every process $\mathcal{D} \in$ F-S in the domain of the prediction map $\mathbf{p}^*$, its image under $\mathbf{p}^*$ is fully specified by the probabilities assigned to atomic propositions on its output given point distributions on its input.

There is a key difference between the prediction map in a classical realist and in an operational CI theory: for classical realist CI theories, this map is unique. To show this, we first prove a normal form for general diagrams in F-S.

Theorem 6.4.2. *Any diagram in the classical realist CI theory* F-S *can be rewritten (using rewrite rules in* F-S*) into the form*

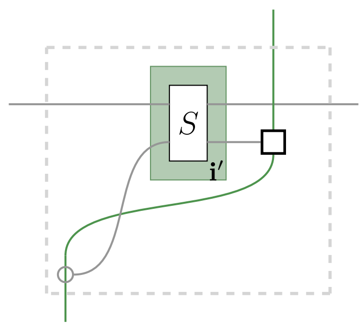

, (6.145)

where S is a substochastic map in SubStoch*.*

Proof. See Appendix C.4. □

We conjecture that this normal form is unique, or equivalently, that the substochastic map S is unique. (To prove this, it would suffice to prove that the normal form description of each generator is unique, since the composition of two diagrams in normal forms has a unique normal form description.)

Note that there is *not* an equivalent normal form for diagrams in operational CI theories, as Theorem 6.4.2 strongly relies on the constraint of Eq. (6.128). This normal form then allows us to prove that the interactions between SubStoch and Func single out a unique prediction map for the full theory F-S.

Theorem 6.4.3. *The prediction map* $\mathbf{p}^*$ *is unique.*

Proof. Consider an arbitrary process in the domain of $\mathbf{p}^*$—that is, an arbitrary causally closed process $\mathcal{D}$. Writing it in normal form, we have

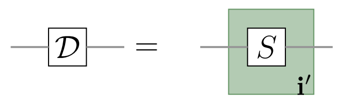

(6.146)

for some substochastic map S. Furthermore, S is unique since $\mathbf{i}'$ is an inclusion map and hence injective. Applying the prediction map, then, one has

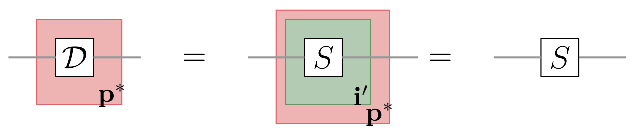

(6.147)

where the last line follows from the fact that $\mathbf{p}^* \circ \mathbf{i}' = \mathbb{1}_{\text{SubStoch}}$. Hence, the prediction map applied to any process in its domain is associated with a unique real matrix, and so $\mathbf{p}^*$ is unique. □

The full picture of a classical realist CI theory is therefore given by a triple of process theories and a triple of DP maps between them:

$$\text{Func} \xrightarrow{\mathbf{e}'} \text{F-S} \underset{\mathbf{p}^*}{\overset{\mathbf{i}'}{\rightleftarrows}} \text{SubStoch}, \tag{6.148}$$

where we use a dashed line to denote that $\mathbf{p}^*$ is partial.

We close this section by noting that it remains to determine the scope of classical realist CI theories. For instance, it is unclear whether Bohmian mechanics can *formally* be cast as such a theory. (Note that this is not specific to our framework; it is also unclear whether it can be formalized within the standard framework of ontological models.) In any case, we note that the central aim of our framework is not to capture the diversity of interpretational views, but rather to make progress on the questions posed in the introduction.

6.5 Inferential equivalence

We now define a notion of inferential equivalence between processes in a causal-inferential theory. This definition can clearly be made in any causal-inferential theory, but we will focus here only on operational CI theories and then on classical realist CI theories. This will let us define quotiented operational CI theories and quotiented classical realist CI theories.

We will discuss how the former relates to the notion of a generalized probabilistic theory, while the latter subsumes the traditional notion of an ontological model.

6.5.1 Inferential equivalence in operational CI theories

Two elements of P-S are inferentially equivalent if and only if they lead to exactly the same predictions, no matter which causally closed diagram they are embedded in. To make such statements diagrammatically, it is useful to introduce the notion of a *tester* for a given process—that is, a special case of a clamp (introduced in Section 6.1.1) whose composition with a given process yields a causally closed diagram. As a simple example, we say that two states of knowledge $\sigma_{\overline{\mathsf{A}\to\mathsf{B}}}$ and $\sigma'_{\overline{\mathsf{A}\to\mathsf{B}}}$ about a transformation procedure from A to B are inferentially equivalent with respect to the prediction map $\mathbf{p}$, denoted

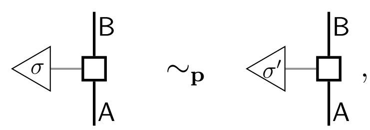

(6.149)

if and only if they make the same predictions for all testers, $\mathcal{T}$, so that

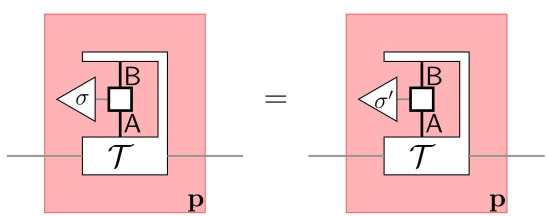

$\forall \mathcal{T} \in$ P-S. (6.150)

As an explicit example from within quantum theory, consider four lists of laboratory instructions, denoted P_1 to P_4, that are designed to prepare the quantum states $|0\rangle$, $|1\rangle$, $|+\rangle$, and $|-\rangle$, respectively. Then, the states of knowledge

$$\frac{1}{2}[P_1] + \frac{1}{2}[P_2] \quad \text{and} \quad \frac{1}{2}[P_3] + \frac{1}{2}[P_4], \tag{6.151}$$

although clearly distinct, are nonetheless inferentially equivalent, as they correspond to the same quantum state (namely the maximally mixed state).

More generally, the notion of inferential equivalence for any type of process in P-S is defined as follows:

Definition 8 (Inferential equivalence for operational CI theories)**.** *Two processes in* P-S, $\mathcal{D}$ *and* $\mathcal{E}$*, are inferentially equivalent with respect to the prediction map* $\mathbf{p}$*, denoted* $\mathcal{D} \sim_{\mathbf{p}} \mathcal{E}$*, if and only if*

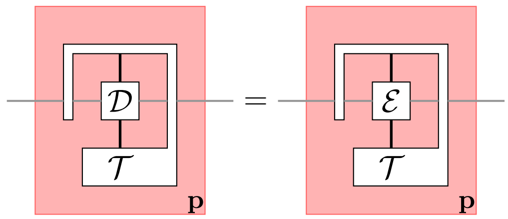

$\forall \mathcal{T} \in$ P-S. (6.152)

In fact, one can test for inferential equivalence purely in terms of probabilities (as opposed to stochastic maps).

Lemma 6.5.1. *One has inferential equivalence* $\mathcal{D} \sim_{\mathbf{p}} \mathcal{E}$ *if and only if*

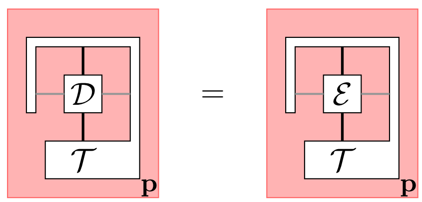

$\forall \mathcal{T} \in$ P-S. (6.153)

Proof. This follows immediately from Definition 8 and Theorem 6.4.1. □

For processes that are causally closed, this condition greatly simplifies:

Lemma 6.5.2. *Two causally closed processes are inferentially equivalent if and only if they are equal as stochastic maps under the application of the prediction map* $\mathbf{p}$*:*

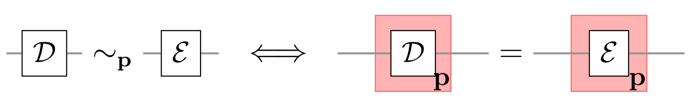

(6.154)

This is a much simpler condition to check, since one need not quantify over all possible testers.

Proof. By definition,

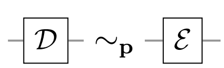

(6.155)

is equivalent to

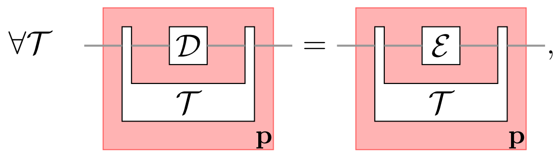

(6.156)

which, by diagram-preservation, is equivalent to

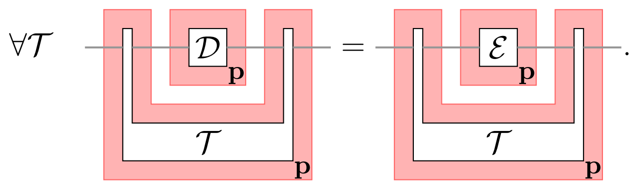

(6.157)

Finally, this is equivalent to

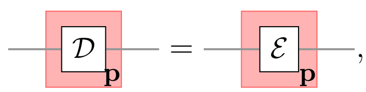

(6.158)

where the $\Rightarrow$ direction follows from the special case where $\mathcal{T}$ is simply the identity on the two inferential systems, and where the $\Leftarrow$ direction follows from the fact that equality is preserved by composition (in this case, with $\mathbf{p}(\mathcal{T})$). $\square$

For the still more restricted set of processes in the image of $\mathbf{i}$ the condition simplifies even further:

Corollary 6.5.2.1. *Two processes in the image of* $\mathbf{i}$: SubStoch $\to$ P-S *are inferentially equivalent if and only if they are equal as substochastic maps in* SubStoch*:*

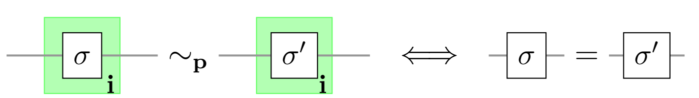

(6.159)

Proof. By Lemma 6.5.2, we have that the LHS of the implication in the corollary is equivalent to the equality

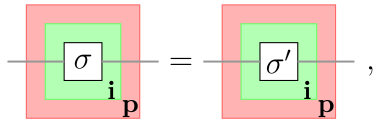

(6.160)

which gives the RHS of the implication in the corollary by Eq. (6.107), namely $\mathbf{p} \circ \mathbf{i} = \mathbb{1}_{\text{SubStoch}}$. ☐

These results imply that every causally closed process is associated with a unique stochastic map.

Lemma 6.5.3. *Every causally closed process $\mathcal{D} \in$ P-S is inferentially equivalent to a unique process in the image of* $\mathbf{i}$*, namely,*

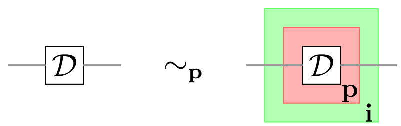

(6.161)

Proof. The constraint on the prediction map $\mathbf{p}$ of Eq. (6.107) immediately implies that

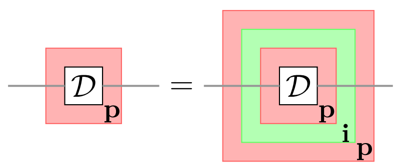

(6.162)

after which Lemma 6.5.2 implies that

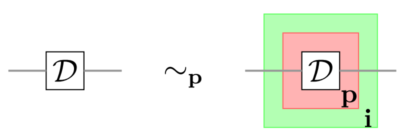

(6.163)

and then Corollary 6.5.2.1 implies that this is the *unique* process in the image of $\mathbf{i}$ in the equivalence class of $\mathcal{D}$. ☐

Quotiented operational CI theories

In many cases, one is only interested in the inferential equivalence class of processes in a causal-inferential theory. In such cases, it is useful to define a new type of theory, wherein one has quotiented[13] with respect to inferential equivalence. We now show how this is done for operational CI theories.

First, we note that the relation $\sim_{\mathbf{p}}$ is preserved under composition:

Lemma 6.5.4. *If $\mathcal{D} \sim_{\mathbf{p}} \mathcal{D}'$, then*

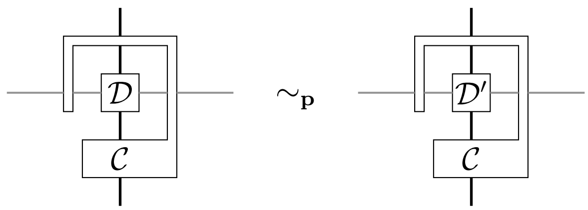

(6.164)

for all clamps $\mathcal{C}$ in P-S.

Proof. Consider, for the sake of contradiction, that there exists some $\mathcal{C}^*$ such that Eq. (6.164) fails. Then, there exists some tester $\mathcal{T}^*$ such that

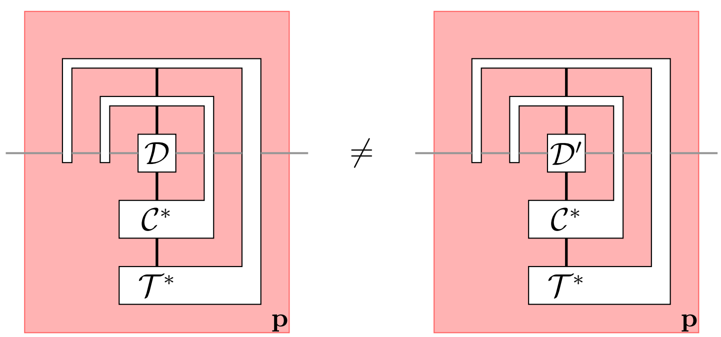

(6.165)

This, however, would imply that the tester

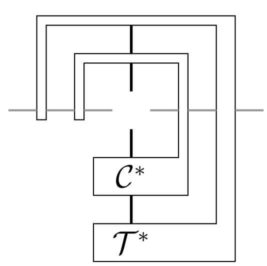

(6.166)

[13] Notions of quotiented operational theories can be found in earlier works, notably including Ref. [69].

generates different inferences for $\mathcal{D}$ and $\mathcal{D}'$, in contradiction with our initial assumption that $\mathcal{D} \sim_{\mathbf{p}} \mathcal{D}'$. □

It follows that $\sim_{\mathbf{p}}$ is a process-theory congruence relation for P-S. That is, if $\mathcal{D} \sim_{\mathbf{p}} \mathcal{D}'$ and $\mathcal{E} \sim_{\mathbf{p}} \mathcal{E}'$ then any valid composite of $\mathcal{D}$ and $\mathcal{E}$ will be inferentially equivalent to the same composite of $\mathcal{D}'$ and $\mathcal{E}'$.

Lemma 6.5.5. *The inferential equivalence relation* $\sim_{\mathbf{p}}$ *defines a process theory congruence relation on* P-S*.*

Proof. Take $\mathcal{D} \sim_{\mathbf{p}} \mathcal{D}'$ and $\mathcal{E} \sim_{\mathbf{p}} \mathcal{E}'$, and consider some arbitrary composition of the non-primed versions. As a particular illustrative example, take

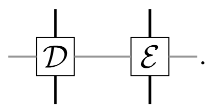
(6.167)

Using Lemma 6.5.4, the fact that $\mathcal{D} \sim_{\mathbf{p}} \mathcal{D}'$, and the fact that the (inferentially) serial composition of $\mathcal{D}$ with $\mathcal{E}$ is a special case of the clamp $\mathcal{C}$ from Eq. (6.164), implies that

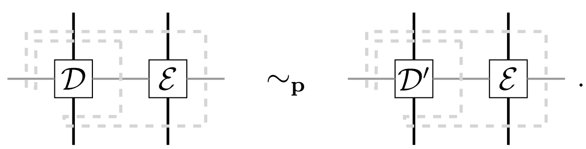
(6.168)

Then, by the same lemma, but now viewing $\mathcal{D}'$ as the clamp and using the fact that $\mathcal{E} \sim_{\mathbf{p}} \mathcal{E}'$, we have:

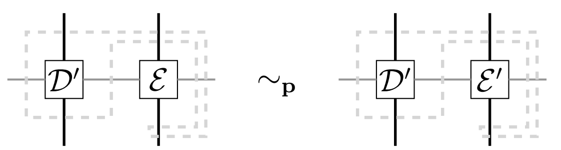
(6.169)

Putting these two together (by transitivity of $\sim_{\mathbf{p}}$) we immediately have:

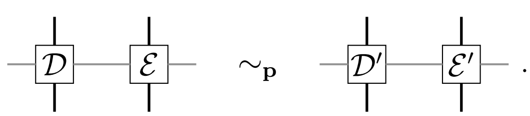
(6.170)

as we require. It is easy to see that identically structured proofs hold for any other way of composing $\mathcal{D}$ and $\mathcal{E}$. □

This lemma is important because it is necessary that inferential equivalence defines a congruence relation in order for quotienting with respect to it to yield a valid process theory.

Definition 9. *We define a quotiented operational CI theory* $\widetilde{\text{P-S}}$ *as the process theory* P-S *quotiented by the congruence relation* $\sim_{\mathbf{p}}$. *That is, it has the same systems as* P-S, *but its processes correspond to equivalence classes of processes in* P-S, *that is, to maximal sets of inferentially equivalent processes. We can moreover define a diagram-preserving map* $\sim_{\mathbf{p}} : \text{P-S} \to \widetilde{\text{P-S}}$, *as*

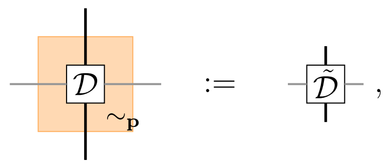
(6.171)

where $\widetilde{\mathcal{D}}$ *is the equivalence class that contains* $\mathcal{D}$. *Composition of equivalence classes is defined by the equivalence class of the composite of an arbitrary choice of representative element for each.*

That this notion of composition is well defined (i.e., independent of the choice of representative elements) follows from Lemma 6.5.5. It then follows that the quotienting map $\sim_{\mathbf{p}} : \text{P-S} \to \widetilde{\text{P-S}}$ is indeed diagram-preserving.

From the above definition, one clearly has that

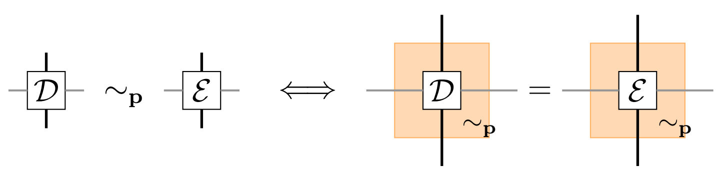
(6.172)

It is worth noting that in our framework, a quotiented operational CI theory is *not* an example of an operational CI theory (unless they are both trivial). This is because if the operational theory is nontrivial, then the quotienting operation *necessarily* loses information: equivalence classes of states of knowledge about procedures are not themselves states of knowledge about procedures. The following example proves the claim. Recall that the closed diagrams in the quotiented theory are (isomorphic to) probabilities, while in the operational CI theory, they constitute a complete description of what one knows and what one is asking in the scenario under consideration. In other words, in any quotiented operational CI theory, the only way for two closed diagrams to be distinct is if they are inferentially inequivalent whereas in any nontrivial unquotiented operational CI theory, there will exist pairs of closed diagrams that are inferentially equivalent and yet still distinct.

Using the quotienting map, the probability associated with a closed diagram can always be decomposed into a sequence of stochastic maps representing one's inferences, by grouping together processes into diagrams that are causally closed, e.g.

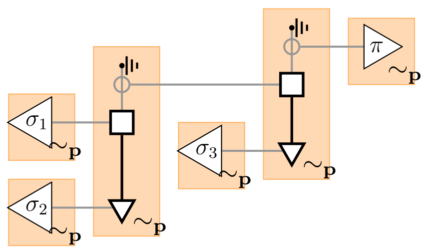

. (6.173)

Clearly, composition with the quotienting map $\sim_{\mathbf{p}}$ can be used to define two new DP maps. The map $\tilde{\mathbf{e}} : \text{Proc} \to \widetilde{\text{P-S}}$ is defined as $\tilde{\mathbf{e}} = \sim_{\mathbf{p}} \circ \mathbf{e}$, and the map $\tilde{\mathbf{i}} : \text{SubStoch} \to \widetilde{\text{P-S}}$ is defined as $\tilde{\mathbf{i}} = \sim_{\mathbf{p}} \circ \mathbf{i}$. We can also introduce a partial diagram-preserving prediction map $\tilde{\mathbf{p}}$ for the quotiented operational CI theory, whose action is given by mapping each process in $\widetilde{\text{P-S}}$ to an element (any element) of P-S in its equivalence class, and then mapping that element to SubStoch via $\mathbf{p}$. All of this can be concisely represented in the following commuting diagram:

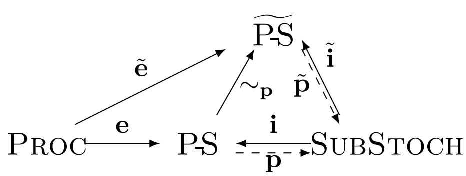

. (6.174)

Subsuming the framework of generalized probabilistic theories

At this point, one can see the relationship between our framework and another well-known framework for operational theories, namely, that of generalized probabilistic theories (GPTs).

A GPT is a minimal framework in which processes are wired together to form circuits that describe an operational scenario and predict the probabilities of the outcomes that one might observe. Perhaps the key feature of a GPT is *tomographic completeness*, which implies that processes within a GPT are taken to represent *equivalence classes* of procedures or events with respect to the operational predictions. That is, two transformations in a

GPT are represented distinctly if and only if there exists a circuit in which they can be embedded to give different probabilities for the outcome of some measurement in that circuit. The set of processes is also assumed to be convex and representable in a (typically finite dimensional) real vector space, to have a unique deterministic effect, and for composition of processes to be bilinear.

In forthcoming work [259], we prove that these properties are satisfied for a natural subset of processes in any quotiented operational CI theory (i.e., the inferentially closed processes), and that consequently the latter can be identified with GPT processes in the traditional sense. Tomographic completeness follows naturally from the quotienting which defines $\widetilde{\text{P-S}}$, and convexity of these processes is inherited from convexity of the inferential theory.

However, the quotiented operational CI theories in our framework are not equivalent to GPTs. There remain important formal and conceptual differences between the two. For example, a quotiented operational CI theory contains both causal and inferential systems and processes, while GPTs contain only a single type of system. It is also worth noting that GPT processes are conventionally viewed as representing equivalence classes of laboratory procedures, while processes in a quotiented operational CI theory have a different interpretation—they represent equivalence classes of *states of knowledge* about laboratory procedures.

6.5.2 Inferential equivalence in classical realist CI theories

Analogously, two elements of F-S are inferentially equivalent if and only if they lead to exactly the same predictions, no matter what causal diagram they are embedded in.

Definition 10 (Inferential equivalence for classical realist CI theories)**.** *Two general elements of* F-S*,* $\mathcal{D}$ *and* $\mathcal{E}$*, are inferentially equivalent with respect to the prediction map* $\mathbf{p}^*$*, denoted* $\mathcal{D} \sim_{\mathbf{p}^*} \mathcal{E}$*, if and only if*

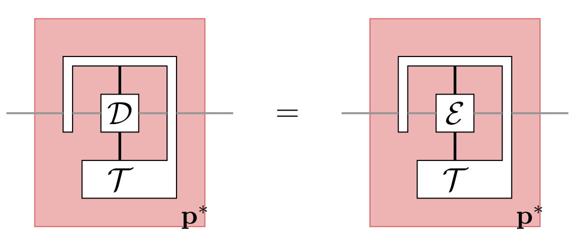

$\forall \mathcal{T} \in$ F-S. (6.175)

The fact that this is a nontrivial relationship may be somewhat surprising. Once it is recognized, however, it is not difficult to come up with examples to illustrate it; we give a simple example below.

Every process in F-S can be associated with a stochastic map, via

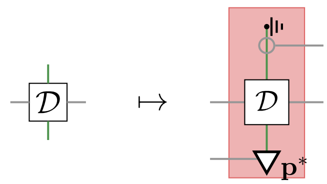

(6.176)

Using Lemma C.3 (stated and proven in Appendix C.5), we can prove the following result, which is an analogue of Lemma 6.5.2, but strengthened to include processes with open causal systems.

Lemma 6.5.6. *Two processes in* F-S *are inferentially equivalent if and only if they are associated with the same substochastic map. That is,*

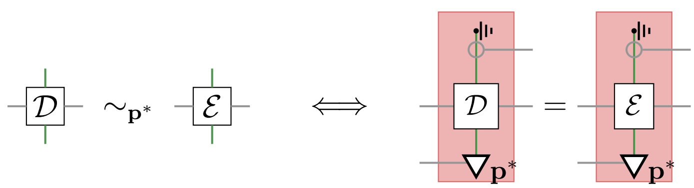

(6.177)

The proof is given in Appendix C.5.

As an explicit example, consider the four bit-to-bit functions $\{f_0, f_1, f_{\text{id}}, f_{\text{flip}}\}$. They are defined by their action on a bit $a \in \{0,1\}$, namely, $f_0(a) = 0$, $f_1(a) = 1$, $f_{\text{id}}(a) = a$, and $f_{\text{flip}}(a) = a \oplus 1$, where $\oplus$ denotes summation modulo 2. Then, the states of knowledge

$$\sigma_c = \frac{1}{2}[f_0] + \frac{1}{2}[f_1] \ \text{ and } \sigma_d = \frac{1}{2}[f_{\text{id}}] + \frac{1}{2}[f_{\text{flip}}] \tag{6.178}$$

are distinct but inferentially equivalent. This is easily seen by the fact that both states of knowledge correspond to the same stochastic map, namely, the completely randomizing bit-to-bit channel:

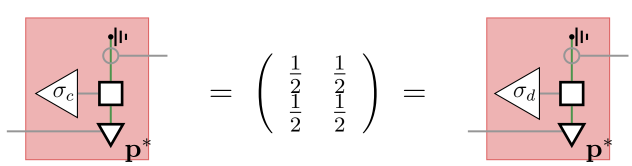

(6.179)

Quotiented classical realist CI theories

In direct analogy with Lemma 6.5.5 and its proof, one can show that $\sim_{\mathbf{p}^*}$ defines a congruence relation, and it follows that one can quotient the classical realist CI theory with respect to this relation.

Definition 11. *We define a quotiented classical realist CI theory* $\widetilde{\text{F-S}}$ *as the process theory* F-S *quotiented by the congruence* $\sim_{\mathbf{p}^*}$*. That is, it has the same systems as* F-S*, but its processes correspond to equivalence classes of processes in* F-S*, that is, to maximal sets of inferentially equivalent processes. We can moreover define a diagram-preserving map* $\sim_{\mathbf{p}^*}\colon \text{F-S} \to \widetilde{\text{F-S}}$*, as*

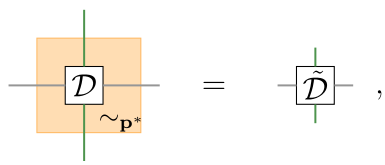

(6.180)

where $\widetilde{\mathcal{D}}$ *is the equivalence class that contains* $\mathcal{D}$*. Composition of equivalence classes is defined by the equivalence class of the composite of an arbitrary choice of representative element for each.*

That this notion of composition is well defined (i.e. independent of the choice of representative elements) follows from the natural analogue of Lemma 6.5.5. It then follows that the map $\sim_{\mathbf{p}^*}\colon \text{F-S} \to \widetilde{\text{F-S}}$ is indeed diagram-preserving.

From the above definition, it clearly follows that

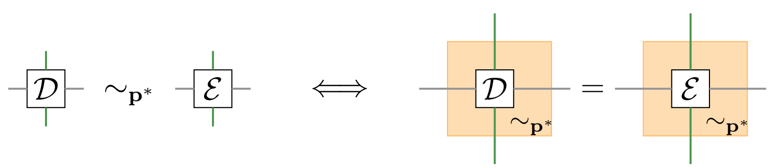

(6.181)

Clearly $\sim_{\mathbf{p}^*}$ can be used to define two new DP maps $\tilde{\mathbf{e}}' : \textsc{Func} \to \widetilde{\text{F-S}}$ and $\tilde{\mathbf{i}}' : \textsc{SubStoch} \to \widetilde{\text{F-S}}$, where $\tilde{\mathbf{e}}' = \sim_{\mathbf{p}^*} \circ \mathbf{e}'$ and $\tilde{\mathbf{i}}' = \sim_{\mathbf{p}^*} \circ \mathbf{i}'$. We can also introduce a prediction map $\tilde{\mathbf{p}^*}$ for the quotiented classical realist CI theory, whose action is given by mapping each process in $\widetilde{\text{F-S}}$ to an element (any element) of F-S in its equivalence class, and then mapping that element to SubStoch via $\mathbf{p}^*$. All of this can be concisely represented in the following commuting diagram:

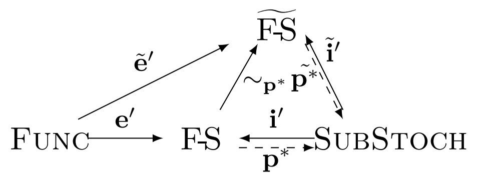

(6.182)

Finally, we derive a simplified normal form for $\widetilde{\text{F-S}}$. First, we show (in Appendix C.5, using some useful identities proven in Appendix C.4) that

Theorem 6.5.7. *Any diagram in* F-S *is always inferentially equivalent to one of the form*

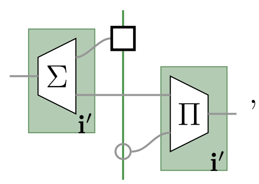

(6.183)

where Σ *is a stochastic map and* Π *is a propositional map.*

Applying the quotienting map (and recalling that it is diagram-preserving and that it leaves processes in SUBSTOCH invariant), this implies that

Corollary 6.5.7.1. *Any diagram in* $\widetilde{\text{F-S}}$ *can be rewritten into the following normal form:*

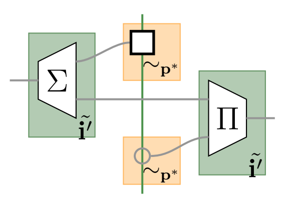

(6.184)

Subsuming the traditional notion of an ontological theory

We can now point out a connection between the notion of a quotiented classical realist CI theory and the traditional notion of an ontological model [139, 281]. Specifically, the notion of a quotiented classical reality CI theory subsumes the type of ontological theory that is presumed as the codomain of the traditional ontological modelling map. In particular, the stochastic processes in the codomain of this map (such as probability distributions and response functions over the ontic state space) are recovered as the substochastic maps defined by our Eq. (6.176).

Note, however, that the notion of a quotiented classical realist CI theory contains both causal and inferential systems, while traditional ontological models concern only a single type of system. We will say much more about representing operational theories in Section 6.6 and onward.

6.5.3 To make an omelette...

A causal-inferential theory allows one to describe a physical scenario while maintaining the distinction between the causal and inferential components of the theory. On the classical realist side, highlighting this distinction helps to identify the root of the puzzlement associated to phenomena such as 'Simpson's paradox' [275] and 'Berkson's paradox' [39], by making plain the sense in which previous frameworks have scrambled causal and inferential notions. Moreover, the conceptual and formal tools it provides can help to break the habits of mind that lead to such confusions, namely, the tendency to slide from statements about conditional probabilities to claims about cause-effect relationships.

On the operational side, the framework we have developed here helps to highlight a type of scrambling of causal and inferential concepts that seems intrinsic to any operational theory, and which is not so apparent in the conventional frameworks. It concerns the nature of a process in Proc, the causal component of an operational theory. Recall that these processes are descriptions of laboratory procedures. Now note that although specifying a laboratory procedure may serve to completely specify *some* degrees of freedom of the devices (usually macroscopic ones), it also generally involves expressing *incomplete* knowledge of the vast majority of its degrees of freedom (the microscopic ones).[14] As such, a process in Proc generally does not stipulate an *actual* causal relation between its inputs and outputs, given that it is consistent with many possibilities for this causal relation. It is for this reason that we have stipulated that a process in Proc describes only *potential* causal influences. The reason that the processes in the causal component of an *operational* CI theory—unlike those in the causal component of a *realist* CI theory—must in part stipulate an agent's uncertainty is because *such uncertainty is inherent in all descriptions of phenomena at the operational level.* At the operational level, therefore, the best one can hope to do is to unscramble the notion of *potential causal influence* from that of inference.

[14]For operational CI theories that admit of a classical realist representation, this fact is reflected in our assumptions about the realist representation map ξ, namely, that a point distribution over causal processes in the operational CI theory need not be mapped by ξ to a point distribution over causal processes in the realist CI theory.

In the rest of this section, we describe a different—and much more significant—example of how our framework reveals a type of scrambling of causal and inferential notions that previously went unnoticed. Specifically, we argue that quotienting with respect to inferential equivalence necessarily erases information about causal relations, so that quotiented theories inevitably incorporate a scrambling of causal and inferential notions.

We hope to make this point in more detail in a subsequent paper [259] whose purpose is to describe the sense in which GPTs are recovered from quotientied operational CI theories. For now, we will give a concrete example involving processes in F-S in order to clarify how causation and inference get scrambled under quotienting.

It is based on the example introduced in Section 6.5.2, involving the two states of knowledge about bit-to-bit functional dynamics described in Eq. (6.178) and that can be recast diagramatically as follows:

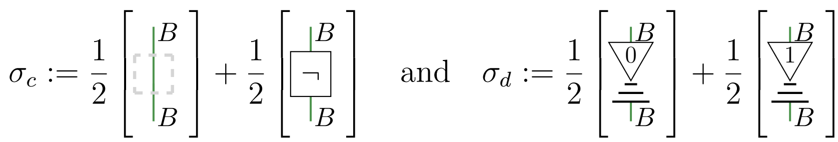

(6.185)

These two states of knowledge refer to completely distinct causal relations between the input and output bit: σ_c has support only on processes with a causal connection between the input and the output, while σ_d has support only on processes that are causally disconnected. Nonetheless, because the stochastic map associated with each of these states of knowledge is given by the completely randomizing bit-to-bit channel, as noted in Eq. (6.179), it follows that they are inferentially equivalent:

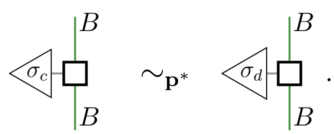

(6.186)

We see, therefore, that two radically different causal structures (causal connection and causal disconnection) lead one to make all of the same inferences about the relevant systems. Similar examples can be constructed in P-S. It is in this sense that the processes in quotiented causal-inferential theories, such as $\widetilde{\text{F-S}}$ and $\widetilde{\text{P-S}}$, exhibit a scrambling of causation and inference.

6.6 Classical realist representations

6.6.1 Classical realist representations of operational CI theories

A classical realist representation of an operational theory is an attempt to provide an underlying realist explanation of the operational statistics. It posits that each system is characterized by an ontic state, which constitutes a complete characterization of its physical attributes and mediates causal influences between the laboratory procedures. Something akin to such a representation in earlier work is the notion of an *ontological model of an operational theory*. (Strictly speaking, the closest analogue to the notion of an ontological model in our framework is not the notion of a classical realist representation that we consider in this section, but the variant thereof corresponding to the ζ map given below in Definition 15. The notion of ontological modelling is mentioned here only to help the reader broadly situate the notion of a classical realist representation.) [15]

Definition 12. *A* classical realist representation *of an operational CI theory* P-S*, by a classical realist CI theory* F-S*, is a diagram-preserving map* $\xi : \text{P-S} \to \text{F-S}$ *, depicted as*

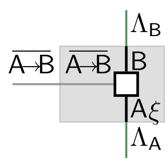

, (6.187)

satisfying **(i)** *the* preservation of predictions*, namely that the diagram*

$$\begin{array}{ccc} \text{P-S} & \overset{\mathbf{p}}{\dashrightarrow} & \textsc{SubStoch} \\ \downarrow \xi & & \Vert \\ \text{F-S} & \overset{\mathbf{p}^*}{\dashrightarrow} & \textsc{SubStoch} \end{array} \tag{6.188}$$

[15] Note that a classical realist representation of an operational scenario describes both its causal and inferential aspects, hence both ontological and epistemological aspects thereof. In this sense, it is clear that the term 'ontological modelling' would not be ideal for this sort of representation because the term suggests that one is concerned with modelling *only* the ontological aspects. Insofar as the sort of representation that was referred to as 'an ontological model of an operational theory' in prior work [139] *also* described both ontological and epistemological aspects (even if these were scrambled somewhat), the term was not ideal for those representations either. In retrospect, a better terminology would have been one that signaled that both ontological *and* epistemological aspects were being described therein, just as the term 'causal-inferential' signals a description incorporating both causal and inferential elements.

commutes, where the double line between the two copies of SUBSTOCH *is an extended equals sign, and* **(ii)** *the* preservation of ignorability

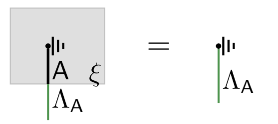

(6.189)

We will sometimes refer to the classical realist representation map ξ *as an* F-S*-representation of* P-S*.*

Note that ξ : P-S $\to$ F-S leaves inferential systems invariant. This can be derived from preservation of predictions, as follows. Start with some inferential system X in the top left of the commuting square (Eq. (6.188)). There are two paths to the bottom right: in one direction, we map via the prediction map $\mathbf{p}$ to the same system X in SUBSTOCH and then we map via the equality to the same system X in the other copy of SUBSTOCH; in the other direction, we map using the classical realist representation ξ to ξ_X in F-S and then map to SUBSTOCH via the prediction map $\mathbf{p}^*$ which leaves us with ξ_X. We then see that the only way that the diagram can commute is if $\xi_X = X$.

The fact that we take a classical realist representation to be diagram-preserving is an immediate consequence of our choice to take diagrams in an operational CI theory to represent one's hypothesis about the fundamental causal and inferential structure in the given scenario. Since an ontological representation is meant to be the most fundamental description of one's scenario, it should respect this hypothesis, with the only difference being that it will generally be a more fine-grained description (e.g., where laboratory procedures are replaced by functional dynamics). We will leave to Appendix C.2 the reason behind our choice to have operational CI diagrams represent fundamental structure, and we also show therein that this choice does not limit the scope of possible classical realist representations in our framework.

A particularly natural class of classical realist representations are those that can be thought of simply as representing every state of knowledge about a procedure by a corresponding state of knowledge about the function that underlies it.[16] Diagrammatically,

[16]One might wonder whether this is sufficiently general given that for a procedure mapping system A to system B, the variable Λ_B might not be a function of Λ_A alone but of Λ_A together with some local auxiliary variable Λ (whose value is drawn from some probability distribution). Such worries are unfounded, however, since every value of Λ defines a function from Λ_A to Λ_B, and a probability distribution over this value induces a probability distribution over the latter function.

these are represented as

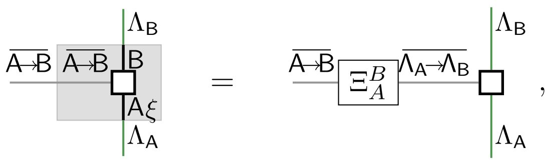

(6.190)

where Ξ_A^B is stochastic. and satisfies a set of compositionality constraints in order for ξ to be diagram-preserving, e.g.,

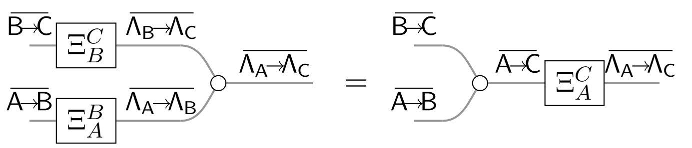

(6.191)

This class of classical realist representations is so natural, in fact, that one might even wish to demand that a classical realist representation be *defined* by such a constraint, although we have not done so here.

Whether it is part of the definition or not, it is often sufficient to focus on this class alone because it turns out that every classical realist representation is inferentially equivalent to one in this class.

Theorem 6.6.1. *Any classical realist representation ξ satisfies*

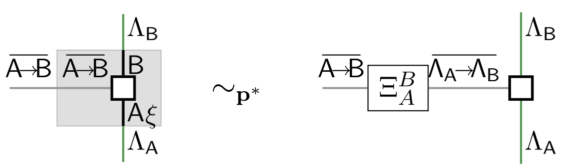

(6.192)

where Ξ_A^B is a stochastic map taking states of knowledge about operational procedures to states of knowledge about functional dynamics.

Proof. The proof is given in Appendix C.6. □

Nonetheless, it is not clear whether or not all classical realist representations, as defined in Definition 12, are of the form of Eq. (6.190). For example, classical realist representations of the form

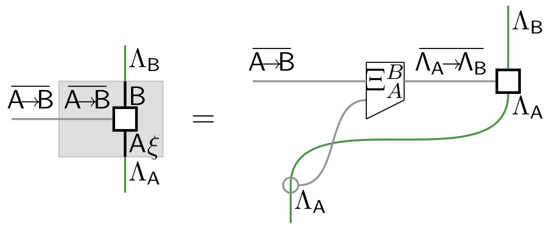

(6.193)

may be consistent with Definition 12, and appear to be more general than those of the form of Eq. (6.190). Such models, however, seem to fail to satisfy an assumption of autonomy—that the fundamental dynamics are independent of their inputs—and this may be grounds for dismissing them as candidates for a classical realist representation. It remains to be seen whether these can indeed be ruled out from our definition, or ruled out as a consequence of some formal notion of autonomy (which one might consider adding to Definition 12).

The question of the existence of a classical realist representation of an operational CI theory is closely connected to the pre-existing question of whether a given operational theory violates Bell-like inequalities. We explore the connection in Section 6.7.1.

6.6.2 Classical realist representations of quotiented operational CI theories

It is also useful to define classical realist representations of quotiented operational CI theories.

Definition 13. *A* classical realist representation *of a quotiented operational CI theory,* $\widetilde{\text{P-S}}$, *by a quotiented classical realist CI theory,* $\widetilde{\text{F-S}}$, *is a diagram-preserving map* $\tilde{\xi} : \widetilde{\text{P-S}} \to \widetilde{\text{F-S}}$, *depicted as*

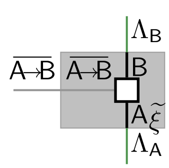

, (6.194)

satisfying **(i)** *the* preservation of predictions, *namely that the diagram*

$$\begin{array}{ccc} \widetilde{\text{P-S}} & \overset{\widetilde{\mathbf{p}}}{\dashrightarrow} & \textsc{SubStoch} \\ \Big\downarrow \widetilde{\xi} & & \Big\| \\ \widetilde{\text{F-S}} & \overset{\widetilde{\mathbf{p}}^*}{\dashrightarrow} & \textsc{SubStoch} \end{array} , \tag{6.195}$$

commutes, where the double line between the two copies of SubStoch *is an extended equals sign, and* **(ii)** *the* preservation of ignorability

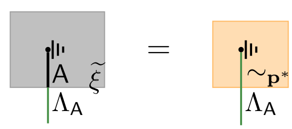

. (6.196)

We will sometime refer to the classical realist representation $\widetilde{\xi} : \widetilde{\text{P-S}} \to \widetilde{\text{F-S}}$ *as an* $\widetilde{\text{F-S}}$-*representation of* $\widetilde{\text{P-S}}$.

Representations of this sort are analogous to the simplex embedding maps introduced in Ref. [262] in the context of prepare-measure scenarios. (These mapped the states and effects of the GPT to probability distributions and response functions over the ontic state space in an ontological theory.)

Proposition 1. *The classical realist representation map* $\widetilde{\xi}$ *can be written as*

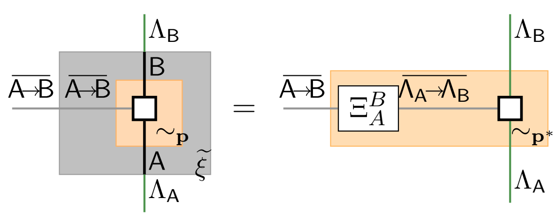

, (6.197)

where Ξ_A^B *is a stochastic map taking states of knowledge about operational procedures to states of knowledge about functional dynamics.*

Proof. The proof is a direct adaptation of the proof of Theorem 6.6.1, but where the starting point, Eq. (6.484), is modified by replacing the inferential equivalence with equality and using the normal form for $\widetilde{\text{F-S}}$ as given by Corollary 6.5.7.1 , and where one uses the form of Lemma C.1 which involves equality rather than inferential equivalence. □

6.6.3 Leibnizianity (formalized)

A natural methodological principle to impose on candidate realist explanations of operational facts is the following [285]:

> If an ontological theory implies the existence of two scenarios that are empirically indistinguishable in principle but ontologically distinct (where both indistinguishability and distinctness are evaluated by the lights of the theory in question), then the ontological theory should be rejected and replaced with one relative to which the two scenarios are ontologically identical.

In Ref. [285], it is argued that this methodological principle was proposed by Leibniz as a version of his principle of the identity of indiscernibles and that it was strongly endorsed (at least implicitly) by Einstein. We shall refer to it here as *Leibniz's methodological principle.*

Although this statement of Leibniz's principle has been argued to motivate the standard definition of generalized noncontextuality, it is necessary to consider an epistemological *generalization* of Leibniz's principle in order to motivate generalized noncontextuality in the rehabilitated form that we will endorse below. Specifically, instead of concerning pairs of scenarios that are empirically indistinguishable, the generalized principle concerns pairs of processes in a causal-inferential theory (such as states of knowledge) that are equivalent in the sense of allowing an agent to make precisely the same inferences. We formalize the new version of the principle as a constraint on realist representations, which we term *Leibnizianity.*

Definition 14 (Leibnizianity of a classical realist representation)**.** *A classical realist representation map* $\xi : \text{P-S} \to \text{F-S}$ *is said to be* Leibnizian *if it preserves inferential equivalence relations. Otherwise, it is said to be* nonLeibnizian*.*

More formally, a classical realist representation map ξ is Leibnizian if, for any pair of inferentially equivalent processes $\mathcal{D}, \mathcal{E} \in$ P-S, one has

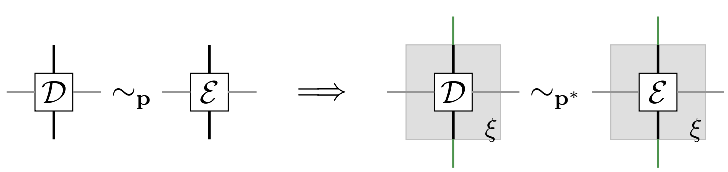

(6.198)

This means that if two processes lead one to make the same inferences when embedded into any diagram within the operational CI theory, then their representations within the classical realist CI theory must be such that they lead one to make all the same inferences when embedded into any diagram within the classical realist CI theory. As an example of an application of the principle, Leibnizianity stipulates that *inferentially equivalent states of knowledge about experimental procedures must be represented by inferentially equivalent states of knowledge about functional dynamics.*

It is straightforward to verify from the definitions that an equivalent (process-theoretic) characterization of Leibnizianity is the following.

Proposition 2. *A classical realist representation* $\xi : \text{P-S} \to \text{F-S}$ *of the unquotiented operational CI theory, is Leibnizian if and only if there exists a classical realist representation* $\widetilde{\xi} : \widetilde{\text{P-S}} \to \widetilde{\text{F-S}}$ *of the quotiented operational CI theory, such that the following diagram commutes:*

$$\begin{array}{ccc} & & \widetilde{\text{P-S}} \\ & \nearrow^{\sim_{\mathbf{p}}} & \\ \text{P-S} & & \downarrow \widetilde{\xi} \\ \downarrow \xi & & \widetilde{\text{F-S}} \\ & \nearrow^{\sim_{\mathbf{p}}} & \\ \text{F-S} & & \end{array} \quad . \tag{6.199}$$

We will ultimately be interested in contemplating the possibility of realist CI theories that are nonclassical alternatives to F-S (see Sec. 6.9.4), so we will find it useful to sometimes refer to a representation in terms of F-S as simply an F-S-representation. A given operational CI theory may admit of both Leibnizian and nonLeibnizian F-S-representations. It will be termed F-S-*Leibnizian-representable* if it admits of at least *one* Leibnizian F-S-representation.

We have just described the implication $\exists$ Leibnizian $\xi \implies \exists\, \widetilde{\xi}$. Whether the implication $\exists$ Leibnizian $\xi \impliedby \exists\, \widetilde{\xi}$ holds remains an open question, but we conjecture that it does:

Conjecture 1. *If a quotiented operational CI theory admits of a classical realist representation (as a quotiented classical realist CI theory), then the unquotiented operational CI theory admits of a Leibnizian classical realist representation (as an unquotiented classical realist CI theory). More formally, if there exists a map* $\widetilde{\xi} : \widetilde{\text{P-S}} \to \widetilde{\text{F-S}}$ *satisfying Definition 13, then there exists a map* $\xi : \text{P-S} \to \text{F-S}$ *satisfying Definition 12, and which makes the diagram of Eq. (6.199) commute.*

Hence, we have

$$\exists \text{ Leibnizian } \xi \quad \overset{?}{\Longleftrightarrow} \quad \exists\, \widetilde{\xi}. \tag{6.200}$$

Note that a classical realist representation $\widetilde{\xi}$ of a *quotiented* operational CI theory cannot itself be said to be either Leibnizian or nonLeibnizian. This is because in quotiented CI theories, there are no distinct but inferentially equivalent processes, for which one could ask whether their representations are inferentially equivalent or not.

In Refs. [281, 285], it was argued that it is the plausibility of the Leibnizian methodological principle that accounts for the plausibility of the principle of generalized noncontextuality defined in Ref. [281]. Indeed, our notion of Leibnizianity turns out to have a close connection to the notion of generalized noncontextuality. We discuss this connection in Section 6.8.1.

6.6.4 Summary of the basic framework

The full set of process theories and DP maps that we have introduced can be summarized by the following diagram (which is color-coded to match the intuitive schematic given further down):

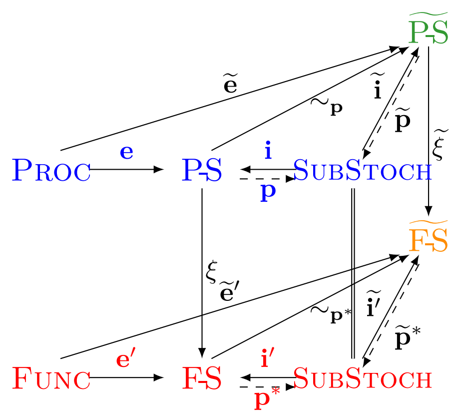

. (6.201)

The top slice of this describes an operational CI theory (in blue) and its quotienting (in green), while the bottom slice describes a classical realist CI theory (in red) and its quotienting (in yellow). The *unquotiented* classical realist and operational CI theories are constructed out of their respective causal theory (Func or Proc) together with the classical theory of inference (SubStoch), which is common to both. The map ξ (if it exists) constitutes a representation of an operational CI theory P-S by a classical realist CI theory F-S, while the map $\widetilde{\xi}$ (if it exists) constitutes a representation of a quotiented operational CI theory $\widetilde{\text{P-S}}$ by a quotiented classical realist CI theory $\widetilde{\text{F-S}}$. If both ξ and $\widetilde{\xi}$ exist and the square Eq. (6.199) commutes, then we say that ξ is a Leibnizian classical realist representation.

Fig. (6.201) is summarized by the following schematic:

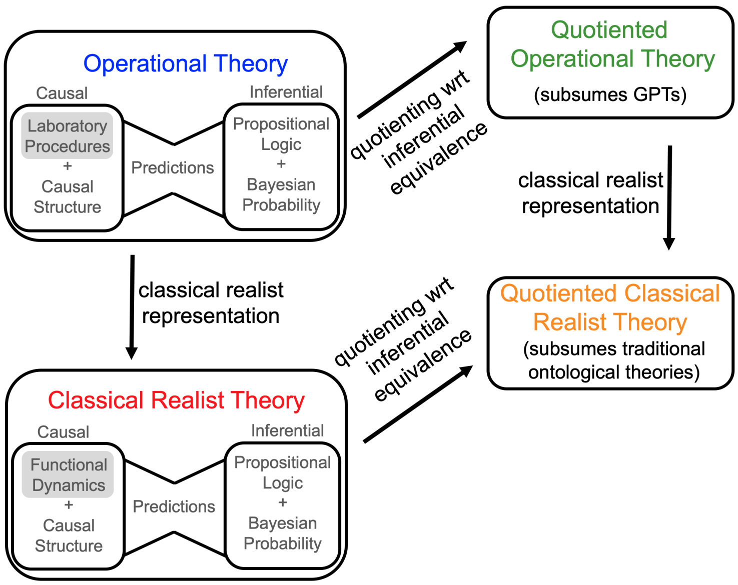

The form of Diagram (6.201) might lead one to wonder about whether there ought to be a diagram-preserving map from Proc to Func. One could define such a map, but it would state that every procedure could be associated with a unique function acting on the ontic states. As noted in Sec. 6.5.3, this is not what we expect of a classical realist model of an operational theory, as laboratory procedures typically constitute a *coarse-grained* description, and as such are not associated with a unique function, but rather a distribution over these.

We refer the reader to Appendix C.1 for a discussion of prior work that is related to (or provided inspiration for) our framework.

6.7 Bell-like no-go theorems

6.7.1 Bell-like inequalities as a consequence of assuming the existence of a classical realist representation

Consider a bipartite Bell experiment where X and Y denote the setting variables, and A and B denote the outcome variables. Suppose one takes the causal structure of a Bell experiment to be given by the following directed acyclic graph, or DAG, where the triangle depicts an unobserved common cause:

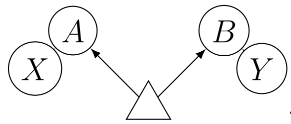

(6.202)

This corresponds to the diagrams

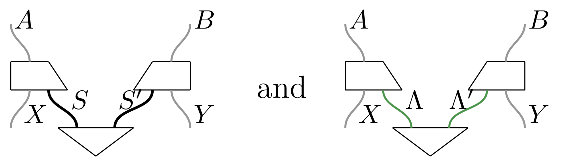

(6.203)

in Proc and Func respectively.

A priori, this assumption about causal structure of the Bell experiment is the natural one. It is motivated by the idea that relativity implies no superluminal causation (not just a prohibition on superluminal signals). We refer to it as the *common-cause hypothesis* [12, 256, 322] regarding the causal structure. In Section 6.7.2, we will consider alternatives to it.

What about the full causal-inferential structure? In P-S and F-S these are

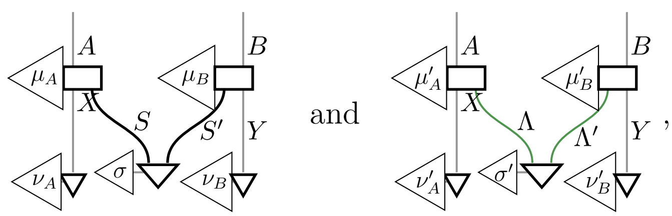

(6.204)

respectively, where we have allowed for arbitrary states of knowledge μ_A, μ_B, ν_A and ν_B, and σ about the procedures (respectively μ'_A, μ'_B, ν'_A and ν'_B, and σ' about the functions), but we have not allowed for any statistical dependencies between the identities of the procedures (nor, therefore, between the identities of the functions). For instance, the factorization of ν_A and σ (and hence of ν'_A and σ') is motivated by the implausibility of nature conspiring to ensure that the mechanism that sets the value of the setting variable is related to the mechanism that fixes the value of the common cause. (This is related to superdeterminism, discussed further in the next section.)

We now discuss how the causal-inferential hypotheses embodied in Eq. (6.204) constrain the possible observations that can be made within the classical realist theory, as well as those that can be made within a given operational theory.

First, we consider the question of what joint distributions over X, Y, A and B can be generated in the causal-inferential structure of Eq. (6.204) within the classical realist CI theory F-S. These are given by the diagram

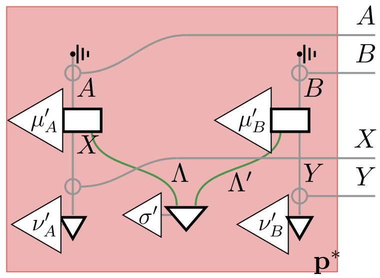

, (6.205)

where one ranges over arbitrary sets Λ and Λ' and probability distributions σ', μ'_A, μ'_B, ν'_A, and ν'_B. We refer to any distribution that arises in this way as F-S-*realizable*. The constraints that pick out the set of F-S-realizable distributions generally come in the form of both equalities and inequalities, and we will term these F-S-*compatibility constraints*. In particular, inequality constraints will be termed F-S-*compatibility inequalities*.

Within our framework, the Bell inequalities (e.g., the Clauser-Horne-Shimony-Holte inequalities for the case where X, Y, A and B are binary) are examples of F-S-compatibility inequalities for the causal-inferential structure of Eq. (6.204) (which, as noted earlier, is the natural choice for the Bell experiment).

Next, we turn to the question of what joint distributions over X, Y, A and B can be generated in the causal-inferential structure of Eq. (6.204) within an operational CI theory P-S. These are given by the diagram

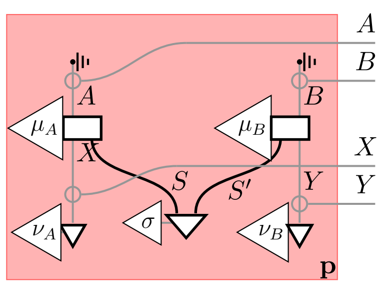

, (6.206)

where one ranges over arbitrary systems S and S' in P-S and probability distributions σ, μ_A, μ_B, ν_A, and ν_B. Note that the set of distributions that can be obtained in this way depends on the prediction map $\mathbf{p}$ of P-S, and so will vary from one operational CI theory to the next. We refer to any distribution that can arise in this way within an operational CI theory P-S as P-S-*realizable*. The constraints that pick out the set of P-S-realizable distributions will be termed P-S-compatibility constraints. In particular, any inequality constraints will be termed P-S-*compatibility inequalities*.

For the case of quantum theory, considered as the operational CI theory P$_{\mathrm{Q}}$-S introduced in Section 6.4.1, the well-known Tsirelson inequalities [74] are examples of P$_{\mathrm{Q}}$-S-compatibility inequalities for the causal-inferential structure of Eq. (6.204).

If a classical realist representation as in Theorem 6.6.1 exists for a given operational theory P-S, it follows that every distribution that is P-S-realizable is also F-S-realizable. Thus, for a given operational CI theory P-S to admit of a classical realist representation, it must be the case that the set of P-S-realizable distributions for any possible causal-inferential structure is included in the F-S-realizable distributions for the same causal-inferential structure.

For the case of quantum theory, there exist distributions [44] that satisfy the Tsirelson inequalities but violate the Bell inequalities, i.e., that are P$_{\mathrm{Q}}$-S-compatible but not F-S-compatible with the causal-inferential structure of Eq. (6.204). It follows that the set of P$_{\mathrm{Q}}$-S-realizable distributions is not included in the set of F-S-realizable distributions, and consequently that P$_{\mathrm{Q}}$-S does not admit of a classical realist representation.

This is how Bell's theorem is conceptualized in our framework.

When Bell's theorem is conceptualized in this way, it is found to have counterparts in causal structures distinct from that of the Bell scenario. For instance, recent work has demonstrated that for the triangle scenario [104, 106, 236, 323], whose DAG and associated diagram in PROC are

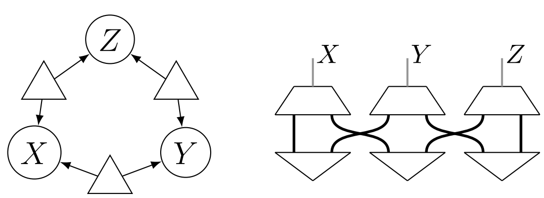

, (6.207)

there is a gap between what is realizable in a causal model where some common causes can be quantum and what is realizable in a causal model where all common causes are classical. A similar result has been shown for the instrumental scenario [63, 307] , whose DAG and associated diagram in PROC are

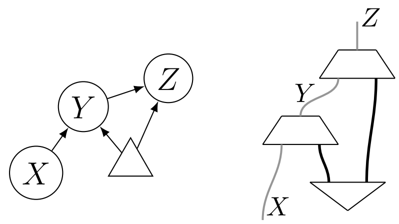

$$. \tag{6.208}$$

These new Bell-like no-go theorems are subsumed in our framework as proofs of the impossibility of a classical realist CI representation based on the correlations predicted by quantum theory for these causal scenarios. The counterpart within our framework to realizability by a *quantum* causal model [12, 29, 106, 142, 321, 323] is P_QS-realizability, while the counterpart within our framework to realizability by a *classical* causal model is F-S-realizability. Thus, the counterpart to these no-go theorems is that for each of these causal structures, one can find distributions over the observed variables that are P_QS-realizable, but not F-S-realizable.

6.7.2 The conventional ways out of Bell-like no-go theorems

We now describe the standard responses to Bell-like no-go theorems, focusing on the specific case of Bell's theorem (rather than those based on, e.g., the instrumental or triangle scenarios).

A common attitude towards Bell's theorem is that it demonstrates that realism must be abandoned (at least in the quantum sphere) hence vindicating an operationalist philosophy of science. As discussed in the introduction, we take the key distinction between a realist and a purely operational account of statistical correlations to be whether or not these accounts provide a causal explanation of the correlations. We believe that any philosophy of science that is antirealist in this sense—namely, which denies the possibility of a causal explanation of statistical correlations—is unsatisfactory. Furthermore, given the possibility (which we will introduce in Section 6.9) for nonclassical realist representations that modify the notions of causation and inference used in one's causal explanation, we are certainly not persuaded by any claim that the standard no-go theorems *necessitate* a retreat from realism in the quantum sphere. For these reasons, we are here interested in reviewing the standard ways of maintaining realism in the face of Bell's theorem, in order to contrast them with the way we propose to do so.

For those unwilling to compromise on realism, the conventional way out is to deny the natural causal-inferential hypothesis of Eq. (6.204). They assume therefore that a *radical* causal-inferential hypothesis underpins the correlations observed in a Bell scenario. In this case, the existence of a classical realist representation does not imply satisfaction of the Bell inequalities, and thus violations of Bell inequalities no longer imply a challenge to the possibility of such a representation. We now describe the two most common positions about the nature of the radical causal hypotheses.

The hypothesis that a proponent of superluminal causation endorses. Those who see superluminal[17] causation as the way out of Bell's theorem take the causal structure of a Bell experiment to be one that allows for causal influences between the wings, even when these are space-like separated. As an example, this influence might be between the setting on the left wing and the outcome of the right, a causal hypothesis that is depicted by the following DAG, or equivalently, by the following circuit diagram:

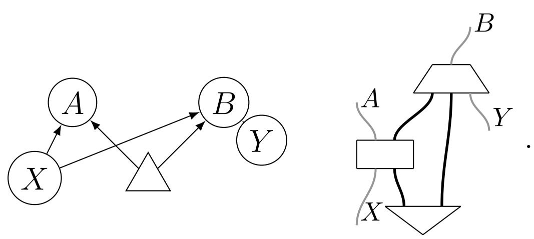

(6.209)

The hypothesis that a proponent of superdeterminism endorses. Those who see superdeterminism as the way out of Bell's theorem take there to be some statistical dependence between a setting variable, say X, and the common cause of the outcomes. This assumption could be encoded in either the causal or the inferential structure. We here opt to encode it as the assumption that X and the common cause of the outcomes are not causally disconnected. This assumption can be depicted by the following DAG, or equivalently, by the following circuit diagram:

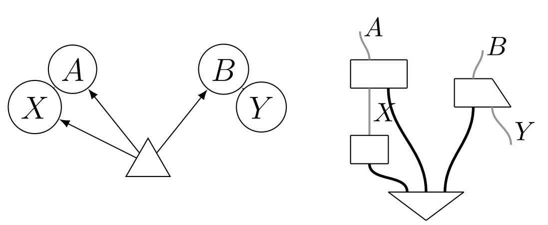

(6.210)

[17]Formally describing the distinction between sub- versus superluminality motivates a minor extension of our framework in which systems come equipped with spatiotemporal labels. We leave this for future work.

At this stage, we would like to head off a possible confusion. Because it is customary within pre-existing operational frameworks to use the standard quantum circuit of Eq. (6.203) as the diagram representing the Bell experiment, it might seem that a proponent of a radical causal hypothesis must be contemplating a realist representation that fails to be diagram-preserving, and thus it might seem that our framework, which assumes diagram preservation, cannot do justice to their view. However, as noted previously and elaborated on in Appendix C.2, our framework stipulates that the diagram representing a given scenario in the causal subtheory of an operational CI theory is a representation of one's hypothesis about the fundamental causal structure, and hence *need not* correspond to the standard quantum circuit of Eq. (6.203). As such, researchers with different realist worldviews, faced with the same experimental scenario and observed statistics, might model these *differently* within the framework of operational causal-inferential theories. For example, the representation of a Bell scenario within a quantum operational CI theory will be constrained by the causal structure of Eq. (6.209) by a proponent of superluminal causation, but will be constrained by the causal structure of Eq. (6.210) by a proponent of superdeterminism. (Furthermore, the manner by which such researchers would formalize quantum theory as an operational CI theory will not be the one described in Section 6.4.1.)

Criticisms of the conventional ways out of Bell-like no-go theorems

We now discuss various reasons why we view these conventional responses to Bell-like no-go theorems as unsatisfactory.

The following are grounds for rejecting the causal-inferential hypothesis of the proponent of superluminal causation:

- superluminal causal influences are in tension with the spirit of relativity theory (even if these influences are constrained in such a way as to be consistent with the impossibility of superluminal signals)
- superluminal causal influences that are constrained in such a way as to be consistent with the impossibility of superluminal signals violate the principle of no fine-tuning (defined in Section 6.10.2). [324].

Meanwhile, the following are grounds for rejecting the causal-inferential hypothesis of the proponent of superdeterminism:

- for any of the various mechanisms that could determine the value of the setting variables (an agent's free choice, a random number generator, a hash of the day's stock prices, etcetera), it is implausible that it would be related to the mechanism that determines the common cause of the outcome variables insofar as this would require a kind of conspiracy of causal determinations

- the fact that any such nontrivial statistical associations between a setting variable and the common cause of the outcome variables must be constrained to be consistent with the observed *lack* of any statistical assocation between that setting variable and the outcome variable at the opposite wing of the experiment implies violation of the principle of no fine-tuning [324].

In fact, in Ref. [324] it was shown that *every* causal hypothesis that can realize the distributions predicted by quantum theory (i.e., Bell-inequality-violating distributions) using a classical causal model (i.e., via F-S) implies a violation of the principle of no fine-tuning.

We take these arguments to be good grounds for rejecting the idea of explaining correlations in Bell scenarios via a radical causal-inferential hypothesis together with a classical realist representation.

Even if one rejects the principle of no fine-tuning, standard tools of model selection (which are sensitive not just to underfitting but overfitting as well) can adjudicate between various hypotheses regarding the right way to operationally model the Bell experiment, and these also rule against a radical causal hypothesis [91].

The no-fine-tuning arguments just given apply equally well to no-go theorems based on causal structures beyond Bell scenarios, e.g., the instrumental and triangle scenarios. That is, one can attempt to resolve the contradiction in each of these cases by hypothesizing that the causal structure is in fact distinct from that depicted in Eq. (6.207) and Eq. (6.208). Such resolutions, however, also generally suffer from a fine-tuning objection insofar as the set of distributions that are $\mathrm{P_Q}$S-realizable for the original causal structure often have measure zero within the set of distributions that are F-S-realizable for the radical causal structure, and they will in some instances also require superluminal causation.

Insofar as these conventional ways out of Bell's theorem require radical causal-inferential assumptions with the aforementioned undesirable features, the natural question becomes: *does there remain any recourse for achieving a realist causal-inferential representation of quantum theory without these unappealing features?* In Section 6.9, we outline a research program for achieving realism *while preserving the conservative causal-inferential hypothesis*, by allowing for intrinsically nonclassical notions of influence and inference.

Before coming to this, however, we discuss how no-go theorems based on the principle of noncontextuality appear within our framework.

6.8 Noncontextuality no-go theorems

6.8.1 Generalized noncontextuality (rehabilitated)

The pre-existing notion of generalized noncontextuality [281] is framed as a constraint on ontological models of operational theories and can be summarized as "*operationally equivalent procedures must be represented identically in the ontological theory*". However, our framework has refined the traditional notions of ontological theories and of operational theories, and the notion of operational equivalence of procedures has been replaced by inferential equivalence of states of knowledge about procedures. The notion of generalized noncontextuality must therefore be refined accordingly.

Generalized noncontextuality is a principle that constrains *a* classical realist representation map, but it is *not* the map ξ that we have been focused on so far. Rather, it is a constraint on a map $\zeta : \text{P-S} \to \widetilde{\text{F-S}}$ from the *unquotiented* operational CI theory P-S to the *quotiented* classical realist CI theory $\widetilde{\text{F-S}}$. The fact that this is a map across the unquotiented-quotiented divide, means that this specific sort of classical realist representation is less fundamental than the map ξ. Nonetheless, it is useful to introduce this map as a formal tool, e.g., in order to make connections to existing literature. Paralleling Definitions 12 and 13, we define ζ as follows:

Definition 15. *A classical realist representation of an unquotiented operational CI theory,* P-S*, by a quotiented classical realist CI theory,* $\widetilde{\text{F-S}}$*, is a diagram-preserving map* $\zeta : \text{P-S} \to \widetilde{\text{F-S}}$ *satisfying* **(i)** *the* preservation of predictions*, namely that the diagram*

$$\begin{array}{ccc} \text{P-S} & \overset{\mathbf{p}}{\dashrightarrow} & \textsc{SubStoch} \\ \downarrow \zeta & & \Vert \\ \widetilde{\text{F-S}} & \overset{\widetilde{\mathbf{p}}^*}{\dashrightarrow} & \textsc{SubStoch} \end{array} , \qquad (6.211)$$

commutes, where the double line between the two copies of SubStoch *is an extended equals sign, and* **(ii)** *the* preservation of ignorability

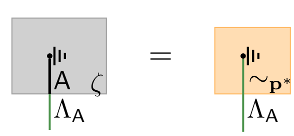

(6.212)

We will sometimes refer to the classical realist representation $\zeta : \text{P-S} \to \widetilde{\text{F-S}}$ *as an* $\widetilde{\text{F-S}}$*-representation of* P-S*.*

Adding the ζ map to the diagram relating P-S, F-S, and their quotiented counterparts yields

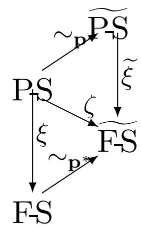

. (6.213)

Note that the ζ map is the closest counterpart in our framework to the pre-existing notion of an *ontological model of an operational theory* [139].

It follows that the closest counterpart in our framework to the pre-existing notion of a *generalized-noncontextual* ontological model is that of a generalized-noncontextual classical realist representation of this sort. This can be defined in terms of the formal notion of Leibnizianity introduced in Definition 14, but applied to ζ (rather than ξ):

Definition 16 (Generalized-noncontextual classical realist representation)**.** *The classical realist representation map* $\zeta : \text{P-S} \to \widetilde{\text{F-S}}$ *is generalized-noncontextual if it preserves inferential equivalence relations.*

In the case of a single causally closed process, this can be summarized as "*inferentially equivalent states of knowledge about experimental procedures must be represented by the same stochastic map.*"

We now give an equivalent (process-theoretic) characterization of a generalized-noncontextual classical realist representation, in analogy to Proposition 2.

Proposition 3. *A classical realist representation map* $\zeta : \text{P-S} \to \widetilde{\text{F-S}}$ *is generalized-noncontextual if and only if there exists a map* $\widetilde{\xi}$ *as defined in Definition 13 such that the upper right triangle in Eq.* (6.213) *commutes, implying that the map* ζ *can be factored as* $\widetilde{\xi} \circ \sim_{\mathbf{p}}$*.*

It follows that if there exists a map $\widetilde{\xi}$, then there exists a generalized-noncontextual map ζ (namely $\zeta = \widetilde{\xi} \circ \sim_{\mathbf{p}}$). In fact, the opposite implication holds as well because the claim that ζ is generalized-noncontextual means *by definition* that there exists a map $\widetilde{\xi}$

such that the upper right triangle in Eq. (6.213) commutes. We can therefore summarize the relationship as follows:

$$\exists \text{ generalized-noncontextual } \zeta \quad \Longleftrightarrow \quad \exists\, \widetilde{\xi}. \tag{6.214}$$

Note that insofar as a quotiented operational CI theory subsumes the notion of a GPT, this fact is the analogue, for the rehabilitated version of generalized noncontextuality of Proposition 1 of Ref. [260], which asserts that an operational theory admits of a generalized-noncontextual ontological model if and only if the GPT defined by it admits of an ontological model.

To re-emphasize a point made in Ref. [260], a *quotiented* operational CI theory is not the sort of thing that can be either generalized-noncontextual or generalized-contextual because information about context is precisely what is eliminated by the quotienting operation. In other words, it is a category mistake to ask whether $\widetilde{\xi}$ is generalized-noncontextual or not. Thus, for any experiment that does not support a Bell-like no-go theorem but does support a noncontextuality no-go theorem, its model within P-S is *always* consistent with *some* classical realist representation ξ given the representational freedom that is afforded by context-dependence. Its model within $\widetilde{\text{P-S}}$, however, might not admit of *any* classical realist representation.

Old proofs of the failure of generalized noncontextuality will imply proofs of the failure of the rehabilitated version of generalized noncontextuality, since all that has changed is how one conceptualizes the mathematics. This is particularly clear given Eq. (6.214) and the close connection between the question of the existence of a $\widetilde{\xi}$ map in our framework and the question of whether a given GPT model of an experiment admits of an ontological model.

In previous work, generalized noncontextuality was defined case-by-case for various types of procedures (e.g. preparations, measurements, transformations [281], and instruments [228, 231, 260]), and it was then stipulated that the natural assumption of generalized noncontextuality is the *universal* version of this assumption, meaning for *all* types of procedures. In contrast, our process-theoretic characterization of generalized noncontextuality applies to *all* types of experimental procedures including more exotic processes such as combs and circuit fragments with arbitrary causal structure, and therefore is a *universal* notion from the get-go.

6.8.2 Failures of generalized noncontextuality imply failures of Leibnizianity

We now discuss the relation between the rehabilitated notion of generalized noncontextuality and the notion of Leibnizianity.

While the rehabilitated notion of generalized noncontextuality is a principle that constrains the representation map ζ, the notion of Leibnizianity is a principle that constrains the representation map ξ. It is like generalized noncontextuality insofar as it can be defined in terms of the commutation of the diagram, but it *cannot* be understood as a notion of independence on context, as we now explain.

We begin by defining the notion of context that is at play in the notion of generalized noncontextuality. For any process within a CI theory, its *context* is the information that determines which element of an inferential equivalence class of processes it is—that is, it is information about a process that is irrelevant for making predictions. Note that both operational CI theories *and classical realist CI theories* have nontrivial contexts: in either case, a full specification of a state of knowledge over the relevant causal processes describes both the equivalence class and the context of a procedure. An explicit example of two processes in P-S that are in the same inferential equivalence class but which differ by context is given by the two states of knowledge about operational procedures described in Eq. (6.151). An explicit example of two processes in F-S that are in the same inferential equivalence class but which differ by context is given by the two states of knowledge about functions described in Eq. (6.179). In this language, what Leibnizianity demands is that an operational process's context can only determine the context of its image under the realist representation map, not the inferential equivalence class of its image under the realist representation map. Even when a given representation map ξ satisfies this condition, the context of an operational process's image under ξ *can* depend on the context of the operational process, and hence, it would be inappropriate to call the map 'context-independent' or 'noncontextual'. Thus, the inapplicability of the term 'noncontextual' for describing the relevant constraint on the representation map ξ is seen to be a consequence of the fact that the fundamental notion of a realist CI theory is an *unquotiented* one, which has contexts.

In contrast, the realist representation map ζ has contexts in its domain but not its co-domain, so that what Leibnizianity demands in this case is that an operational process's context cannot determine anything about its image under ζ. Thus, a map ζ can be either generalized-noncontextual or generalized-contextual, depending on whether or not the image of an operational process under this map is independent of the context of the operational process.

We turn now to the formal relationships between the two notions. Combining Eq. (6.200) and Eq. (6.214), we infer that

$$\exists \text{ Leibnizian } \xi \overset{?}{\Longleftrightarrow} \exists \text{ gen.-noncontextual } \zeta \Longleftrightarrow \exists \, \widetilde{\xi}. \tag{6.215}$$

Accordingly, Conjecture 1 can be reformulated as follows:

Conjecture 2. *If an operational CI theory admits of a generalized-noncontextual classical realist representation as a quotiented classical realist CI theory, then it admits of a Leibnizian classical realist representation as an unquotiented classical realist CI theory. More formally, if there exists a map* $\zeta : \text{P-S} \to \widetilde{\text{F-S}}$ *satisfying Definition 15 and which makes the upper right triangle in the diagram of* (6.213) *commute, then there exists a map* $\xi : \text{P-S} \to \text{F-S}$ *satisfying Definition 12 and which makes the square in the diagram of* (6.213) *commute.*

The existence of a Leibnizian classical realist representation ξ implies the existence of a generalized-noncontextual classical realist representation ζ. Contrapositively, the nonexistence of such a ζ implies the nonexistence of such a ξ, and therefore every no-go theorem for generalized-noncontextual classical realist representations yields a no-go theorem for Leibnizian classical realist representations.

What about implications in the other direction? If Conjecture 2 is false, then there may be proofs of the impossibility of a Leibnizian classical realist representation that are not proofs of the impossibility of a generalized-noncontextual classical realist representation. In this case, there may be novel no-go theorems for classical realist representations of operational quantum theory. By contrast, if the conjecture is true, then every no-go theorem based on Leibnizianity yields a no-go theorem based on generalized noncontextuality.

6.8.3 Reframing the standard no-go theorem for generalized noncontextuality

To illustrate how the notion of noncontextuality appears in our framework, we consider a simple prepare-measure scenario. Let the setting variable for the measurement be denoted by Y and the outcome be denoted by B. In many discussions of noncontextuality, the preparation device is imagined to have just a setting variable. In order to achieve more symmetry between the preparation device and the measurement device, however, it is convenient to consider a preparation device that has not only a setting variable, but an outcome as well [171]. We denote the setting by X and the outcome by A. The causal structure of a prepare-measure is presumed to be the following

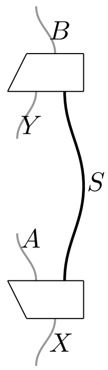

and

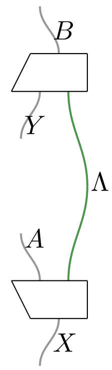

(6.216)

where we have shown the representations in Proc and Func respectively.

The full causal-inferential structure is represented as follows in P-S and F-S respectively:

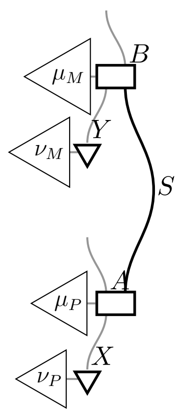

and

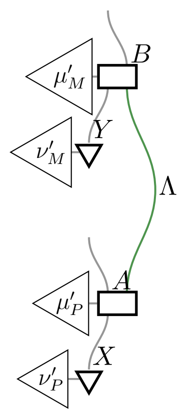

(6.217)

where we have allowed for arbitrary states of knowledge μ_P, μ_M, ν_P, and ν_M about the procedures (respectively μ'_P, μ'_M, ν'_P, and ν'_M about the functions),

The P-S-realizable joint distributions over X, Y, A and B for the causal-inferential structure of Eq. (6.217) are those given by the diagram

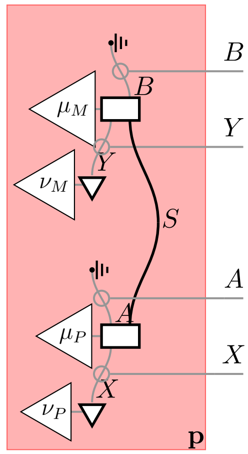

(6.218)

where one ranges over arbitrary system S in P-S and probability distributions μ_P, μ_M, ν_P, and ν_M.

The F-S-realizable distributions over X, Y, A and B are those given by

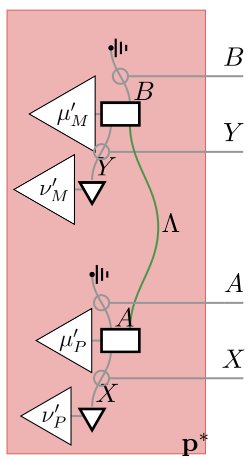

(6.219)

where one ranges over an arbitrary set Λ and probability distributions μ'_P, μ'_M, ν'_P, and ν'_M.

It turns out that for this causal-inferential structure, the only restrictions on F-S-realizable distributions over X, Y, A and B are that X and Y must be independent and that A and Y are independent. So there is no opportunity for a Bell-like no-go result for this causal-inferential structure.

Nonetheless, one *can* prove a noncontextuality no-go result in such structures. The reason is that proving a noncontextuality no-go is not about proving the impossibility of an F-S-representation of P-S, as is the case for proving a Bell no-go result. Rather, it is about proving the impossibility of finding an F-S-representation of P-S *that is Leibnizian*, as given in Definition 14.

In the case of quantum theory, that is, P_Q-S, one can prove such a result using the causal-inferential structure of (6.219), by leveraging no-go theorems for generalized noncontextuality in prepare-measure scenarios. One can of course also consider the consequences of Leibnizianity for scenarios beyond the prepare-measure variety.

Whether an operational theory P-S admits of a generalized-noncontextual $\widetilde{\text{F-S}}$-representation coincides with several pre-existing notions of classical explainability. This follows from the fact that the existence of a generalized-noncontextual $\widetilde{\text{F-S}}$-representation of P-S implies the existence of an $\widetilde{\text{F-S}}$-representation of $\widetilde{\text{P-S}}$ and the latter coincides with the following two notions of classical explainability: simplex-embeddability of the GPT describing prepare-measure experiments [262, 272] and the existence of a positive quasiprobability representation of the GPT [103, 260, 283].

Additionally, an operational theory P-S that fails to admit of a generalized-noncontextual classical realist representation provides advantages for information processing relative to those that do admit of such a representation [14, 60, 144, 154, 191, 192, 234, 248, 249, 255, 263, 286]. This bolsters the notion that *failing* to admit of a generalized-noncontextual $\widetilde{\text{F-S}}$-representation is a good notion of *non*classicality.

In light of the relationships proven above between Leibnizianity and generalized noncontextuality, each of these results concerning generalized noncontextuality can be repurposed as a motivation for assuming that the existence of a Leibnizian F-S-representation is a good notion of classicality for an operational CI theory P-S.

To summarize, our framework yields a new perspective on the relationship between Bell-like and noncontextuality (or Leibnizianity) no-go theorems. Both types of no-go theorems concern the representability of an operational CI theory P-S in terms of a classical realist CI theory F-S. A Bell-like no-go theorem is a demonstration that there does not exist *any* classical realist representation map $\xi : \text{P-S} \to \text{F-S}$ as in Definition 12. A noncontextuality no-go theorem, on the other hand, is a demonstration that there does not exist such a map *that is Leibnizian*—that is, one wherein the inferential equivalence relations are preserved, as in Definition 14. Hence we see that Bell-like no-go theorems are based on a weaker assumption. Nonetheless, in our view, the stronger assumption of Leibnizianity is just as plausible. Furthermore, noncontextuality no-go theorems have greater breadth of applicability than their Bell-like counterparts since they can be proven for a broader set of causal-inferential structures—even those involving just a single causal system.

6.8.4 The conventional way out of noncontextuality no-go theorems

What is the conventional response to the lack of the existence of a generalized-noncontextual classical realist representation ζ of quantum theory, considered as the operational CI theory $\mathrm{P_QS}$? (Or equivalently, to the lack of the existence of a classical realist representation $\widetilde{\xi}$ of quantum theory, considered as the quotiented operational CI theory $\widetilde{\mathrm{P_QS}}$?) For those who are unwilling to compromise on the standard notion of a realist representation, the typical response is to endorse a failure of generalized noncontextuality.

Endorsing such a failure requires a renouncement of Leibnizianity. Thus, to anyone who is committed to Leibniz's principle, this 'way out' of the no-go results will be unappealing.

There is, however, the possibility of an alternative to the conventional response, one that aims to salvage Leibnizianity within a realist representation. The idea is the one already noted at the end of Section 6.7.2, namely, to underlie operational quantum theory with a realist causal-inferential theory wherein the causal and inferential components are intrinsically nonclassical. The next sections take up this research program.

6.9 Beyond classical realism

In the conclusions of the last two sections, we criticized the conventional ways out of Bell-like and noncontextuality no-go theorems on the grounds that the price they must pay to salvage the standard notion of realism—violating the principle of no-superluminal causation, violating the principle of no fine-tuning, and abandoning the Leibnizian methodological principle—is too high. We also noted that this motivates a new type of research program, wherein one seeks to salvage these principles by considering novel notions of realist representations wherein the causal and inferential components thereof become intrinsically nonclassical. The hope is that a realist causal-inferential theory of this type will have enough in common with its classical counterpart that a representation in terms of it can nonetheless be judged to provide satisfactory *explanations* of the operational phenomena. (In previous work, this idea has been described as 'achieving realism while going beyond the standard ontological models framework' [284].) Up until now, the constraints that such a representation must satisfy have been articulated only vaguely, if at all. The framework of causal-inferential theories allows us to say much more about the nature of such a representation and hence about how to further this research program.

Suppose that a nonclassical analogue of FUNC is denoted XFUNC and a nonclassical analogue of SubStoch is denoted XSubStoch and that the nonclassical realist causal-inferential theory defined by the interaction of these is denoted XF-XS. We have:

$$\text{XFunc} \xrightarrow{\mathbf{e}'} \text{XF-XS} \underset{\mathbf{p}^*}{\overset{\mathbf{i}'}{\leftrightarrows}} \text{XSubStoch}\ , \tag{6.220}$$

where the maps $\mathbf{e}'$, $\mathbf{i}'$ and $\mathbf{p}^*$ play the same roles that they do in F-S.

The question becomes: what properties must the causal-inferential theory XF-XS have in order to be considered realist, that is, such that representability in terms of such a theory can be considered to provide a *causal explanation* of the observed correlations? These properties will help to identify what alternatives there might be to representing physical systems by sets, propositions about these by subsets, states of knowledge by distributions over a set and causal determination by functions from one set to another.

Here, some readers might worry that all such alternatives should be ruled out by the very meanings of the terms in question: might it be that one *cannot even speak* about systems, logical propositions and states of knowledge about systems, and causal influences among systems, *unless* systems are described by sets, propositions by subsets, states of knowledge by distributions, and causal determination by functions? In short, some might argue that the definitions of notions of inference and causation are *analytic*, and therefore immune to revision. But this is not the case, a point we make by an analogy with nonEuclidean geometry.

The idea is that a putative nonclassical realist causal-inferential theory XF-XS will stand to the classical realist causal-inferential theory F-S as a nonEuclidean geometry stands to Euclidean geometry.[18] Just as the meanings of the terms 'point' and 'line' in a nonEuclidean geometry are determined from the axioms of that geometry rather than corresponding to the common-sense notions, so too will the meaning of various causal and inferential concepts within a given nonclassical realist causal-inferential theory XF-XS be determined by the specific axioms of that process theory (i.e., the diagrammatic rewrite rules) rather than corresponding to the conventional ones. In this sense, we are embracing the attitude towards mathematical structure that is characteristic of category theory and that contrasts with the attitude of set theory wherein everything is built up from concepts concerning sets. In particular, the fact that systems in a nonclassical realist causal-inferential theory are not associated with sets does not imply that such systems cannot be the locus of

[18]Hilary Putnam famously used this analogy to describe how a quantum logic ought to be conceptualized relative to classical logic [232]. We are simply extending the analogy to describe how probabilistic inference and causal influence ought to be conceptualized as well.

causal influences or the subject of propositions and states of knowledge. Note that the attitude towards scientific realism that such a research program presumes is cognate with the philosophical position of structural realism [172].

Any attempt to provide a nonclassical generalization of the notions of causation and inference, however, is highly constrained insofar as it will need to preserve those features of these notions which one judges to be essential. These constraints are the analogue of constraints on nonEuclidean geometries that arise from what one takes to be essential to the notions of 'point' and 'line'. They are constraints on the process theory which, if violated, might well lead one to question whether the theory really is describing causal influences and inferences after all. Note that there is no way to be certain about the appropriateness of such constraints a priori, because without concrete modification of the classical theory having been proposed and shown to be coherent and useful, it is difficult to know which of the features of the classical theory are essential.

6.9.1 Constraints a causal-inferential theory must satisfy to be considered realist

Given that we have distinguished the notions of operational and realist causal-inferential theories, it is clear that (in our view) a generic causal-inferential theory—in particular an *operational* CI theory—does not contain enough structure to be deemed worthy of the title 'realist'. While an operational CI theory can predict observations, it does not itself provide a *realist explanation* of those predictions. Thus, we do not consider P-S to be an instance of XF-XS. In this section, we highlight the additional structure possessed by a classical realist CI theory F-S over and above that possessed by an operational CI theory P-S. This structure helps to identify the properties one should demand of a causal-inferential theory XFunc in order that it be deemed 'realist'.

Because P-S and F-S are built out of the same inferential subtheory, SubStoch, the contrast between them reduces to the contrast between the causal subtheories out of which they are built, Proc and Func respectively, and to differences in how these interact with SubStoch.

In Func, aspects of the causal structure are encoded not just in the shape of the circuit but also in the identities of the functions. For example, if a process corresponds to a function that is independent of its argument, then there is no causal connection between the input and the output of that process. In other words, the function associated with some process specifies the causal structure *internal* to the process.

In Proc, on the other hand, the internal structure of a process is not specified. A given process does not necessarily even have a causal influence from its inputs to its outputs—it is only that there is *potential* for such a causal influence. Some information about the internal causal structure may be inferred from the image of this process under the prediction map, but this does not generally provide a full specification of the internal causal structure of the processes (e.g., we saw in Section 6.5.1 that inferentially equivalent processes can correspond to different causal structures).

To summarize, Proc encodes *potential* causal influences, while Func encodes *actual* causal influences. What we have termed a *causal theory* is meant as an umbrella for these two notions.[19]

Formally, the issue is that the interpretation of processes in a process theory is derived primarily from their interactions with other processes—through the nontrivial equalities that involve them. But Proc is a free process theory, with no nontrivial equalities; hence, processes in Proc have an interpretation that is impoverished relative to those in Func.

Thus a first criterion for a CI theory to be deemed realist is the following:

1. The causal subtheory of the CI theory must have enough nontrivial equalities such that its processes represent actual causal influences rather than potential causal influences.

The exact formalization of this remains to be determined, but we give constraints on how to do so later in this section. A minimal requirement is that XFunc is not a free process theory.

We now turn to a comparison of the interaction between Proc and SubStoch and the interaction between Func and SubStoch.

Func and SubStoch exhibit strong forms of interaction. For example, one can define propositions about (or equivalently, directly gain information about) *any* causal system in F-S, since the generator in Eq. (6.124) is defined for all systems in a classical realist CI theory. In contrast, the interaction between Proc and SubStoch is very limited, insofar as one cannot define propositions about (or directly learn information about) any systems that are nonclassical. This leads to our next criterion for a CI theory to be deemed realist:

[19] It is worth noting that the assumption of diagram-preservation for the classical realist representation ξ ensures the preservation of the structure of *potential causal influences*, not that of actual causal influences. That is, if there is no potential causal influence between a pair of systems in some given diagram of Proc, then the image of this diagram under ξ must be such that there is no actual causal influence between these in Func. If, however, there *is* a potential causal influence between a pair of systems in some given diagram of Proc, then there may or may not be an actual causal influence between these in Func.

2. Propositions must be able to attach to all systems in the CI theory.

Formally, the fact that F-S satisfies this criterion has important consequences. Chiefly, it is a prerequisite for one to introduce the equality Eq. (6.128), which allows the translation of a proposition about the output of a causal mechanism into a proposition about its input and the identity of the mechanism.

By contrast, there is no equality analogous to Eq. (6.128) in operational CI theories. Specifying states of knowledge about the causal processes in an operational CI theory does not yield statistical predictions until one specifies a prediction map. Indeed, nearly all of the non-generic features of an operational CI theory are buried within the choice of prediction map.[20]

We therefore elevate this feature into a criterion of its own:

3. It should be possible to propagate knowledge claims through any causal mechanism. Formally, there must exist an analogue of Eq. (6.128).

This last criterion is central to the idea of a realist causal-inferential theory. A commitment to realism means that the systems can mediate causal influences, and that one can understand every valid inference as a consequence of knowledge propagation *through* these causal mediaries. In particular, if it is the existence of a causal pathway between two variables that accounts for the inferences that can be made between these, then it must be possible to understand these inferences as decomposable into a sequence of inferences, stepping through systems along the causal pathway. For example, in a Bell scenario (as in Diagram (6.206)), updating one's knowledge of the outcome at the left wing, A (which depends on background knowledge about X), leads to an updating of one's knowledge of the outcome on the right wing, B (which depends on background knowledge about Y), via the mediary of systems S and S'. Specifically, updating one's knowledge of A leads to an updating of one's knowledge of S, which in turn leads to an updating of one's knowledge of S', which in turn leads to an updating of one's knowledge of B. The ability of systems to encode information and to be mediaries in a sequence of refinements of knowledge is key, we argue, for a given theory to be described as realist.

The equality in Eq. (6.128) leads to a great deal of the structure of F-S and ultimately to the uniqueness of the prediction map, as in Theorem 6.4.3. It seems an essential part of any

[20]Note that if an operational CI theory P-S *does* admit of a classical realist representation in terms of F-S, then this representation serves to provide an explanation for the prediction map of P-S in terms of the unique prediction map of F-S and the representation map ξ.

fundamental theory of nature that the predictions one makes should be uniquely determined by a complete causal-inferential description of one's scenario within that fundamental theory.

Hence, we have another criterion for a CI theory to be deemed realist:

4. The CI theory must have a unique prediction map.

Our final criterion for a CI theory to be deemed realist also relies on the fact that one can attach propositions to any system, as per Criterion 2. Recall that two processes are inferentially equivalent if no matter in what causally-closed circuit they are embedded, they lead one to make the same predictions concerning the classical variables that are the (inferential) inputs and outputs of that circuit. This is an 'external' characterization of inferential equivalence. Within a classical realist theory, Lemma 6.5.6 showed that there is a second equivalent characterization of inferential equivalence that is 'intrinsically' defined. This intrinsic characterization hinges on a particular mapping, Eq. (6.176), from arbitrary processes in F-S to processes in SubStoch. This allowed us to characterize the inferential equivalence of processes in F-S via the stochastic maps that naturally describe the intrinsic relationship between the causal and inferential inputs and outputs of the process, without reference to external scenarios involving the process.

If Criterion 2 is satisfied within a nonclassical realist theory, then one can also define a mapping analogous to that in Eq. (6.176), where the mapping takes every process $\mathcal{D}$ in XF-XS to a related process in XSubStoch, such that the latter provides an intrinsic description of inferences from the causal and inferential inputs of $\mathcal{D}$ to the causal and inferential outputs of $\mathcal{D}$. These are analogous to the stochastic maps just described in the case of classical realist theories. The criterion, then, is that one can define this intrinsic characterization of inferential equivalence, and that it furthermore be equivalent to the external characterization:

5. Two arbitrary processes in a CI theory are inferentially equivalent if and only if the external and intrinsic characterizations of inferential equivalence coincide. Formally, there must exist an analogue of Eq. (6.176) and Lemma 6.5.6.

The challenge moving forward, then, is to find mathematical structures XFunc and XSubStoch that respect all of the desiderata required for a CI theory to be deemed realist.

Concepts that should have analogues in any realist theory— One expects any putative nonclassical theory of causation to contain analogues of most, if not all, of the standard notions that arise in the framework of classical causal models: common causes, causal mediaries, d-separation, evaluation of counterfactuals, etcetera. Preliminary work towards establishing how the evaluation of counterfactuals is formally achieved within F-S is provided in Section 6.10.2.

One also expects that any putative nonclassical theory of inference should contain analogues of most, if not all, of the standard notions that arise in Boolean logic and Bayesian probability theory: logical connectives, implication, conditional independence, sufficient statistics, etcetera. [21]

6.9.2 Nonclassical realist representations

Having described what it means for a causal inferential theory to embody a satisfactory notion of realism, we can now describe the notion of a *nonclassical realist representation*: namely, a representation of an operational CI theory in terms of a nonclassical realist CI theory. The definition is the analogue of Definition 12, but where the image of the map is a nonclassical realist CI theory rather than a classical one. This definition (given below) is the sense in which we have now formalized the idea of 'achieving realism while going beyond the standard ontological models framework'.

For an arbitrary operational CI theory P-S, we can seek to find a nonclassical realist CI theory XF-XS in terms of which P-S can be represented. That is, we can ask if there is a realist representation map $\xi : \text{P-S} \to \text{XF-XS}$ that can be defined in an analogous way to Def. 12.

Note, however, that whereas the inferential subtheories of F-S and of P-S were identical (namely, SubStoch), the inferential subtheory of XF-XS is allowed to be something more general than that of P-S, namely, what we have denoted XSubStoch. Consequently, one needs to modify the condition of preservation of empirical predictions (the commutation of Eq. (6.188)). Rather than the two inferential subtheories being equal, they will be related by some sort of map $\phi : \text{SubStoch} \to \text{XSubStoch}$ whose defining properties is a subject for future work [22]. Given a satisfactory definition of ϕ, one can define:

[21] Note that we have not yet incorporated all of these notions in SubStoch at a diagrammatic level, although it is clear how all of these can be determined nondiagrammatically. Incorporating these into the diagrammatic formalism (or demonstrating that they are dispensable without compromising the usefulness of the theory of inference) is therefore a topic for further research. Incorporating them explicitly would enable the study of how these inferential features interact with causal processes.

[22] In the cases of primary interest to us, we expect that ϕ will be an inclusion map. Hence, it will still be possible and meaningful to make classical inferences within XSubStoch.

Definition 17. *A* nonclassical realist representation *of an unquotiented operational CI theory* P-S *by an unquotiented nonclassical realist CI theory* XF-XS *is a diagram-preserving map* $\xi : \text{P-S} \to \text{XF-XS}$ *satisfying* **(i)** *the* preservation of predictions*, namely that the diagram*

$$\begin{array}{ccc} \text{P-S} & \overset{\mathbf{p}}{\dashrightarrow} & \textsc{SubStoch} \\ \downarrow \xi & & \downarrow \phi \\ \text{XF-XS} & \overset{\mathbf{p}'}{\dashrightarrow} & \textsc{XSubStoch} \end{array} \quad . \tag{6.221}$$

commutes, and **(ii)** *the* preservation of ignorability

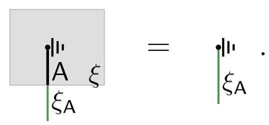

(6.222)

We will sometimes refer to the nonclassical realist representation map ξ *as an* XF-XS*-representation of* P-S*.*

Analogously, one can also extend the notion of a representation of a *quotiented* operational CI theory $\widetilde{\text{P-S}}$ in terms of a quotiented classical realist CI theory $\widetilde{\text{F-S}}$ to that of a representation in terms of a quotiented *nonclassical* CI theory $\widetilde{\text{XF-XS}}$. Again, the only non-trivial aspect of this generalization is in defining the map $\phi : \textsc{SubStoch} \to \textsc{XSubStoch}$.

One can summarize the notions of nonclassical realist representations via the analogue of Eq. (6.201):

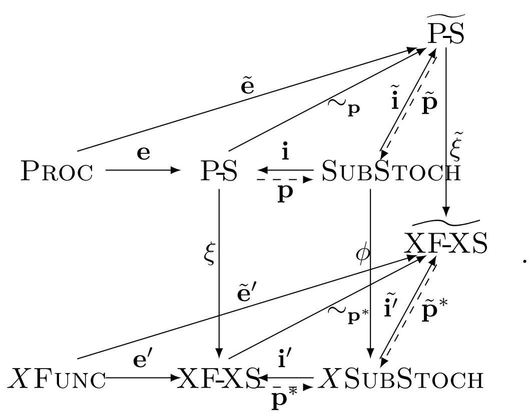

(6.223)

6.9.3 A new way out of Bell-like no-go theorems

Recall that a Bell-like no-go theorem arises whenever one finds a causal structure in which the set of P-S-realizable probability distributions is not contained in the set of F-S-realizable probability distributions—that is, not contained among those that can be generated by a *classical* realist representation.

The possibility of *nonclassical* realist representations provides a novel way out of such no-go theorems. Rather than asking if the observed experimental statistics are F-S-realizable, one can instead ask if they are XF-XS-realizable, that is, representable using a map $\xi : \text{P-S} \to \text{XF-XS}$ into some *nonclassical* realist CI theory XF-XS.

If this can be done, it seems appropriate to claim that such a realist representation has salvaged locality. More precisely, such a representation has provided a means of being conservative with respect to causal structure—in particular, not requiring superluminal influences—by being radical with respect to the nature of the realist CI theory.

We have not yet provided an explicit proposal for a realist CI theory XF-XS that can reproduce the quantum predictions while providing a satisfactory realist explanation of them. However, our work in formalizing the notion of a nonclassical realist theory, e.g., via the formal criteria given in Section 6.9.1, constitutes a first concrete step in this direction.

6.9.4 A new way out of noncontextuality no-go theorems

A nonclassical realist CI theory XF-XS necessarily includes a notion of inferential equivalence—because XF-XS is assumed to provide a unique prediction map, one simply evaluates equivalences relative to it. It follows that one can define Leibnizianity for nonclassical realist representations much as it was defined for classical realist representations (in Section 6.6.3):

Definition 18 (Leibnizianity of a nonclassical realist representation)**.** *A nonclassical realist representation map* $\xi : \text{P-S} \to \text{XF-XS}$ *is said to be* Leibnizian *if it preserves inferential equivalence relations.*

Consequently, it makes just as much sense to ask whether a given operational CI theory admits of a *nonclassical* realist representation that is Leibnizian as it did to ask that question of a classical realist representation. This is a key benefit of our new process-theoretic definition of Leibnizianity.

As in the case of classical realist representations, we can give an equivalent characterization in terms of a commuting square. A nonclassical realist representation map

$\xi : \text{P-S} \to \text{XF-XS}$ is Leibnizian if and only if there exists a map $\widetilde{\xi} : \widetilde{\text{P-S}} \to \widetilde{\text{XF-XS}}$ such that the following diagram commutes:

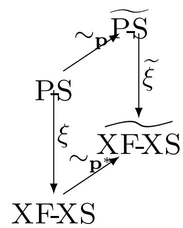

(6.224)

where $\mathbf{p}^*$ is the unique prediction map in XF-XS.

One can extend the notion of generalized-noncontextuality to nonclassical realist representations in a similar fashion. Specifically, the map $\zeta : \text{P-S} \to \widetilde{\text{XF-XS}}$ is defined to be generalized-noncontextual if the triangle in upper right of the following diagram commutes:

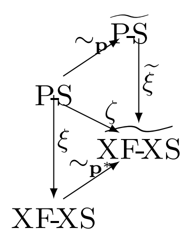

(6.225)

The fact that the definitions are all process-theoretic implies that we have abstract notions of generalized noncontextuality and Leibnizianity that apply to nonclassical realist representations and that satisfy analogues of Eq. (6.215), that is,

$$\exists \text{ Leibnizian } \xi \overset{?}{\Longleftarrow\!\!\!\Longrightarrow} \exists \text{ gen.-noncontextual } \zeta \iff \exists\, \widetilde{\xi}. \tag{6.226}$$

The possibility of nonclassical realist representations therefore holds the potential for a novel way out of noncontextuality no-go theorems, a way out that does not compromise on the Leibnizian methodological principle.

6.10 Discussion

6.10.1 Finding a satisfactory ontology and epistemology *for quantum theory*

The long-term aim of this work is to generate a compelling interpretation of quantum theory—one that satisfies the spirit of locality and Leibnizianity. We turn now to the special case of nonclassical realist representations *of quantum theory*.

Recall from Section 6.4.1 that quantum mechanics can be cast as an operational CI theory $\mathrm{P_QS}$ having the structure

$$\textsc{Proc}_Q \xrightarrow{\mathbf{e}} \mathrm{P_QS} \underset{\mathbf{p}_Q}{\overset{\mathbf{i}}{\rightleftarrows}} \textsc{SubStoch} \ . \tag{6.227}$$

Our ultimate objective is to identify a *quantum* realist CI theory[23]

$$\mathrm{Q}\textsc{Func} \xrightarrow{\mathbf{e}'} \mathrm{QF\text{-}QS} \underset{\mathbf{p}^*}{\overset{\mathbf{i}'}{\rightleftarrows}} \mathrm{Q}\textsc{SubStoch} \ , \tag{6.228}$$

which satisfies the constraints articulated in Section 6.9.1 (so that it can meaningfully be said to be a realist theory) and where QF-QS further provides a Leibnizian representation of $\mathrm{P_QS}$.

If such a nonclassical realist representation map $\xi : \mathrm{P_QS} \to \mathrm{QF\text{-}QS}$ is found, then it follows that for any given causal-inferential structure, the set of QF-QS-realizable distributions includes the $\mathrm{P_QS}$-realizable distributions. Hence, one obtains a way out of Bell-like no-go theorems that is more satisfactory than the conventional ways out insofar as it need not involve any superluminal influences (thereby salvaging the spirit of locality), and insofar as it need not avail itself of any fine-tuning of parameters. Since we have further required that the realist representation map $\xi : \mathrm{P_QS} \to \mathrm{QF\text{-}QS}$ must be *Leibnizian* (as implied by the existence of the map $\tilde{\xi}$), one also obtains a way out of the noncontextuality no-go theorems that is more satisfactory than the conventional way out, insofar as it salvages the Leibnizian methodological principle (and thereby the spirit of generalized noncontextuality).

Although there has been some work on interpreting some of the formalism of quantum theory as a nonclassical generalization of Bayesian inference, the kind of classical theory

[23]Naturally, the 'Q' in the notation where we previously put 'X' refers to the fact that we are aiming specifically for a *quantum* generalization of Func, SubStoch, and F-S.

that served as the target of this generalization was an ontological model. As noted in Section 6.5.2, however, the notion of an ontological model corresponds to a *quotiented* classical realist CI theory. And, as argued in Section 6.5.3, such a theory involves a partial scrambling of causal and inferential concepts. This is problematic because it is likely that the constraints on putative quantum generalizations of classical theories are *only* clear if causal and inferential notions are cleanly separated in the latter, and hence only if these quantum generalizations are pursued at the level of the *unquotiented* theory.

As an example, consider how the project of finding a nonclassical generalization of Bayesian inference was pursued in Ref. [182], which built upon ideas proposed in Refs. [177, 180]. The focus was on finding intrinsically quantum counterparts to the notions of joint, marginal and conditional probability distributions, as well as counterparts to the relations that hold between these, such as the counterpart of marginalization, the law of total probability, and the formula for Bayesian inversion. However, a conditional probability distribution, or equivalently, a stochastic map, represents an *inferential equivalence class* of states of knowledge about functional dynamics, and often involves a partial scrambling of causal and inferential concepts (as illustrated in Section 6.5.3). The fact that the focus of this earlier work was on a mathematical object that scrambled causal and inferential concepts may explain why there are outstanding problems with the approach, such as those described in Ref. [182] and in Ref. [152]. A state of knowledge about functional dynamics—unlike the inferential equivalence class of such objects—involves no such scrambling. It is consequently *this* object that is more appropriate to focus on and for which to seek an intrinsically quantum counterpart.

It is possible that the proposals for quantum generalizations of propositional logic which were pursued under the banner of 'quantum logic' [148, 149] also suffer from having mistaken inferential equivalence for identity. Certainly, we believe that conventional approaches, such as the one that takes the counterpart of a Boolean lattice to be an orthomodular lattice, are unlikely to yield success in the research program described here. This is because such approaches are informed solely by the structure of projectors on Hilbert space and this may well merely be describing aspects of the *quotiented* quantum realist CI theory, while it is only the *unquotiented* theory QF-QS that one can hope to decompose into a causal subtheory QFunc and an inferential subtheory QSubStoch (where the structure concerning propositional logic lives).

We now highlight some prior work that is likely to be useful in developing an intrinsically quantum notion of a realist CI theory.

On the causal side, recent work on developing an intrinsically quantum notion of a causal model [12, 29, 88] is likely to provide a good starting point for finding the correct quantum generalization of Func. In particular, the notion of decomposing a unitary gate into a more refined circuit that includes 'dots' (isomorphisms wherein a Hilbert space is decomposed into a direct sum of tensor products), introduced in Ref. [12] and studied in depth in Ref. [190], is likely to be incorporated in some way into QFunc.

In pursuing the correct quantum generalization of SubStoch, recent work developing a synthetic approach to probability theory (formalized as 'Markov categories') [72, 109, 114, 155] is likely to be useful. This is because if the Bayes subtheory of SubStoch can be characterized more abstractly, the possibilities for quantum generalization should become more evident. In particular, the work of Ref. [85], which is in the same spirit as Refs. [72, 109], may provide an important piece of the puzzle (in spite of not having the benefit of a proper unscrambling of causal and inferential notions). Specifically, the *logical broadcasting* map described therein may be the counterpart in QSubStoch of the copy operation in SubStoch.

Similar comments may well apply to prior work in the field of quantum logic, namely, that in spite of suffering from some causal-inferential scrambling, specific insights from that research program could prove useful in finding the counterpart within QSubStoch to various notions within the subtheory Boole of SubStoch.

To close, we note that there has been a great deal of interest in whether certain mathematical objects in the quantum formalism—most notably quantum states—have an ontological or an epistemological status [59, 101, 112, 113, 139, 176–178, 230, 282]. Although disentangling ontology and epistemology is certainly critical to the project of unscrambling Jaynes' omelette, it is worth noting that in some cases this question presumes a false dichotomy. To see this, note that even in a *classical* realist CI theory, certain mathematical objects play multiple roles—for example, functions appearing in Func describe the causal influence that one variable has on another, while the same functions in SubStoch (now represented as deterministic stochastic maps) describe how learning about one variable leads to updating one's knowledge of another. It seems likely, therefore, that certain mathematical objects in a quantum realist CI theory will also have counterparts in both the causal and inferential subtheories. Indeed, a single-system unitary is likely to be such an object, sometimes describing the nature of a causal influence in the causal subtheory and sometimes the nature of how one updates one's knowledge in the inferential subtheory. The question about whether a given mathematical object in the quantum formalism has an ontological or epistemological status, therefore, must sometimes be refined to take into account the context in which the mathematical object appears.

6.10.2 Subsuming the framework of classical causal modeling

We have considered two distinct classes of causal theories, namely, Proc and Func. The primary technical distinction between these two is that Func has equalities, while Proc does not. We have seen various consequences of this extra structure on Func, e.g., the uniqueness of the prediction map in F-S. Conceptually, the primitive type of process in Proc (a list of lab instructions) constitutes an extremely minimal description of the causal mechanism relating its inputs to its outputs. Meanwhile, the primitive type of process in Func (a functional dependence) constitutes a much more informative description.

In fact, causal dependences in a classical theory can be *defined* in terms of functional dependence of one variable on another. This is done, for instance, in *structural equation models* [216], and it is the notion of classical causation that we endorse here. Hence, one might expect that structural equation models could be subsumed in our framework within F-S, which allows for both a description of the functional (hence causal) dependences among variables, as well as a specification of one's knowledge about exogenous variables. Similarly, the notion of a probabilistic causal model (or 'causal Bayesian network') [216], wherein the functional dependences and the states of knowledge of the exogenous variables are not specified individually, but are folded together into a conditional probability distribution, is likely to be subsumed in our framework within the quotiented theory $\widetilde{\text{F-S}}$. In future work, we hope to explore the relationship between our framework and various notions of classical causal models, and to argue that in some regards, our framework is more general than the standard one.

To accommodate all of the purposes to which classical causal models are put (in particular, considering the consequences of interventions and evaluating counterfactuals), it will be useful to introduce a distinct type of causal theory of functional dynamics, embedded within Func, which we will term PreFunc. The systems and processes in PreFunc are the same as in Func, but the process theory is defined without equalities. In particular, the composition of two functions $f(\cdot)$ and $g(\cdot)$ in sequence in PreFunc is *not* strictly equal to the function $f(g(\cdot))$. One can then define a DP map from PreFunc to Func that induces an equivalence relation on PreFunc, namely, two diagrams in PreFunc are equivalent if they define the same function when the component functions in the diagram are composed.

Remark 17. *The transition from* Func *to* PreFunc *can be viewed as an example of a very general construction on process theories. First, one defines a forgetful functor from the category* ProcessTheory *to a new category (which we will call* ProcessSet*) where a particular process theory is mapped to its underlying set of processes thereby forgetting the compositional structure of the process theory. We can then define a free functor that is left adjoint to the forgetful functor. The composition of these two functors then defines a* *comonad on* ProcTheory *which, in particular, takes* Func *to* PreFunc*. This is closely related to [219, Ex. 4.2.2].*

To better understand the differences between PreFunc and Func, consider the following pair of diagrams, where the gate $\bullet\!\!-\!\!\oplus$ represents a classical controlled NOT operation:

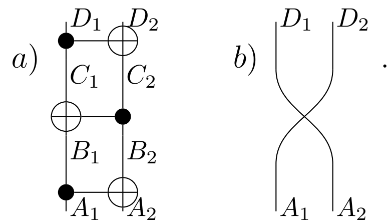

(6.229)

In Func, the process described by diagram (a) and that described by diagram (b) are strictly equal. The two diagrams are merely distinct manners of specifying the overall input-output functionality of the effective function from A_1 and A_2 to D_1 and D_2. In PreFunc, however, diagrams represent 'histories' of processes, rather than merely representing input-output functionalities. These two diagrams viewed within PreFunc are therefore not equal to one another, but rather are merely equivalent in the sense defined just above.

Despite the fact that (a) and (b) are equal within Func, it is clear that the interventions possible on each of them are distinct. To formally describe all possible interventions in a given scenario, it is essential that one works within PreFunc, wherein (a) and (b) are merely equivalent; e.g., this allows one to consider an intervention on B_1, B_2, C_1 or C_2.

In order to provide a fully formal diagrammatic treatment of the interventional aspects of the framework of classical causal models [216], it will likely be useful to take the causal theory to be PreFunc rather than Func. We will address this project more explicitly in future work.

Insofar as our work reveals that stochastic matrices (equivalently, conditional probability distributions) relating a cause to its effect generically scramble together causal and inferential concepts, this is true even for the notion of a *do-conditional*, which is defined as the conditional probability distribution of an effect variable given a cause variable when the value of the cause is intervened upon, rather than being determined by its natural causal parents. From the perspective of our work, the only object that does not lose any information about what is known about the causal influence of one variable on another is the probability distribution over the function that relates the one variable to the other, while the do-conditional merely describes an inferential equivalence class of such objects. This and related ideas will also be explored in future work.

6.10.3 More future directions

There are many natural extensions of our work, and many ways in which it is likely to shed light on other research programs. We now discuss some of these research directions, beyond those highlighted in the discussion sections or introduced throughout this chapter.

We first note a straightforward supplementation to the notion of an operational CI theory. Recall that an operational CI theory shares the inferential subtheory in common with the classical realist CI theory, but the causal subtheory Proc is distinct from Func. Nonetheless, because there is a distinction within Proc between classical systems (the settings and outcomes of procedures) and general systems, one can imagine a supplementation of Proc wherein the classical systems and all processes thereon have all of the structure of Func. Although the inferential consequences of this structure can in principle be obtained by encoding it in the prediction map, the framework is more useful if the additional equalities are present within Proc itself. This supplementation is likely to be particularly useful for the study of computational complexity in general operational theories [20, 25, 119, 165, 173].

It should also be straightforward to formulate our framework using the language of category theory; category-theoretic tools might then provide guidance on which extensions of our framework are most easily formalized next, and might provide technical tools (e.g., for going beyond the finiteness assumptions that we have made). Making connections to the string diagrammatic representation of double-categories [208] may be a useful first step.

The process-theoretic formulation of a CI theory makes it easy to incorporate extra structure into the systems. Of particular interest would be to equip systems with the action of particular groups in order to be able to represent symmetries explicitly in our formalism. This is essential for an understanding of unspeakable information [30, 218] and for leveraging this to prove new no-go theorems and find new types of nonclassicality. Tools from Refs. [267, 268] provide a useful starting point for this project.

Additionally, it would be useful to complete the project begun in Section 6.10.2, namely, that of determining how various results in the framework of classical causal models [216] can be recast using the formalism of classical realist causal-inferential theories, and to explore to what extent the additional causal-inferential unscrambling that our framework provides may be beneficial to the field of causal inference. It will also be interesting to consider how the notion of an *operational* CI theory compares to the notion of a causal model with latent systems that can be quantum or GPT [106, 108, 142] and whether our framework offers some advantages relative to these.

At present, our framework describes only the reasoning of a single agent. It would be interesting to incorporate the reasoning of multiple agents. This project will require integrating insights such as pooling of states of knowledge [183, 288]. Relatedly, it would be interesting to consider what insight our framework can add to puzzles regarding the fact that agents can themselves be considered as physical systems. Such puzzles include the scenario of Wigner's friend [317] and variants thereof [105]. On a related note, it would be interesting to study how the available causal mechanisms in a CI theory determine the precise manner in which any agent, considered as a physical system, can gather information about its environment—and hence, what sort of theory of inference is most adaptive for it. One might expect that such considerations will constrain the interplay between the causal and inferential subtheories of any realist CI theory.

One could also seek to attack the problem of reconstructing quantum theory from novel axioms using our framework as an alternative to existing frameworks for reconstructing quantum theory [21, 45, 70, 71, 76, 90, 132, 136, 138, 146, 197, 269, 300, 306]. This would be particularly interesting if one could reconstruct quantum theory as a *realist* CI theory rather than a generic operational CI theory. Our framework may also may provide new insights into axioms that single out quantum correlations [42, 111, 128, 143, 187, 215] (including, e.g., the constraints articulated in Section 6.10.1).

One can naturally define postquantumness [19, 50, 140, 145, 174, 250, 254, 326] of correlations in our framework. For a given causal structure, any distribution that is P-S-realizable by an operational theory P-S, but that is not $\mathrm{P_Q}$-S-realizable, is said to be postquantum. Our framework may also allow for new ways of studying postquantumness; e.g., if one were to develop a notion of a *quantum* realist causal-inferential theory, then, for a given operational CI theory, one could seek to determine which experimental scenarios manifest postquantumness in the sense of failing to admit of a quantum realist representation or failing to admit of a Leibnizian quantum realist representation.

Epistemically restricted classical statistical theories, such as those described in Refs. [282, 282], if conceptualized as operational CI theories, are theories that admit of a Leibnizian classical realist representation. In this sense, if the world were governed by such a theory, there would be no problem to providing satisfactory realist explanations of observations, and one would have no need to consider any departure from the classical notion of realism F-S. Nonetheless, it might be interesting to try and cast such theories as examples of *nonclassical* realist CI theories themselves, that is, as defining a triple (XFunc, XSubStoch, XF-XS) that differs from the classical triple (Func, SubStoch, F-S). Ideally, this would be done such that the epistemic restriction emerges as a consequence of assumptions about the underlying reality, as opposed to being a supplementary assumption. Even though such theories are classical insofar as they also admit a Leibnizian classical realist representation, this project might nonetheless constitute a useful warm-up for the project of characterizing

QF-QS. On the one hand, by exploring other realist CI theories we will gain insight into how the causal and inferential subtheories constrain one another, and, on the other hand, there are many formal similarities between epistemically restricted statistical theories and quantum theory (indeed, they often constitute subtheories of quantum theory).

We also discussed in Section 6.6.3 how, if Conjecture 1 is false then there are no-go theorems for Leibnizian classical realist representations of operational quantum theory beyond the no-go theorems based on generalized noncontextuality. It is therefore important to settle the question of the status of Conjecture 1.

Although one should not demand (or even expect) Leibnizianity to be a strong enough principle to single out a unique nonclassical realist representation for operational quantum theory, results in Chapter 3 are suggestive that this may in fact be possible. (In particular, that chapter proves that there is a unique classical realist representation of any odd-dimensional stabilizer subtheory, namely, that given by Refs. [133, 282].)

We have presented a partial development of a graphical calculus for Boolean propositional logic. We leave for future work the problems of developing this into a complete graphical calculus, extending it to incorporate predicate logic, and generalizing it to nonclassical logics. Similarly, there are additional tools from Bayesian probability theory that would be useful to incorporate into SubStoch such as postselection and Bayesian inversion. Both of these projects are likely to help with the eventual development of QSubStoch.

Another research direction concerns the development of a resource theory [81] of nonclassicality. We have here argued that the distinction between classical and nonclassical is best understood as a distinction concerning the sort of realism required to provide an explanation of operational predictions. Within any proposal for a nonclassical realist CI theory XF-XS which subsumes the classical realist CI theory F-S, therefore, one can hope to formulate a resource theory of nonclassicality of processes. In this way, the research program described here could clarify the notion of nonclassicality inherent in 'common-cause boxes' (i.e., Bell scenarios), studied in Refs. [246, 258, 322], and the nonclassicality inherent in contextuality scenarios (i.e., scenarios that imply a noncontextuality no-go theorem, but not a Bell-like no-go theorem). label is CI

References

[1] E. T. Jaynes, in Complexity, Entropy, and the Physics of Information, edited by W. H. Zurek (Addison-Wesley, 1990) p. 381.

[2] Faacets. http://www.faacets.com.

[3] S. Aaronson and D. Gottesman. Improved simulation of stabilizer circuits. *Phys. Rev. A*, 70:052328, Nov. 2004.

[4] S. Abramsky, R. S. Barbosa, M. Karvonen, and S. Mansfield. A comonadic view of simulation and quantum resources. In *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–12. IEEE, 2019.

[5] S. Abramsky and B. Coecke. A categorical semantics of quantum protocols. In *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science, 2004.*, pages 415–425. IEEE, 2004.

[6] A. Acín, R. Augusiak, D. Cavalcanti, C. Hadley, J. K. Korbicz, M. Lewenstein, L. Masanes, and M. Piani. Unified Framework for Correlations in Terms of Local Quantum Observables. *Phys. Rev. Lett.*, 104:140404, Apr. 2010.

[7] A. Acín, N. Brunner, N. Gisin, S. Massar, S. Pironio, and V. Scarani. Device-Independent Security of Quantum Cryptography against Collective Attacks. *Phys. Rev. Lett.*, 98:230501, June 2007.

[8] A. Acín, N. Gisin, and L. Masanes. From bell's theorem to secure quantum key distribution. *Phys. Rev. Lett.*, 97:120405, Sept. 2006.

[9] A. Acin and L. Masanes. Certified randomness in quantum physics. *Nature*, 540(7632):213–219, Dec. 2016.

[10] D. Aharonov. A simple proof that toffoli and hadamard are quantum universal. *quant-ph/0301040*, 2003.

[11] S. W. Al-Safi and A. J. Short. Simulating all Nonsignaling Correlations via Classical or Quantum Theory with Negative Probabilities. *Phys. Rev. Lett.*, 111:170403, Oct. 2013.

[12] 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, July 2017.

[13] A. Ambainis. A new protocol and lower bounds for quantum coin flipping. *Journal of Computer and System Sciences*, 68(2):398–416, 2004. Special Issue on STOC 2001.

[14] A. Ambainis, M. Banik, A. Chaturvedi, D. Kravchenko, and A. Rai. Parity Oblivious *d*-Level Random Access Codes and Class of Noncontextuality Inequalities. *arXiv:1607.05490*, July 2016.

[15] H. Anwar, E. T. Campbell, and D. E. Browne. Qutrit magic state distillation. *New J. Phys.*, 14(6):063006, 2012.

[16] J.-D. Bancal, S. Pironio, A. Acín, Y.-C. Liang, V. Scarani, and N. Gisin. Quantum non-locality based on finite-speed causal influences leads to superluminal signalling. *Nat. Phys.*, 8(12):867–870, Oct. 2012.

[17] J.-D. Bancal, N. Sangouard, and P. Sekatski. Noise-Resistant Device-Independent Certification of Bell State Measurements. *Physical Review Letters*, 121(25):250506, Dec. 2018.

[18] D. Bankova, B. Coecke, M. Lewis, and D. Marsden. Graded entailment for compositional distributional semantics. *arXiv:1601.04908*, 2016.

[19] H. Barnum, M. A. Graydon, and A. Wilce. Composites and categories of euclidean jordan algebras. 2016.

[20] H. Barnum, C. M. Lee, and J. H. Selby. Oracles and query lower bounds in generalised probabilistic theories. *Foundations of physics*, 48(8):954–981, 2018.

[21] H. Barnum, M. P. Müller, and C. Ududec. Higher-order interference and single-system postulates characterizing quantum theory. *New J. Phys.*, 16(12):123029, 2014.

[22] J. Barrett. Nonsequential positive-operator-valued measurements on entangled mixed states do not always violate a bell inequality. *Phys. Rev. A*, 65:042302, Mar. 2002.

[23] J. Barrett. Information processing in generalized probabilistic theories. *Phys. Rev. A*, 75:032304, Mar 2007.

[24] J. Barrett, E. G. Cavalcanti, R. Lal, and O. J. E. Maroney. No ψ-Epistemic Model Can Fully Explain the Indistinguishability of Quantum States. *Phys. Rev. Lett.*, 112:250403, June 2014.

[25] J. Barrett, N. de Beaudrap, M. J. Hoban, and C. M. Lee. The computational landscape of general physical theories. *npj Quantum Information*, 5(1):1–10, 2019.

[26] J. Barrett and N. Gisin. How much measurement independence is needed to demonstrate nonlocality? *Phys. Rev. Lett.*, 106:100406, Mar. 2011.

[27] J. Barrett, L. Hardy, and A. Kent. No Signaling and Quantum Key Distribution. *Phys. Rev. Lett.*, 95:010503, June 2005.

[28] J. Barrett, N. Linden, S. Massar, S. Pironio, S. Popescu, and D. Roberts. Nonlocal correlations as an information-theoretic resource. *Physical Review A*, 71(2):022101, Feb. 2005.

[29] J. Barrett, R. Lorenz, and O. Oreshkov. Quantum Causal Models. June 2019.

[30] S. D. Bartlett, T. Rudolph, and R. W. Spekkens. Reference frames, superselection rules, and quantum information. *Reviews of Modern Physics*, 79(2):555, 2007.

[31] D. Beckman, D. Gottesman, M. A. Nielsen, and J. Preskill. Causal and localizable quantum operations. *Phys. Rev. A*, 64:052309, Oct. 2001.

[32] A. Belén Sainz, M. J. Hoban, P. Skrzypczyk, and L. Aolita. Bipartite post-quantum steering in generalised scenarios. July 2019.

[33] J. S. Bell. On the Einstein Podolsky Rosen paradox. *Physics*, 1:195–200, Nov. 1964.

[34] J. S. Bell. On the Problem of Hidden Variables in Quantum Mechanics. *Rev. Mod. Phys.*, 38:447–452, July 1966.

[35] J. S. Bell. La nouvelle cuisine. In *Quantum Mechanics, High Energy Physics And Accelerators: Selected Papers Of John S Bell (With Commentary)*, pages 910–928. World Scientific, 1995.

[36] J. S. Bell. *Speakable and Unspeakable in Quantum Mechanics: Collected Papers on Quantum Philosophy.* Cambridge University Press, June 2004.

[37] C. H. Bennett. Quantum cryptography using any two nonorthogonal states. *Phys. Rev. Lett.*, 68:3121–3124, May 1992.

[38] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters. Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels. *Phys. Rev. Lett.*, 70:1895–1899, Mar. 1993.

[39] J. Berkson. Limitations of the application of fourfold table analysis to hospital data. *Biometrics Bulletin*, 2(3):47–53, 1946.

[40] C. Branciard, D. Rosset, N. Gisin, and S. Pironio. Bilocal versus nonbilocal correlations in entanglement-swapping experiments. *Phys. Rev. A*, 85(3):032119, 2012.

[41] C. Branciard, D. Rosset, Y.-C. Liang, and N. Gisin. Measurement-Device-Independent Entanglement Witnesses for All Entangled Quantum States. *Physical Review Letters*, 110(6):060405, Feb. 2013.

[42] G. Brassard, H. Buhrman, N. Linden, A. A. Méthot, A. Tapp, and F. Unger. Limit on Nonlocality in Any World in Which Communication Complexity Is Not Trivial. *Phys. Rev. Lett.*, 96:250401, June 2006.

[43] S. Bravyi and A. Kitaev. Universal quantum computation with ideal clifford gates and noisy ancillas. *Phys. Rev. A*, 71(2):022316, 2005.

[44] N. Brunner, D. Cavalcanti, S. Pironio, V. Scarani, and S. Wehner. Bell nonlocality. *Rev. Mod. Phys.*, 86:419–478, Apr. 2014.

[45] A. Budiyono and D. Rohrlich. Quantum mechanics as classical statistical mechanics with an ontic extension and an epistemic restriction. *Nat. Commun.*, 8(1):1306, 2017.

[46] C. Budroni and A. Cabello. Bell inequalities from variable-elimination methods. *J. Phys. A*, 45(38):385304, 2012.

[47] H. Buhrman, R. Cleve, S. Massar, and R. de Wolf. Nonlocality and communication complexity. *Rev. Mod. Phys.*, 82:665–698, Mar. 2010.

[48] F. Buscemi. All entangled quantum states are nonlocal. *Phys. Rev. Lett.*, 108:200401, May 2012.

[49] F. Buscemi. Degradable channels, less noisy channels, and quantum statistical morphisms: An equivalence relation. *Probl Inf Transm*, 52:201–213, 2016.

[50] A. Cabello. Specker's fundamental principle of quantum mechanics. *arXiv:1212.1756 [physics, physics:quant-ph]*, Dec. 2012. arXiv: 1212.1756.

[51] E. T. Campbell, H. Anwar, and D. E. Browne. Magic-state distillation in all prime dimensions using quantum reed-muller codes. *Phys. Rev. X*, 2(4):041021, 2012.

[52] L. Catani and D. E. Browne. Spekkens' toy model in all dimensions and its relationship with stabiliser quantum mechanics. *New J. Phys.*, 19(7):073035, 2017.

[53] L. Catani and D. E. Browne. State-injection schemes of quantum computation in spekkens' toy theory. *Phys. Rev. A*, 98:052108, Nov. 2018.

[54] L. Catani and M. Leifer. A mathematical framework for operational fine tunings. *arXiv:2003.10050*, 2020.

[55] D. Cavalcanti and P. Skrzypczyk. Quantum steering: A review with focus on semidefinite programming. *Reports on Progress in Physics*, 80(2):024001, 2017.

[56] D. Cavalcanti, P. Skrzypczyk, G. H. Aguilar, R. V. Nery, P. H. S. Ribeiro, and S. P. Walborn. Detection of entanglement in asymmetric quantum networks and multipartite quantum steering. *Nature Communications*, 6:7941, Aug. 2015.

[57] D. Cavalcanti, P. Skrzypczyk, and I. Šupić. All entangled states can demonstrate nonclassical teleportation. *Phys. Rev. Lett.*, 119:110501, Sept. 2017.

[58] E. G. Cavalcanti, M. J. W. Hall, and H. M. Wiseman. Entanglement verification and steering when alice and bob cannot be trusted. *Phys. Rev. A*, 87:032306, Mar. 2013.

[59] C. M. Caves, C. A. Fuchs, and R. Schack. Quantum probabilities as bayesian probabilities. *Physical review A*, 65(2):022305, 2002.

[60] A. Chailloux, I. Kerenidis, S. Kundu, and J. Sikora. Optimal bounds for parity-oblivious random access codes. *New J. Phys.*, 18(4):045003, 2016.

[61] A. Chailloux, I. Kerenidis, and J. Sikora. Lower bounds for quantum oblivious transfer. *Quantum Info. Comput.*, 13(1-2):158–177, Jan. 2013.

[62] R. Chaves. Polynomial Bell Inequalities. *Phys. Rev. Lett.*, 116:010402, Jan. 2016.

[63] R. Chaves, G. Carvacho, I. Agresti, V. Di Giulio, L. Aolita, S. Giacomini, and F. Sciarrino. Quantum violation of an instrumental test. *Nature Physics*, 14(3):291–296, 2018.

[64] R. Chaves, D. Cavalcanti, and L. Aolita. Causal hierarchy of multipartite Bell nonlocality. *Quantum*, 1:23, Aug. 2017.

[65] R. Chaves, R. Kueng, J. B. Brask, and D. Gross. Unifying Framework for Relaxations of the Causal Assumptions in Bell's Theorem. *Phys. Rev. Lett.*, 114:140403, Apr. 2015.

[66] R. Chaves, L. Luft, and D. Gross. Causal structures from entropic information: geometry and novel scenarios. *New J. Phys.*, 16(4):043001, 2014.

[67] G. Chiribella, G. M. D'Ariano, and P. Perinotti. Quantum Circuit Architecture. *Phys. Rev. Lett.*, 101:060401, Aug. 2008.

[68] G. Chiribella, G. M. D'Ariano, and P. Perinotti. Theoretical framework for quantum networks. *Phys. Rev. A*, 80:022339, Aug. 2009.

[69] G. Chiribella, G. M. D'Ariano, and P. Perinotti. Probabilistic theories with purification. *Phys. Rev. A*, 81:062348, June 2010.

[70] G. Chiribella, G. M. D'Ariano, and P. Perinotti. Informational derivation of quantum theory. *Phys. Rev. A*, 84:012311, July 2011.

[71] G. Chiribella and R. Spekkens. Quantum theory: Informational foundations and foils. *Springer Netherlands*, 2015.

[72] K. Cho and B. Jacobs. Disintegration and bayesian inversion, both abstractly and concretely. *arxiv org/abs/1709.00322*, 2017.

[73] D. Chruściński and G. Sarbicki. Entanglement witnesses: Construction, analysis and classification. *Journal of Physics A: Mathematical and Theoretical*, 47(48):483001, 2014.

[74] B. S. Cirel'Son. Quantum generalizations of Bell's inequality. *Letters in Mathematical Physics*, 4(2):93–100, Mar. 1980.

[75] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt. Proposed Experiment to Test Local Hidden-Variable Theories. *Phys. Rev. Lett.*, 23:880–884, Oct. 1969.

[76] R. Clifton, J. Bub, and H. Halvorson. Characterizing quantum theory in terms of information-theoretic constraints. *Foundations of Physics*, 33(11):1561–1591, Nov. 2003.

[77] B. Coecke. The logic of entanglement. *0402014*, 2004.

[78] B. Coecke. Quantum picturalism. *Contemporary physics*, 51(1):59–83, 2010.

[79] B. Coecke. Terminality implies non-signalling. *arXiv:1405.3681*, 2014.

[80] B. Coecke, T. Fritz, and R. W. Spekkens. A mathematical theory of resources. *Information and Computation*, 250:59–86, 2016. Quantum Physics and Logic.

[81] B. Coecke, T. Fritz, and R. W. Spekkens. A mathematical theory of resources. *Info. Comp.*, 250:59–86, Oct. 2016.

[82] B. Coecke and A. Kissinger. *Picturing Quantum Processes: A First Course in Quantum Theory and Diagrammatic Reasoning*. Cambridge University Press, 2017.

[83] B. Coecke and M. Lewis. A compositional explanation of the 'pet fish' phenomenon. In *International Symposium on Quantum Interaction*, pages 179–192. Springer, 2015.

[84] B. Coecke, M. Sadrzadeh, and S. Clark. Mathematical Foundations for a Compositional Distributional Model of Meaning. Mar. 2010.

[85] B. Coecke and R. W. Spekkens. Picturing classical and quantum bayesian inference. *Synthese*, 186(3):651–696, 2012.

[86] R. Colbeck. Quantum And Relativistic Protocols For Secure Multi-Party Computation. *arXiv:0911.3814*, 2009.

[87] R. Colbeck and R. Renner. Free randomness can be amplified. *Nat. Phys.*, 8:450 EP –, 05 2012.

[88] F. Costa and S. Shrapnel. Quantum causal modelling. *New J. Phys.*, 18(6):063032, 2016.

[89] T. S. Cubitt, D. Leung, W. Matthews, and A. Winter. Improving Zero-Error Classical Communication with Entanglement. *Phys. Rev. Lett.*, 104:230503, June 2010.

[90] B. Dakic and C. Brukner. Quantum Theory and Beyond: Is Entanglement Special? Nov. 2009.

[91] P. Daley, K. Resch, and R. W. Spekkens. Experimentally adjudicating between different causal explanations of bell inequality violations using standard model selection techniques. Forthcoming.

[92] G. M. D'Ariano, G. Chiribella, and P. Perinotti. *Quantum theory from first principles: an informational approach.* Cambridge University Press, 2017.

[93] A. Datta, S. T. Flammia, and C. M. Caves. Entanglement and the power of one qubit. *Phys. Rev. A*, 72:042316, Oct. 2005.

[94] J. I. de Vicente. On nonlocality as a resource theory and nonlocality measures. *J. Phys. A*, 47(42):424017, Oct. 2014.

[95] N. Delfosse, P. Allard Guerin, J. Bian, and R. Raussendorf. Wigner function negativity and contextuality in quantum computation on rebits. *Phys. Rev. X*, 5:021003, Apr. 2015.

[96] B. d'Espagnat. The Quantum Theory and Reality. *Scientific American*, 241(5):158–181, 1979.

[97] C. Dhara, G. Prettico, and A. Acín. Maximal quantum randomness in Bell tests. *Phys. Rev. A*, 88:052116, Nov. 2013.

[98] W. Dür, G. Vidal, and J. I. Cirac. Three qubits can be entangled in two inequivalent ways. *Phys. Rev. A*, 62:062314, Nov. 2000.

[99] B. Eastin and E. Knill. Restrictions on transversal encoded quantum gate sets. *Phys. Rev. Lett.*, 102(11):110502, 2009.

[100] A. Einstein, B. Podolsky, and N. Rosen. Can Quantum-Mechanical Description of Physical Reality Be Considered Complete? *Physical Review*, 47(10):777–780, May 1935.

[101] J. Emerson. Quantum Chaos and Quantum-Classical Correspondence. Nov. 2002.

[102] J. Farinholt. An ideal characterization of the clifford operators. *J. Phys. A: Math. Theor.*, 47(30):305303, 2014.

[103] C. Ferrie and J. Emerson. Frame representations of quantum mechanics and the necessity of negativity in quasi-probability representations. *J. Phys. A*, 41(35):352001, July 2008.

[104] T. C. Fraser and E. Wolfe. Causal compatibility inequalities admitting quantum violations in the triangle structure. *Phys. Rev. A*, 98:022113, Aug. 2018.

[105] D. Frauchiger and R. Renner. Quantum theory cannot consistently describe the use of itself. *Nat. Commun.*, 9(1):3711, Sept. 2018.

[106] T. Fritz. Beyond bell's theorem: correlation scenarios. *New J. Phys.*, 14(10):103001, 2012.

[107] T. Fritz. Beyond bell's theorem: correlation scenarios. *New J. Physics*, 14(10):103001, 2012.

[108] T. Fritz. Beyond bell's theorem ii: Scenarios with arbitrary causal structure. *Communications in Mathematical Physics*, 341(2):391–434, 2016.

[109] T. Fritz. A synthetic approach to markov kernels, conditional independence and theorems on sufficient statistics. *Advances in Mathematics*, 370:107239, 2020.

[110] T. Fritz and P. Perrone. Bimonoidal structure of probability monads. *Electronic Notes in Theoretical Computer Science*, 341:121–149, 2018.

[111] T. Fritz, A. B. Sainz, R. Augusiak, J. B. Brask, R. Chaves, A. Leverrier, and A. Acín. Local orthogonality as a multipartite principle for quantum correlations. *Nat. Commun.*, 4(1):1–7, Aug. 2013. Number: 1 Publisher: Nature Publishing Group.

[112] C. A. Fuchs. Qbism, the perimeter of quantum bayesianism. *arXiv:1003.5209*, 2010.

[113] C. A. Fuchs and R. Schack. Quantum-bayesian coherence. *Reviews of modern physics*, 85(4):1693, 2013.

[114] R. Furber and B. Jacobs. Towards a categorical account of conditional probability. *arXiv:1306.0831*, 2013.

[115] R. Gallego and L. Aolita. Resource Theory of Steering. *Physical Review X*, 5(4):041008, Oct. 2015.

[116] R. Gallego and L. Aolita. Nonlocality free wirings and the distinguishability between Bell boxes. *Phys. Rev. A*, 95(3), Mar. 2017.

[117] R. Gallego, L. E. Würflinger, R. Chaves, A. Acín, and M. Navascués. Nonlocality in sequential correlation scenarios. *New J. Phys.*, 16(3):033037, 2014.

[118] E. F. Galvão. Discrete wigner functions and quantum computational speedup. *Phys. Rev. A*, 71:042302, Apr. 2005.

[119] A. J. Garner. Interferometric computation beyond quantum theory. *Foundations of Physics*, 48(8):886–909, 2018.

[120] J. Geller and M. Piani. Quantifying non-classical and beyond-quantum correlations in the unified operator formalism. *J. Phys. A*, 47(42):424030, Oct. 2014.

[121] A. Gheorghiu and C. Heunen. Ontological models for quantum theory as functors. *arXiv:1905.09055*, 2019.

[122] A. Gheorghiu, E. Kashefi, and P. Wallden. Robustness and device independence of verifiable blind quantum computing. *New Journal of Physics*, 17(8):083040, Aug. 2015.

[123] K. S. Gibbons, M. J. Hoffman, and W. K. Wootters. Discrete phase space based on finite fields. *Phys. Rev. A*, 70(6):062101, 2004.

[124] S. Gogioso. Fantastic quantum theories and where to find them. *arXiv:1703.10576*, 2017.

[125] S. Gogioso and F. Genovese. Infinite-dimensional categorical quantum mechanics. *arXiv:1605.04305*, 2016.

[126] S. Gogioso and F. Genovese. Quantum field theory in categorical quantum mechanics. *arXiv:1805.12087*, 2018.

[127] S. Gogioso and C. M. Scandolo. Categorical probabilistic theories. *arXiv:1701.08075*, 2017.

[128] T. Gonda, R. Kunjwal, D. Schmid, E. Wolfe, and A. B. Sainz. Almost Quantum Correlations are Inconsistent with Specker's Principle. *Quantum*, 2:87, Aug. 2018.

[129] T. Gonda and R. W. Spekkens. Monotones in general resource theories, 2019.

[130] D. Gottesman. Stabilizer codes and quantum error correction. 1997.

[131] D. Gottesman. The heisenberg representation of quantum computers. 1998.

[132] P. Goyal. Information-geometric reconstruction of quantum theory. *Phys. Rev. A*, 78(5):052120, 2008.

[133] D. Gross. Hudson's theorem for finite-dimensional quantum systems. *J. Math. Phys.*, 47(12):122107, 2006.

[134] G. Gutoski. Properties of local quantum operations with shared entanglement. *arXiv:0805.2209*, 2008.

[135] M. J. W. Hall. Local deterministic model of singlet state correlations based on relaxing measurement independence. *Phys. Rev. Lett.*, 105:250404, Dec. 2010.

[136] L. Hardy. Quantum Theory From Five Reasonable Axioms. *eprint arXiv:quant-ph/0101012*, Jan. 2001.

[137] L. Hardy. A formalism-local framework for general probabilistic theories including quantum theory. *arXiv:1005.5164*, 2010.

[138] L. Hardy. Reformulating and reconstructing quantum theory. *arXiv:1104.2066*, 2011.

[139] N. Harrigan and R. W. Spekkens. Einstein, Incompleteness, and the Epistemic View of quantum states. *Found. Phys.*, 40(2):125–157, 2010.

[140] J. Hefford and S. Gogioso. Hyper-decoherence in density hypercubes. *arXiv:2003.08318*, 2020.

[141] C. W. Helstrom. Quantum detection and estimation theory. *J. Stat. Phys.*, 1(2):231–252, 1969.

[142] J. Henson, R. Lal, and M. F. Pusey. Theory-independent limits on correlations from generalized bayesian networks. *New J. Phys.*, 16(11):113043, 2014.

[143] J. Henson and A. B. Sainz. Macroscopic noncontextuality as a principle for almost-quantum correlations. *Phys. Rev. A*, 91(4):042114, Apr. 2015.

[144] M. J. Hoban, E. T. Campbell, K. Loukopoulos, and D. E. Browne. Non-adaptive measurement-based quantum computation and multi-party Bell inequalities. *New J. Phys.*, 13(2):023014, 2011.

[145] M. J. Hoban and A. B. Sainz. A channel-based framework for steering, non-locality and beyond. *New Journal of Physics*, 20(5):053048, May 2018.

[146] P. A. Höhn and C. S. P. Wever. Quantum theory from questions. *Phys. Rev. A*, 95(1):012102, 2017.

[147] G. Hooft. Free will in the theory of everything. *arXiv:1709.02874*, 2017.

[148] C. A. Hooker. *The Logico-Algebraic Approach to Quantum Mechanics: Vol. II: Contemporary Consolidation.* Springer Science & Business Media, 1979.

[149] C. A. Hooker. *The Logico-Algebraic Approach to Quantum Mechanics: Volume I: Historical Evolution*, volume 5. Springer Science & Business Media, 2012.

[150] K. Horodecki, A. Grudka, P. Joshi, W. Kłobus, and J. Łodyga. Axiomatic approach to contextuality and nonlocality. *Physical Review A*, 92(3):032104, Sept. 2015.

[151] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki. Quantum entanglement. *Rev. Mod. Phys.*, 81(2):865, 2009.

[152] D. Horsman, C. Heunen, M. F. Pusey, J. Barrett, and R. W. Spekkens. Can a quantum state over time resemble a quantum state at a single time? *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 473(2205):20170395, 2017.

[153] E. Hostens, J. Dehaene, and B. De Moor. Stabilizer states and clifford operations for systems of arbitrary dimensions and modular arithmetic. *Phys. Rev. A*, 71:042315, Apr. 2005.

[154] M. Howard, J. Wallman, V. Veitch, and J. Emerson. Contextuality supplies the 'magic' for quantum computation. *Nature*, 510(7505):351–355, June 2014.

[155] B. Jacobs, A. Kissinger, and F. Zanasi. Causal inference by string diagram surgery. In *International Conference on Foundations of Software Science and Computation Structures*, pages 313–329. Springer, 2019.

[156] A. Joyal and R. Street. The geometry of tensor calculus, i. *Advances in mathematics*, 88(1):55–112, 1991.

[157] J. Kaniewski and S. Wehner. Device-independent two-party cryptography secure against sequential attacks. *New J. Phys.*, 18(5):055004, 2016.

[158] M. Karvonen. Categories of empirical models. *arXiv:1804.01514*, 2018.

[159] G. M. Kelly. Many-variable functorial calculus. i. In *Coherence in categories*, pages 66–105. Springer, 1972.

[160] A. Kissinger, M. Hoban, and B. Coecke. Equivalence of relativistic causal structure and process terminality. *arXiv:1708.04118*, 2017.

[161] E. Knill and R. Laflamme. Power of one bit of quantum information. *Phys. Rev. Lett.*, 81:5672–5675, Dec. 1998.

[162] S. Kochen and E. Specker. The problem of hidden variables in quantum mechanics. *J. Math. & Mech.*, 17:59–87, 1967. Also available from the Indiana Univ. Math. J.

[163] A. Krishna, R. W. Spekkens, and E. Wolfe. Deriving robust noncontextuality inequalities from algebraic proofs of the Kochen-Specker theorem: the Peres-Mermin square. *arXiv:1704.01153*, Apr. 2017.

[164] Krishna, Anirudh. Experimentally Testable Noncontextuality Inequalities Via Fourier-Motzkin Elimination. Master's thesis, University of Waterloo, 2015.

[165] M. Krumm and M. P. Müller. Quantum computation is the unique reversible circuit model for which bits are balls. *npj Quantum Information*, 5(1):1–8, 2019.

[166] R. Kunjwal. Contextuality beyond the Kochen-Specker theorem. *arXiv:1612.07250*, 2016.

[167] R. Kunjwal. Beyond the Cabello-Severini-Winter framework: Making sense of contextuality without sharpness of measurements. *Quantum*, 3:184, Sept. 2019.

[168] R. Kunjwal, M. Lostaglio, and M. F. Pusey. Anomalous weak values and contextuality: Robustness, tightness, and imaginary parts. *Phys. Rev. A*, 100:042116, Oct. 2019.

[169] R. Kunjwal and R. W. Spekkens. From the Kochen-Specker Theorem to Noncontextuality Inequalities without Assuming Determinism. *Phys. Rev. Lett.*, 115:110403, Sept. 2015.

[170] R. Kunjwal and R. W. Spekkens. From statistical proofs of the kochen-specker theorem to noise-robust noncontextuality inequalities. *Phys. Rev. A*, 97:052110, May 2018.

[171] R. Kunjwal and R. W. Spekkens. From statistical proofs of the kochen-specker theorem to noise-robust noncontextuality inequalities. *Physical Review A*, 97(5):052110, 2018.

[172] J. Ladyman and D. Ross. *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press, 2007.

[173] C. M. Lee and J. Barrett. Computation in generalised probabilisitic theories. *New J. Phys.*, 17(8):083001, 2015.

[174] C. M. Lee and J. H. Selby. A no-go theorem for theories that decohere to quantum mechanics. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 474(2214):20170732, 2018.

[175] C. M. Lee and R. W. Spekkens. Causal inference via algebraic geometry: feasibility tests for functional causal structures with two binary observed variables. *arXiv:1506.03880*, June 2015.

[176] M. S. Leifer. Quantum dynamics as an analog of conditional probability. *Phys. Rev. A*, 74(4):042310, 2006.

[177] M. S. Leifer. Conditional density operators and the subjectivity of quantum operations. In *AIP Conference Proceedings*, volume 889, pages 172–186. American Institute of Physics, 2007.

[178] M. S. Leifer. Is the quantum state real? an extended review of psi-ontology theorems. *arXiv:1409.1570*, 2014.

[179] M. S. Leifer and O. J. E. Maroney. Maximally Epistemic Interpretations of the Quantum State and Contextuality. *Phys. Rev. Lett.*, 110:120401, Mar. 2013.

[180] M. S. Leifer and D. Poulin. Quantum graphical models and belief propagation. *Annals of Physics*, 323(8):1899–1946, 2008.

[181] M. S. Leifer and R. W. Spekkens. Pre- and Post-Selection Paradoxes and Contextuality in Quantum Mechanics. *Phys. Rev. Lett.*, 95:200405, Nov. 2005.

[182] 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, Nov. 2013.

[183] M. S. Leifer and R. W. Spekkens. A bayesian approach to compatibility, improvement, and pooling of quantum states. *Journal of Physics A: Mathematical and Theoretical*, 47(27):275301, June 2014.

[184] Y.-C. Liang, R. W. Spekkens, and H. M. Wiseman. Specker's parable of the overprotective seer: A road to contextuality, nonlocality and complementarity. *Physics Reports*, 506(1):1–39, 2011.

[185] P. Lillystone and J. Emerson. A contextual ψ-epistemic model of the n-qubit stabilizer formalism. 2019.

[186] P. Lillystone, J. J. Wallman, and J. Emerson. Contextuality and the single-qubit stabilizer subtheory. *Phys. Rev. Lett.*, 122:140405, Apr. 2019.

[187] N. Linden, S. Popescu, A. J. Short, and A. Winter. Quantum Nonlocality and Beyond: Limits from Nonlocal Computation. *Physical Review Letters*, 99(18):180502, Oct. 2007. Publisher: American Physical Society.

[188] P. Lipka-Bartosik and P. Skrzypczyk. The operational advantages provided by non-classical teleportation. Aug. 2019.

[189] H.-K. Lo and H. F. Chau. Is quantum bit commitment really possible? *Phys. Rev. Lett.*, 78:3410–3413, Apr. 1997.

[190] R. Lorenz and J. Barrett. Causal and compositional structure of unitary transformations. Jan. 2020.

[191] M. Lostaglio. Certifying quantum signatures in thermodynamics and metrology via contextuality of quantum linear response. *Phys. Rev. Lett.*, 125:230603, Dec 2020.

[192] M. Lostaglio and G. Senno. Contextual advantage for state-dependent cloning. *arXiv:1905.08291*, 2019.

[193] S. Mansfield and E. Kashefi. Quantum advantage from sequential-transformation contextuality. *Physical review letters*, 121(23):230401, 2018.

[194] A. Mari and J. Eisert. Positive wigner functions render classical simulation of quantum computation efficient. *Phys. Rev. Lett.*, 109:230503, Dec. 2012.

[195] O. J. E. Maroney. How statistical are quantum states? *arXiv:1207.6906*, July 2012.

[196] I. Marvian. Inaccessible information in probabilistic models of quantum systems, non-contextuality inequalities and noise thresholds for contextuality. *arXiv:2003.05984*, 2020.

[197] L. Masanes and M. P. Mueller. A derivation of quantum theory from physical requirements. *New J. Phys.*, 13(6):063001, June 2011.

[198] T. Maudlin. Bell's Inequality, Information Transmission, and Prism Models. In *Philosophy of Science Association*, number 1, pages 404–417, 1992.

[199] T. Maudlin. *Quantum Non-Locality and Relativity : Metaphysical Intimations of Modern Physics*. Blackwell Publishers, 2002.

[200] D. Mayers. Unconditionally secure quantum bit commitment is impossible. *Phys. Rev. Lett.*, 78:3414–3417, Apr. 1997.

[201] D. Mayers and A. Yao. Self testing quantum apparatus. *Quantum Information & Computation*, 4(4):273–286, 2004.

[202] M. D. Mazurek, M. F. Pusey, R. Kunjwal, K. J. Resch, and R. W. Spekkens. An experimental test of noncontextuality without unphysical idealizations. *Nat. Comm.*, 7(1), June 2016.

[203] M. D. Mazurek, M. F. Pusey, K. J. Resch, and R. W. Spekkens. Experimentally bounding deviations from quantum theory in the landscape of generalized probabilistic theories. *arXiv:1710.05948*, 2017.

[204] P.-A. Melliès. Functorial boxes in string diagrams. In *International Workshop on Computer Science Logic*, pages 1–30. Springer, 2006.

[205] N. D. Mermin. Hidden variables and the two theorems of John Bell. *Rev. Mod. Phys.*, 65:803–815, July 1993.

[206] A. Montanaro and R. de Wolf. A Survey of Quantum Property Testing. *Theory of Computing*, pages 1–81, July 2016.

[207] A. Montina. Exponential complexity and ontological theories of quantum mechanics. *Phys. Rev. A*, 77(2):022104, 2008.

[208] D. J. Myers. String diagrams for double categories and equipments. *arXiv:1612.02762*, 2016.

[209] M. Navascués, Y. Guryanova, M. J. Hoban, and A. Acín. Almost quantum correlations. *Nature communications*, 6(1):1–7, 2015.

[210] M. Navascues, S. Pironio, and A. Acin. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New J. Phys.*, 10(7):073013, 2008.

[211] G. Nebe, E. M. Rains, and N. J. A. Sloane. *Self-dual codes and invariant theory*, volume 17. Springer, 2006.

[212] M. A. Nielsen and I. L. Chuang. Quantum computation and quantum information. *Phys. Today*, 54(2):60, 2001.

[213] T. Norsen. Bell locality and the nonlocal character of nature. *quant-ph/0601205*, 2006.

[214] E. Patterson, D. I. Spivak, and D. Vagner. Wiring diagrams as normal forms for computing in symmetric monoidal categories. *arXiv preprint arXiv:2101.12046*, 2021.

[215] M. Pawłowski, T. Paterek, D. Kaszlikowski, V. Scarani, A. Winter, and M. Żukowski. Information causality as a physical principle. *Nature*, 461(7267):1101–1104, Oct. 2009. Number: 7267 Publisher: Nature Publishing Group.

[216] J. Pearl. *Causality*. Cambridge university press, 2009.

[217] R. Penrose. Applications of negative dimensional tensors. *Combinatorial mathematics and its applications*, 1:221–244, 1971.

[218] A. Peres and P. Scudo. Unspeakable quantum information. eprint. *quantph/0201017*, 2002.

[219] P. Perrone. Notes on category theory with examples from basic mathematics. *arXiv:1912.10642*, 2019.

[220] M. Piani. Channel steering. *J. Opt. Soc. Am. B*, 32(4):A1–A7, Apr. 2015.

[221] M. Piani and J. Watrous. Necessary and Sufficient Quantum Information Characterization of Einstein-Podolsky-Rosen Steering. *Physical Review Letters*, 114(6):060404, Feb. 2015.

[222] S. Pironio, A. Acín, S. Massar, A. B. de la Giroday, D. N. Matsukevich, P. Maunz, S. Olmschenk, D. Hayes, L. Luo, T. A. Manning, and C. Monroe. Random numbers certified by Bell's theorem. *Nature*, 464:1021 EP –, 04 2010.

[223] S. Pironio, V. Scarani, and T. Vidick. Focus on device independent quantum information. *New Journal of Physics*, 18(10):100202, Oct. 2016.

[224] S. Popescu and D. Rohrlich. Quantum nonlocality as an axiom. *Foundations of Physics*, 24(3):379–385, Mar. 1994.

[225] D. Poulin. Stabilizer formalism for operator quantum error correction. *Phys. Rev. Lett.*, 95:230504, Dec. 2005.

[226] M. F. Pusey. Stabilizer notation for spekkens' toy theory. *Foundations of Physics*, 42(5):688–708, 2012.

[227] M. F. Pusey. Negativity and steering: A stronger Peres conjecture. *Phys. Rev. A*, 88(3):032313, Sept. 2013.

[228] M. F. Pusey. Anomalous Weak Values Are Proofs of Contextuality. *Phys. Rev. Lett.*, 113:200401, Nov. 2014.

[229] M. F. Pusey. The robust noncontextuality inequalities in the simplest scenario. *arXiv:1506.04178*, June 2015.

[230] M. F. Pusey, J. Barrett, and T. Rudolph. On the reality of the quantum state. *Nature Physics*, 8(6):475–478, 2012.

[231] M. F. Pusey and M. S. Leifer. Logical pre- and post-selection paradoxes are proofs of contextuality. volume 195, pages 295–306. EPTCS, 2015.

[232] H. Putnam. *Is Logic Empirical?*, pages 216–241. Springer Netherlands, Dordrecht, 1969.

[233] G. Pütz, D. Rosset, T. J. Barnea, Y.-C. Liang, and N. Gisin. Arbitrarily Small Amount of Measurement Independence Is Sufficient to Manifest Quantum Nonlocality. *Physical Review Letters*, 113(19):190402, Nov. 2014.

[234] R. Raussendorf. Contextuality in measurement-based quantum computation. *Phys. Rev. A*, 88:022322, Aug. 2013.

[235] R. Raussendorf, D. E. Browne, N. Delfosse, C. Okay, and J. Bermejo-Vega. Contextuality and wigner-function negativity in qubit quantum computation. *Phys. Rev. A*, 95:052334, May 2017.

[236] M.-O. Renou, E. Baumer, S. Boreiri, N. Brunner, N. Gisin, and S. Beigi. Genuine quantum nonlocality in the triangle network. *Phys. Rev. Lett.*, 123:140401, Sept. 2019.

[237] M. O. Renou, J. Kaniewski, and N. Brunner. Self-testing entangled measurements in quantum networks. *Physical Review Letters*, 121(25):250507, Dec. 2018.

[238] F. Riesz. Démonstration nouvelle d'un théorème concernant les opérations fonctionnelles linéaires. In *Annales scientifiques de l'École Normale Supérieure*, volume 31, pages 9–14, 1914.

[239] M. Ringbauer, B. Duffus, C. Branciard, E. G. Cavalcanti, A. G. White, and A. Fedrizzi. Measurements on the reality of the wavefunction. *Nat Phys*, 11(3):249–254, Mar. 2015. Article.

[240] D. Riste, S. Poletto, M.-Z. Huang, A. Bruno, V. Vesterinen, O.-P. Saira, and L. DiCarlo. Detecting bit-flip errors in a logical qubit using stabilizer measurements. *Nat. Commun.*, 6(1):1–6, 2015.

[241] C. D. B. A. M. T. C. Roberto D. Baldijao, Rafael Wagner. Noncontextuality as a meaning of classicality in quantum darwinism. 2021.

[242] D. Rosset, C. Branciard, N. Gisin, and Y.-C. Liang. Entangled states cannot be classically simulated in generalized Bell experiments with quantum inputs. *New Journal of Physics*, 15(5):053025, May 2013.

[243] D. Rosset, F. Buscemi, and Y.-C. Liang. Resource theory of quantum memories and their faithful verification with minimal assumptions. *Phys. Rev. X*, 8:021033, May 2018.

[244] D. Rosset, R. Ferretti-Schöbitz, J.-D. Bancal, N. Gisin, and Y.-C. Liang. Imperfect measurement settings: Implications for quantum state tomography and entanglement witnesses. *Physical Review A*, 86(6):062325, Dec. 2012.

[245] D. Rosset, A. Martin, E. Verbanis, C. C. W. Lim, and R. Thew. Practical measurement-device-independent entanglement quantification. *Physical Review A*, 98(5):052332, Nov. 2018.

[246] D. Rosset, D. Schmid, and F. Buscemi. Characterizing nonclassicality of arbitrary distributed devices. *arXiv:1911.12462*, 2019.

[247] D. Rosset, D. Schmid, and F. Buscemi. Type-independent characterization of spacelike separated resources. *Phys. Rev. Lett.*, 125:210402, Nov. 2020.

[248] D. Saha and A. Chaturvedi. Preparation contextuality as an essential feature underlying quantum communication advantage. *Phys. Rev. A*, 100:022108, Aug. 2019.

[249] D. Saha, P. Horodecki, and M. Pawłowski. State independent contextuality advances one-way communication. *New J. Phys.*, 21(9):093057, Sept. 2019.

[250] A. B. Sainz, N. Brunner, D. Cavalcanti, P. Skrzypczyk, and T. Vértesi. Postquantum steering. *Physical review letters*, 115(19):190403, 2015.

[251] A. B. Sainz, Y. Guryanova, A. Acín, and M. Navascués. Almost-quantum correlations violate the no-restriction hypothesis. *Phys. Rev. Lett.*, 120(20):200402, 2018.

[252] V. Scarani. The Device-Independent Outlook on Quantum Physics. *Acta Physica Slovaca*, 62(4):347, 2012.

[253] V. Scarani, N. Gisin, N. Brunner, L. Masanes, S. Pino, and A. Acín. Secrecy extraction from no-signaling correlations. *Phys. Rev. A*, 74:042339, Oct. 2006.

[254] D. Schmid, H. Du, M. Mudassar, G. C.-d. Wit, D. Rosset, and M. J. Hoban. Postquantum common-cause channels: the resource theory of local operations and shared entanglement. *arXiv:2004.06133*, 2020.

[255] D. Schmid, H. Du, J. H. Selby, and M. F. Pusey. The stabilizer subtheory has a unique noncontextual model. *arXiv:2101.06263*, 2021.

[256] D. Schmid, T. C. Fraser, R. Kunjwal, A. B. Sainz, E. Wolfe, and R. W. Spekkens. Why standard entanglement theory is inappropriate for the study of Bell scenarios. Apr. 2020.

[257] D. Schmid, K. Ried, and R. W. Spekkens. Why initial system-environment correlations do not imply the failure of complete positivity: A causal perspective. *Phys. Rev. A*, 100:022112, Aug. 2019.

[258] D. Schmid, D. Rosset, and F. Buscemi. The type-independent resource theory of local operations and shared randomness. *Quantum*, 4:262, Apr. 2020.

[259] D. Schmid, J. Selby, and R. W. Spekkens. Generalized probabilistic theories in the framework of causal-inferential theories. Forthcoming.

[260] D. Schmid, J. H. Selby, M. F. Pusey, and R. W. Spekkens. A structure theorem for generalized-noncontextual ontological models. *arXiv:2005.07161*, 2020.

[261] D. Schmid, J. H. Selby, and R. W. Spekkens. Unscrambling the omelette of causation and inference: The framework of causal-inferential theories. 2020.

[262] D. Schmid, J. H. Selby, E. Wolfe, R. Kunjwal, and R. W. Spekkens. Characterization of noncontextuality in the framework of generalized probabilistic theories. *PRX Quantum*, 2:010331, Feb. 2021.

[263] D. Schmid and R. W. Spekkens. Contextual advantage for state discrimination. *Phys. Rev. X*, 8(1):011015, 2018.

[264] D. Schmid, R. W. Spekkens, and E. Wolfe. All the noncontextuality inequalities for arbitrary prepare-and-measure experiments with respect to any fixed set of operational equivalences. *Phys. Rev. A*, 97:062103, June 2018.

[265] E. Schrodinger. Discussion of probability relations between separated systems. *Mathematical Proceedings of the Cambridge Philosophical Society*, 31(4):555–563, 1935.

[266] P. Sekatski, J.-D. Bancal, S. Wagner, and N. Sangouard. Certifying the Building Blocks of Quantum Computers from Bell's Theorem. *Physical Review Letters*, 121(18):180505, Nov. 2018.

[267] J. Selby and B. Coecke. Leaks: quantum, classical, intermediate and more. *Entropy*, 19(4):174, 2017.

[268] J. H. Selby and C. M. Lee. Compositional resource theories of coherence. *arXiv:1911.04513*, 2019.

[269] J. H. Selby, C. M. Scandolo, and B. Coecke. Reconstructing quantum theory from diagrammatic postulates. *arXiv:1802.00367*, 2018.

[270] J. H. Selby, D. Schmid, E. Wolfe, A. B. Sainz, R. Kunjwal, and R. W. Spekkens. Contextuality without incompatibility. 2021.

[271] P. Selinger. A survey of graphical languages for monoidal categories. In *New structures for physics*, pages 289–355. Springer, 2010.

[272] F. Shahandeh. Contextuality of general probabilistic theories and the power of a single resource. 2019.

[273] F. Shahandeh, M. J. W. Hall, and T. C. Ralph. Measurement-Device-Independent Approach to Entanglement Measures. *Physical Review Letters*, 118(15):150505, Apr. 2017.

[274] A. Shimony. Bell's Theorem. In *The Stanford Encyclopedia of Philosophy*, 2017.

[275] E. H. Simpson. The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 13(2):238–241, 1951.

[276] P. Skrzypczyk and N. Linden. Robustness of measurement, discrimination games, and accessible information. *Phys. Rev. Lett.*, 122:140403, Apr. 2019.

[277] P. Skrzypczyk, M. Navascués, and D. Cavalcanti. Quantifying Einstein-Podolsky-Rosen Steering. *Physical Review Letters*, 112(18):180404, May 2014.

[278] P. Skrzypczyk, I. Supic, and D. Cavalcanti. All sets of incompatible measurements give an advantage in quantum state discrimination. *Phys. Rev. Lett.*, 122:130403, Apr. 2019.

[279] R. W. Spekkens. Distinguishing inference from influence: reassessing how to formalize various causal notions appearing in the bell literature. Forthcoming.

[280] R. W. Spekkens. Why i am not a psi-ontologist. *PIRSA*.

[281] R. W. Spekkens. Contextuality for preparations, transformations, and unsharp measurements. *Phys. Rev. A*, 71:052108, May 2005.

[282] R. W. Spekkens. Evidence for the epistemic view of quantum states: A toy theory. *Phys. Rev. A*, 75:032110, Mar. 2007.

[283] R. W. Spekkens. Negativity and Contextuality are Equivalent Notions of Nonclassicality. *Phys. Rev. Lett.*, 101:020401, July 2008.

[284] R. W. Spekkens. *Quasi-Quantization: Classical Statistical Theories with an Epistemic Restriction*, pages 83–135. Springer Netherlands, Dordrecht, 2016.

[285] R. W. Spekkens. The ontological identity of empirical indiscernibles: Leibniz's methodological principle and its significance in the work of Einstein. *arXiv:1909.04628*, Aug. 2019.

[286] 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, Jan. 2009.

[287] R. W. Spekkens and T. Rudolph. Degrees of concealment and bindingness in quantum bit commitment protocols. *Phys. Rev. A*, 65:012310, Dec. 2001.

[288] R. W. Spekkens and H. M. Wiseman. Pooling quantum states obtained by indirect measurements. *Phys. Rev. A*, 75:042104, Apr. 2007.

[289] J. Sperling and W. Vogel. The Schmidt number as a universal entanglement measure. *Physica Scripta*, 83(4):045002, 2011.

[290] P. Spirtes, C. N. Glymour, R. Scheines, and D. Heckerman. *Causation, prediction, and search*. MIT press, 2000.

[291] I. Šupić and J. Bowles. Apr. 2019.

[292] I. Šupić and M. J. Hoban. Self-testing through EPR-steering. *New Journal of Physics*, 18(7):075006, July 2016.

[293] I. Šupić, P. Skrzypczyk, and D. Cavalcanti. Measurement-device-independent entanglement and randomness estimation in quantum networks. *Physical Review A*, 95(4):042340, Apr. 2017.

[294] I. Šupić, P. Skrzypczyk, and D. Cavalcanti. Quantifying non-classical teleportation. Apr. 2018.

[295] R. Takagi and B. Regula. General resource theories in quantum mechanics and beyond: Operational characterization via discrimination tasks. *Phys. Rev. X*, 9:031053, Sept. 2019.

[296] R. Takagi, B. Regula, K. Bu, Z.-W. Liu, and G. Adesso. Operational advantage of quantum resources in subchannel discrimination. *Phys. Rev. Lett.*, 122:140402, Apr. 2019.

[297] A. Tavakoli, J. Kaniewski, T. Vertesi, D. Rosset, and N. Brunner. Self-testing quantum states and measurements in the prepare-and-measure scenario. *Physical Review A*, 98(6):062307, Dec. 2018.

[298] B. M. Terhal. Quantum error correction for quantum memories. *Rev. Mod. Phys.*, 87:307–346, Apr. 2015.

[299] B. F. Toner and D. Bacon. Communication Cost of Simulating Bell Correlations. *Phys. Rev. Lett.*, 91:187904, Oct. 2003.

[300] S. Tull. A categorical reconstruction of quantum theory. *arXiv:1804.02265*, 2018.

[301] R. Uola, A. C. S. Costa, H. C. Nguyen, and O. Gühne. Quantum Steering. Mar. 2019.

[302] R. Uola, T. Kraft, and A. A. Abbott. Quantification of quantum dynamics with input-output games, 2019.

[303] R. Uola, T. Kraft, J. Shang, X.-D. Yu, and O. Gühne. Quantifying quantum resources with conic programming. *Phys. Rev. Lett.*, 122:130404, Apr. 2019.

[304] W. van Dam and M. Howard. Noise thresholds for higher-dimensional systems using the discrete wigner function. *Phys. Rev. A*, 83:032310, Mar. 2011.

[305] J. van de Wetering. Quantum theory is a quasi-stochastic process theory. 2017.

[306] J. van de Wetering. An effect-theoretic reconstruction of quantum theory. *arXiv:1801.05798*, 2018.

[307] T. Van Himbeeck, J. Bohr Brask, S. Pironio, R. Ramanathan, A. B. Sainz, and E. Wolfe. Quantum violations in the Instrumental scenario and their relations to the Bell scenario. *Quantum*, 3:186, Sept. 2019.

[308] U. Vazirani and T. Vidick. Fully Device-Independent Quantum Key Distribution. *Phys. Rev. Lett.*, 113:140501, Sept. 2014.

[309] V. Veitch, C. Ferrie, D. Gross, and J. Emerson. Negative quasi-probability as a resource for quantum computation. *New J. Phys.*, 14(11):113011, 2012.

[310] V. Veitch, S. H. Mousavian, D. Gottesman, and J. Emerson. The resource theory of stabilizer quantum computation. *New J. Phys.*, 16(1):013009, 2014.

[311] G. Vidal. Entanglement monotones. *J. Mod. Optic.*, 47(2-3):355–376, Feb. 2000.

[312] G. Vidal. Efficient classical simulation of slightly entangled quantum computations. *Phys. Rev. Lett.*, 91:147902, Oct. 2003.

[313] D. S. Wang, A. G. Fowler, and L. C. L. Hollenberg. Surface code quantum computing with error rates over 1 *Phys. Rev. A*, 83:020302, Feb. 2011.

[314] J. Watrous. *The theory of quantum information.* Cambridge University Press, 2018.

[315] R. F. Werner. Quantum states with einstein-podolsky-rosen correlations admitting a hidden-variable model. *Phys. Rev. A*, 40:4277–4281, Oct. 1989.

[316] R. F. Werner. Comment on what Bell did. *J. Phys. A*, 47(42):424011, 2014.

[317] E. P. Wigner. *Remarks on the Mind-Body Question*, pages 247–260. Springer Berlin Heidelberg, Berlin, Heidelberg, 1995.

[318] A. Wilce. A shortcut from categorical quantum theory to convex operational theories. *arXiv:1803.00707*, 2018.

[319] H. M. Wiseman. The two Bell's theorems of John Bell. *J. Phys. A*, 47(42):424001, 2014.

[320] H. M. Wiseman, S. J. Jones, and A. C. Doherty. Steering, entanglement, nonlocality, and the einstein-podolsky-rosen paradox. *Phys. Rev. Lett.*, 98:140402, Apr. 2007.

[321] E. Wolfe, A. Pozas-Kerstjens, M. Grinberg, D. Rosset, A. Acín, and M. Navascues. Quantum Inflation: A General Approach to Quantum Causal Compatibility. *arXiv:1909.10519*, 2019.

[322] E. Wolfe, D. Schmid, A. B. Sainz, R. Kunjwal, and R. W. Spekkens. Quantifying Bell: the Resource Theory of Nonclassicality of Common-Cause Boxes. *Quantum*, 4:280, June 2020.

[323] E. Wolfe, R. W. Spekkens, and T. Fritz. The Inflation Technique for Causal Inference with Latent Variables. *arXiv:1609.00672*, Sept. 2016.

[324] 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(3):033002, 2015.

[325] S. A. Yadavalli and R. Kunjwal. Contextuality in entanglement-assisted one-shot classical communication. *arXiv:2006.00469*, 2020.

[326] K. Życzkowski. Quartic quantum theory: an extension of the standard quantum mechanics. *Journal of Physics A: Mathematical and Theoretical*, 41(35):355302, 2008.

Appendices

A Appendices for Chapter 2

A.1 Proof of noncontextuality no-go theorem for MESD

Herein we provide an alternative proof of our no-go theorem, Eq. (2.17); that is, of the fact that the inequality

$$s_q \leq 1 - \frac{c_q}{2} \tag{6.230}$$

must be satisfied for any s_q and c_q arising in a noncontextual model that reproduces the data in Table 2.1 and respects the operational equivalence of Eq. (2.12). While the proof provided in the main text uses an intuitive argument that is native to the task of state discrimination, the argument in this section abstracts away from the specific problem at hand, and as such extends naturally to the more general method required for proving Eq. (2.35) (as discussed in Appendix A.2).

First, we allow the ontological model to have an ontic state space of arbitrary form, and we allow the response functions to be outcome-indeterministic. Second, we show that for any such model, there exists a simpler ontological model which is equally general, but which has only 8 ontic states and has response functions that are purely outcome-deterministic. Third, we show that two of these ontic states are superfluous if B_d is optimal for state discrimination. Fourth, we show that the forms of the epistemic states are greatly constrained by their perfectly predictable responses on the corresponding measurements. Fifth, we parametrize the set of possible epistemic states as probability distributions over the remaining 6 ontic states in accordance with these constraints. Sixth, we calculate the values of s_q and c_q in terms of these response functions and epistemic states. Finally, we impose preparation noncontextuality and eliminate the unobserved variables to obtain the optimal tradeoff between s_q and c_q.

As one ranges over the ontic states in our ontological model, the vector $(\xi_{\phi|B_\phi}(\lambda), \xi_{\psi|B_\psi}(\lambda), \xi_{g_\phi|B_d}(\lambda))$ of valid probability assignments to our three binary basis measurements defines a unit cube. The most obvious ontological model would have one λ for each possible probability assignment (including the indeterministic ones), defining an ontic state space isomorphic to the unit cube. The epistemic states in such an outcome-indeterministic model would be arbitrary normalized probability densities over this set of ontic states (that is, all the interior points of the cube).

However, we can always simplify matters without loss of generality by decomposing each non-extremal probability value assignment into extremal assignments. (These extremal points are outcome-deterministic if and only if there are no nontrivial constraints from measurement noncontextuality, but this is indeed the case here.)

Let us define a variable κ which runs over the eight extremal points in the cube of ontic states. Then, there exists a $p(\kappa|\lambda)$ such that $\xi_{k|M}(\lambda) = \sum_\kappa \xi_{k|M}(\kappa)p(\kappa|\lambda)$. We can thus write any observable probability $p(k|M, P)$ as

$$p(k|M, P) = \int_\Lambda \xi_{k|M}(\lambda)\mu_P(\lambda)d\lambda = \sum_\kappa \xi_{k|M}(\kappa)\mu_P(\kappa) \tag{6.231}$$

where $\mu_P(\kappa) \equiv \int_\Lambda d\lambda\; p(\kappa|\lambda)\mu_P(\lambda)$. This construction lets us write observed probabilities in terms of extremal value assignments by effectively moving uncertainty into the new state distributions $\mu_P(\kappa)$.

We sometimes simplify the notation by letting the distributions and response functions be vectors of probabilities indexed by the ontic states κ; e.g.

$$p(k|M, P) = \sum_\kappa \xi_{k|M}(\kappa)\mu_P(\kappa) = \vec{\xi}_{k|M} \cdot \vec{\mu}_P. \tag{6.232}$$

We thus convert an outcome-indeterministic model over a continuum of ontic states (the unit cube) to an outcome-deterministic model over just 8 ontic states (its vertices), without any loss of generality. The vertices κ_1 to κ_8 correspond to the deterministic triples

$$\begin{aligned}&\Big(\xi_{\phi|B_\phi}(\kappa), \xi_{\psi|B_\psi}(\kappa), \xi_{g_\phi|B_d}(\kappa)\Big) \in \\ &\qquad \{(0,0,0), (0,0,1), (0,1,0), ..., (1,1,1)\},\end{aligned} \tag{6.233}$$

so the three response functions are

$$\vec{\xi}_{\phi|B_\phi} = \begin{bmatrix}0\\0\\0\\0\\1\\1\\1\\1\end{bmatrix}, \quad \vec{\xi}_{\psi|B_\psi} = \begin{bmatrix}0\\0\\1\\1\\0\\0\\1\\1\end{bmatrix}, \quad \vec{\xi}_{g_\phi|B_d} = \begin{bmatrix}0\\1\\0\\1\\0\\1\\0\\1\end{bmatrix}. \tag{6.234}$$

In fact, if we assume B_d is optimal, the fourth and fifth of these value assignments will never occur. Consider for example the triple $(1,0,0)$ (which occurs for κ_5). Since $\xi_{\phi|B_\phi}(\kappa_5) = 1$, the state cannot have been $\bar{\phi}$. Since $\xi_{\psi|B_\psi}(\kappa_5) = 0$, the state cannot have been ψ. Thus, we know the state must have been ϕ or $\bar{\psi}$; in either case, the winning strategy is for B_d to return the outcome g_ϕ. Therefore the winning strategy has $\xi_{g_\phi|B_d}(\kappa_5) = 1$, and thus the triple $(1,0,0)$ never occurs in the winning strategy. Similar logic applies to the triple $(0,1,1)$, and hence we need not consider these two assignments [24]. The remaining value assignments are

$$\begin{aligned}&\left(\xi_{\phi|B_\phi}(\kappa), \xi_{\psi|B_\psi}(\kappa), \xi_{g_\phi|B_d}(\kappa)\right) \in \\ &\{(0,0,0),(0,0,1),(0,1,0),(1,0,1),(1,1,0),(1,1,1)\}.\end{aligned} \tag{6.235}$$

Thus six ontic states are sufficient for describing our experiment: one for each remaining deterministic assignment. It follows that the vectors representing each of the three response functions are:

$$\vec{\xi}_{\phi|B_\phi} = \begin{bmatrix}0\\0\\0\\1\\1\\1\end{bmatrix}, \quad \vec{\xi}_{\psi|B_\psi} = \begin{bmatrix}0\\0\\1\\0\\1\\1\end{bmatrix}, \quad \vec{\xi}_{g_\phi|B_d} = \begin{bmatrix}0\\1\\0\\1\\0\\1\end{bmatrix}. \tag{6.236}$$

[24] These assumptions for B_d ensure that the relationship we derive between s_q and c_q will saturate the bound on s_q implied by any noncontextual model. If we had not used this argument, we would obtain the same relationship, but only as a bound on s_q, not as the saturating equality. This is easily verified explicitly, e.g. by taking $\epsilon = 0$ in Appendix A.2 below. However, including two more ontic states requires considerably more algebra.

We can constrain the most general form of the epistemic states using the perfect predictability of measurements B_ϕ and B_ψ on their corresponding states. Namely, recalling Eq. (B.3) and the form of the response functions, $\xi_{\phi|B_\phi}(\kappa)$, $\xi_{\psi|B_\psi}(\kappa)$, $\xi_{\bar{\psi}|B_\psi}(\kappa) \equiv 1-\xi_{\psi|B_\psi}(\kappa)$, and $\xi_{\bar{\phi}|B_\phi}(\kappa) \equiv 1-\xi_{\phi|B_\phi}(\kappa)$, we can see that our epistemic states must have the form

$$\vec{\mu}_\phi = \begin{bmatrix} 0 \\ 0 \\ 0 \\ a_4 \\ a_5 \\ a_6 \end{bmatrix}, \ \vec{\mu}_{\bar{\phi}} = \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \ \vec{\mu}_\psi = \begin{bmatrix} 0 \\ 0 \\ b_3 \\ 0 \\ b_5 \\ b_6 \end{bmatrix}, \ \vec{\mu}_{\bar{\psi}} = \begin{bmatrix} b_1 \\ b_2 \\ 0 \\ b_4 \\ 0 \\ 0 \end{bmatrix}, \tag{6.237}$$

where normalization requires that $a_4 + a_5 + a_6 = 1$, and so on.

The definitions of c_q and s_q in Eqs. (2.8) and (2.10) translated into our ontological model become

$$c_q = \vec{\mu}_\psi \cdot \vec{\xi}_{\phi|B_\phi} = \vec{\mu}_\phi \cdot \vec{\xi}_{\psi|B_\psi} = 1 - \vec{\mu}_{\bar{\psi}} \cdot \vec{\xi}_{\phi|B_\phi} = 1 - \vec{\mu}_{\bar{\phi}} \cdot \vec{\xi}_{\psi|B_\psi}, \tag{6.238}$$

$$s_q = \vec{\mu}_\phi \cdot \vec{\xi}_{g_\phi|B_d} = 1 - \vec{\mu}_\psi \cdot \vec{\xi}_{g_\phi|B_d} = 1 - \vec{\mu}_{\bar{\phi}} \cdot \vec{\xi}_{g_\phi|B_d} = \vec{\mu}_{\bar{\psi}} \cdot \vec{\xi}_{g_\phi|B_d}. \tag{6.239}$$

Taking these dot products using the vectors in Eq. (6.236) and Eq. (6.237) gives

$$c_q = b_5 + b_6 = a_5 + a_6 = 1 - b_4 = 1 - a_3 \tag{6.240}$$

and

$$s_q = a_4 + a_6 = 1 - b_6 = 1 - a_2 = b_2 + b_4. \tag{6.241}$$

Because the epistemic states must be normalized, it follows that $b_5 + b_6 = 1 - b_3$, $a_5 + a_6 = 1 - a_4$, $a_4 + a_6 = 1 - a_5$, and $b_2 + b_4 = 1 - b_1$. Substituting these four expressions, we obtain

$$c_q = 1 - b_3 = 1 - a_4 = 1 - b_4 = 1 - a_3 \tag{6.242}$$

and

$$s_q = 1 - a_5 = 1 - b_6 = 1 - a_2 = 1 - b_1, \tag{6.243}$$

and hence $b_3 = a_4 = b_4 = a_3$ and $a_5 = b_6 = a_2 = b_1$.

Let us take $s_q = 1 - a_2$ and $c_q = 1 - a_3$. If there were no more constraints, then a_2 and a_3 could range from 0 to 1 independently, and s_q and c_q could take any values from 0 to 1. By imposing preparation noncontextuality, however, we have

$$\vec{\mu}_{\frac{1}{2}} = \frac{1}{2}\begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \\ a_6 \end{bmatrix} = \frac{1}{2}\begin{bmatrix} b_1 \\ b_2 \\ b_3 \\ b_4 \\ b_5 \\ b_6 \end{bmatrix}, \tag{6.244}$$

This implies $b_i = a_i$ for all i. Since $a_1 + a_2 + a_3 = 1$ from normalization, $a_1 = b_1$ from preparation noncontextuality, and $b_1 = a_2$ as derived above, we also have $2a_2 + a_3 = 1$ and hence $c_q = 2a_2$. Finally, writing s_q in terms of c_q yields

$$s_q = 1 - \frac{c_q}{2}. \tag{6.245}$$

A.2 Proof of noncontextuality inequality for MESD

Herein we prove our noncontextuality inequality, Eq. (2.35); that is, we prove that

$$s \leq 1 - \frac{c - \epsilon}{2} \tag{6.246}$$

must be satisfied for any s, c, and ϵ arising in a noncontextual model that reproduces data in Table 2.2 and respects Eq. (2.29).

First, we use the arguments of Appendix A.1 to write down an ontological model with 8 ontic states and purely outcome-deterministic response functions. Second, we parametrize the set of possible epistemic states for this second model in accordance with preparation noncontextuality. Third, we calculate expressions for s, c, and ϵ in terms of these response functions and epistemic states. These manipulations reduce the problem to a small set of linear equalities and inequalities over unobserved and observed variables. Finally, we eliminate the unobserved variables to obtain inequalities concerning only the observed variables s, c, and ϵ.

Exactly as before, we can convert a general, outcome-indeterministic model over a continuum of ontic states (the unit cube) to an outcome-deterministic model over just

8 ontic states (its vertices), without any loss of generality. (As before, this is simply a mathematical construction, and in no way commits us to a fundamental principle of outcome-determinism.) The vertices of the unit cube, κ_1 to κ_8, again correspond to the deterministic triples

$$\begin{aligned}&\Big(\xi_{\phi|M_\phi}(\kappa), \xi_{\psi|M_\psi}(\kappa), \xi_{g_\phi|M_d}(\kappa)\Big) \in \\ &\qquad \{(0,0,0),(0,0,1),(0,1,0),...,(1,1,1)\},\end{aligned} \tag{6.247}$$

and the three response functions are again

$$\vec{\xi}_{\phi|M_\phi} = \begin{bmatrix}0\\0\\0\\0\\1\\1\\1\\1\end{bmatrix}, \quad \vec{\xi}_{\psi|M_\psi} = \begin{bmatrix}0\\0\\1\\1\\0\\0\\1\\1\end{bmatrix}, \quad \vec{\xi}_{g_\phi|M_d} = \begin{bmatrix}0\\1\\0\\1\\0\\1\\0\\1\end{bmatrix}. \tag{6.248}$$

(In a more general situation in which measurement noncontextuality is also leveraged, there will be linear constraints on this set of response functions, and the extremal response functions will no longer all be outcome-deterministic. In this case, one can still explicitly enumerate the finite set of extremal response functions by taking the intersection of the linear constraints with the above cube of value assignments. These extremal points modify the specific form of Eq. (6.248), and our methods would proceed largely unchanged.)

Each preparation generates a probability distribution over κ, so we can write the epistemic states as

$$\vec{\mu}_{P_\phi} = \begin{bmatrix}a_1\\a_2\\a_3\\a_4\\a_5\\a_6\\a_7\\a_8\end{bmatrix}, \quad \vec{\mu}_{P_\psi} = \begin{bmatrix}b_1\\b_2\\b_3\\b_4\\b_5\\b_6\\b_7\\b_8,\end{bmatrix}, \quad \vec{\mu}_{P_{\bar{\phi}}} = \begin{bmatrix}c_1\\c_2\\c_3\\c_4\\c_5\\c_6\\c_7\\c_8,\end{bmatrix}, \quad \vec{\mu}_{P_{\bar{\psi}}} = \begin{bmatrix}d_1\\d_2\\d_3\\d_4\\d_5\\d_6\\d_7\\d_8,\end{bmatrix}, \tag{6.249}$$

where the parameters in each vector are positive and sum to 1.

Dot products between a vector in Eq. (6.248) and a vector in Eq. (6.249) can produce any set of observable statistics, and thus constitute a general ontological model for our measurements and preparations. The values of (s, c, ϵ) that we can observe in a noncontextual model with our assumed symmetries, however, are restricted by the above constraints, all of which we repeat here for convenience.

Eqs. (2.2) and (2.3) imply that for all four preparations,

$$\forall \kappa : 0 \leq [\vec{\mu}_P]_\kappa \leq 1 \tag{6.250}$$

and

$$\sum_k [\vec{\mu}_P]_k = 1. \tag{6.251}$$

Eq. (2.29) gives

$$\vec{\mu}_{P_\phi} + \vec{\mu}_{P_{\bar{\phi}}} = \vec{\mu}_{P_\psi} + \vec{\mu}_{P_{\bar{\psi}}}. \tag{6.252}$$

Eqs. (2.30)-(2.32) are, respectively,

$$\vec{\mu}_{P_\phi} \cdot \vec{\xi}_{g_\phi|M_d} = 1 - \vec{\mu}_{P_\psi} \cdot \vec{\xi}_{g_\phi|M_d} = s. \tag{6.253}$$

$$\vec{\mu}_{P_\psi} \cdot \vec{\xi}_{\phi|M_\phi} = \vec{\mu}_{P_\phi} \cdot \vec{\xi}_{\psi|M_\psi} = c, \tag{6.254}$$

$$\vec{\mu}_{P_\psi} \cdot \vec{\xi}_{\psi|M_\psi} = \vec{\mu}_{P_\phi} \cdot \vec{\xi}_{\phi|M_\phi} = 1 - \epsilon, \tag{6.255}$$

Eq. (2.33) gives

$$\epsilon \leq c \leq 1 - \epsilon. \tag{6.256}$$

Eqs. (6.250)-(6.256) define a set of constraints over the variables $s, c, \epsilon, a_i, b_i, c_i$, and d_i (where $i \in \{1, 2, ..., 8\}$). Although the parameters a_i, b_i, c_i, d_i in our epistemic states are not observable, constraints upon them (Eqs. (6.250) and (6.251)) have consequences for the set of possible triples (s, c, ϵ). Finding the set of inequalities over only (s, c, ϵ) that is implied by the full set of linear equalities and inequalities of Eqs. (6.250)-(6.256) is algebraically tedious by hand, but straightforward using the well-known Fourier Motzkin Elimination algorithm, which returns our result

$$s \leq 1 - \frac{c - \epsilon}{2}. \tag{6.257}$$

It is worth noting that the technique for deriving noncontextuality inequalities we have introduced here, insofar as it reduces to a convex hull problem, is an instance of the problem of quantifier elimination. Recent work in quantum foundations has seen increasing use of quantifier elimination algorithms, in noncontextuality [163, 164] as well as other scenarios. Fourier-Motzkin elimination, which is appropriate for problems wherein the dependence on the variables to be eliminated is linear, has been used to derive Bell inequalities [46], and also recently, to derive Bell-like inequalities for novel causal scenarios [62, 66, 323]. In Ref. [323], where the problem is reduced to what is known as the classical marginals problem—that of determining whether a given set of distributions on various subsets of a set of variables can arise as the marginals of a single joint distribution over all of the variables—this problem can be solved by performing quantifier elimination on the probabilities in the joint distributions using convex hull algorithms. Nonlinear quantifier elimination using cylindrical algebraic decomposition has also found application in deriving Bell-like inequalities in simple scenarios [62, 175]. We anticipate that these more general techniques for quantifier elimination will ultimately also find applications to the derivation of noncontextuality inequalities.

A.3 Noisy quantum realization which violates our noncontextuality inequality

We now sketch a quantum realization of the MESD scenario for any given values of c and ϵ satisfying the assumed symmetries and operational equivalences and violating our noncontextuality inequality for all values of c and ϵ. (The ideal quantum realization of the MESD scenario, given earlier, was defined only for $\epsilon = 0$.)

There is no general technique for finding the set of all data tables achievable in quantum theory for some prepare-and-measure scenario. For some cases (e.g., Bell tests), this set can be approximated efficiently via the Navascues-Pironio-Acin hierarchy [210]. For situations with multiple preparations or additional constraints, no such method exists yet.

However, we can apply our understanding from Section 2.6.1 to construct a quantum model which recovers Eq. (2.36), which we conjecture is optimal for qubits. Namely, because we want to find the maximum value of s consistent with a given c and ϵ, we should attribute as much of the confusability as possible to noise in the M_ϕ and M_ψ measurements, and only attribute the remainder of the confusability to nonorthogonality of the states. As such, in this section we allow the effects E_ϕ, E_ψ, $E_{\bar{\phi}}$, and $E_{\bar{\psi}}$ to be noisy POVM elements (unlike in Appendix A.1, where E_ϕ denoted a projector onto $|\phi\rangle$, and so on).

Imagine P_ϕ prepares state $|0\rangle$ on the Bloch sphere and P_ψ prepares a pure state $|\theta\rangle$

rotated by $\theta \in [0, \pi]$ with respect to $|0\rangle$ in the $X - Z$ plane. We will specify the value of θ later. Within this plane, the effect E_ϕ must lie on the green line shown in Fig. 6.1, since only these effects imply $\langle 0| E_\phi |0\rangle = 1 - \epsilon$.

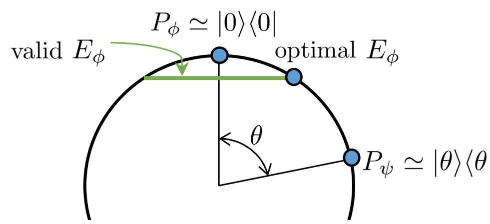

Figure 6.1: Sketch of the quantum model which yields Eq. (2.36).

The choice of E_ϕ that yields the maximum confusability is the one on the green line, closest to $|\theta\rangle$ (but not closer to $|\theta\rangle$ than to $|0\rangle$, since that would imply that $c \geq 1 - \epsilon$). The remainder of the confusability must then be attributed to the nonzero inner product between the two pure states, so θ is fixed by $\langle\theta| E_\phi |\theta\rangle = c$. Now that the two states are specified, calculating the optimal (Helstrom) probability is a simple quantum calculation whose result gives Eq. (2.36), that is

$$s = \frac{1}{2}(1 + \sqrt{1 - \epsilon + 2\sqrt{\epsilon(1-\epsilon)c(c-1)} + c(2\epsilon - 1)}). \tag{6.258}$$

The remaining states and effects are completely fixed by the assumed symmetries and operational equivalence. For a general pair of c and ϵ, this quantum model outperforms the optimal noncontextual model, as seen in Fig. 2.5.

A.4 Full set of noncontextuality inequalities for MESD without symmetries

As promised in Section 2.5.1, we now derive the full set of noncontextuality inequalities for our operational MESD scenario when the symmetries of Eqs. (2.30)-(2.32) are not assumed. In Table 6.1 we show a general data table for 3 binary measurements and 4 preparations which respect our operational equivalence. There are 9 free parameters, since the probabilities in the last column are fixed by those in the first three.

	P_ϕ	P_ψ	$P_{\bar{\phi}}$	$P_{\bar{\psi}}$
$\phi\vert M_\phi$	$1-\epsilon_\phi$	c_ψ	$\epsilon_{\bar{\phi}}$	$1-\epsilon_\phi+\epsilon_{\bar{\phi}}-c_\psi$
$\psi\vert M_\psi$	c_ϕ	$1-\epsilon_\psi$	$1-c_{\bar{\phi}}$	$c_\phi-c_{\bar{\phi}}+\epsilon_\psi$
$g_\phi\vert M_d$	s_ϕ	$1-s_\psi$	$1-s_{\bar{\phi}}$	$s_\phi-s_{\bar{\phi}}-s_\psi$

Table 6.1: Data table for our operational scenario with no symmetries assumed. There are 9 free parameters.

The procedure from Appendix A.2 yields the following set of inequalities over the 9 free parameters, which are necessary and sufficient for the data to have been generated by a noncontextual model respecting operational equivalence Eq. (2.29):

$$
\begin{aligned}
&0 \leq s_\phi \leq 1 \\
&0 \leq s_{\bar{\phi}} \leq 1 \\
&0 \leq s_\psi \leq 1 \\
&0 \leq \epsilon_\phi \leq c_\phi \leq 1 - \epsilon_\phi \\
&0 \leq \epsilon_{\bar{\phi}} \leq c_{\bar{\phi}} \leq 1 - \epsilon_{\bar{\phi}} \\
&0 \leq \epsilon_\psi \leq c_\psi \leq 1 - \epsilon_\psi \\
&0 \leq s_\phi - s_{\bar{\phi}} + s_\psi \leq 1 \\
&0 \leq c_\phi - c_{\bar{\phi}} + \epsilon_\psi \\
&0 \leq c_\psi + s_{\bar{\phi}} - s_\psi + \epsilon_\phi \\
&0 \leq c_\psi - s_{\bar{\phi}} + s_\psi + \epsilon_\phi \\
&0 \leq -c_\psi + s_\phi + s_\psi + \epsilon_{\bar{\phi}} \\
&0 \leq c_\phi + s_{\bar{\phi}} - s_\psi + \epsilon_\psi \\
&0 \leq -c_{\bar{\phi}} + s_\phi + s_\psi + \epsilon_\psi \\
&0 \leq c_\phi - s_{\bar{\phi}} + s_\psi + \epsilon_\psi \\
&0 \leq -c_{\bar{\phi}} + c_\psi + \epsilon_\phi + \epsilon_\psi \\
&0 \leq c_\phi - c_\psi + \epsilon_{\bar{\phi}} + \epsilon_\psi \\
&0 \leq 2 - c_\psi - s_\phi - s_\psi + \epsilon_{\bar{\phi}} \\
&0 \leq 2 - c_{\bar{\phi}} - s_\phi - s_\psi + \epsilon_\psi \\
&0 \leq -c_\phi + c_{\bar{\phi}} + c_\psi + \epsilon_\phi - \epsilon_{\bar{\phi}} - \epsilon_\psi \\
&0 \leq 1 - c_\phi + c_{\bar{\phi}} - c_\psi - \epsilon_\phi + \epsilon_{\bar{\phi}} - \epsilon_\psi
\end{aligned} \tag{6.259}
$$

Of course, these inequalities reproduce Eq. (2.35) if the symmetries are now imposed.

In deriving these inequalities, we have assumed the logical labeling of Eq. (2.33). If one drops the labeling condition, then the resulting inequalities are identical to the facets of the Bell polytope discussed in Section 2.8 (but have no practical relevance to minimum error state discrimination).

B Appendices for Chapter 3

B.1 The stabilizer subtheory

We here expand on the exposition of the stabilizer subtheory from the main text (with some redundancy for completeness).

The stabilizer subtheory is built around the Clifford group, whose elements will be referred to as Clifford unitaries. To define these, we first introduce the *Weyl operators* (also called generalized Pauli operators). Consider a d-dimensional quantum system, and define the computational basis $\{|0\rangle, \ldots, |d-1\rangle\}$ in its Hilbert space $\mathcal{H}$. Each basis element is labelled by an element of $\mathbb{Z}_d$[25], which we refer to as the configuration space. Writing $\omega = \exp(\frac{2\pi i}{d})$, we define the translation operator X and boost operator Z via

$$X|x\rangle = |x+1\rangle \tag{6.260}$$

$$Z|x\rangle = \omega^x |x\rangle . \tag{6.261}$$

Note that here and throughout, all arithmetic is within $\mathbb{Z}_d$. These can be viewed as discrete position and momentum translation operators, respectively, for a particle on a ring. From these, the single-system Weyl operators are defined as

$$W_{p,q} = Z^p X^q, \tag{6.262}$$

where $p, q \in \mathbb{Z}_d$. Note that these are often defined with an additional phase factor $\omega^{\gamma_{pq}}$; however, the choice of this phase is irrelevant for the definition of the stabilizer subtheory, so we set γ_{pq} to zero. (We highlight this irrelevance by introducing the stabilizer subtheory using superoperators, for which any choice of phase cancels.)

The Weyl operators are unitaries whose associated superoperators, $\mathcal{W}_{p,q}(\cdot) := W_{p,q}(\cdot)W_{p,q}^\dagger$, form a group with composition law

$$\mathcal{W}_{p,q}\mathcal{W}_{p',q'} = \mathcal{W}_{p+p',q+q'}, \tag{6.263}$$

and inverse

$$\mathcal{W}_{p,q}^{-1} = \mathcal{W}_{p,q}^\dagger = \mathcal{W}_{-p,-q}. \tag{6.264}$$

(Note that the Weyl operators themselves do not form a group as the above equations only hold up to a particular phase factor.)

[25] When d is prime, $\mathbb{Z}_d$ has the structure of a finite algebraic field. For non-prime d, things are somewhat more complicated [133], but the results in this work still hold.

It will be useful later to note that the Weyl operators are orthonormal with respect to a rescaled Hilbert-Schmidt inner product:

$$\frac{1}{d}\,\mathsf{tr}[W_{p,q}W^{\dagger}_{p',q'}] = \delta_{p,p'}\delta_{q,q'}. \tag{6.265}$$

The Clifford unitaries are defined as unitaries which—up to a phase—map Weyl operators to other Weyl operators under conjugation. Equivalently, their associated superoperators map Weyl superoperators to other Weyl superoperators under conjugation. That is, $\mathcal{U}$ is a Clifford unitary superoperator if for every p, q, one has

$$\mathcal{U}\mathcal{W}_{p,q}\mathcal{U}^{\dagger} = \mathcal{W}_{p',q'}. \tag{6.266}$$

Let us now define the *phase space* $V := \mathbb{Z}_d \times \mathbb{Z}_d$, which is a module[26] equipped with the symplectic product $[\cdot,\cdot] : V \times V \to \mathbb{Z}_d$ given by

$$\left[\begin{pmatrix} p \\ q \end{pmatrix}, \begin{pmatrix} p' \\ q' \end{pmatrix}\right] := pq' - qp'. \tag{6.267}$$

Note that each Weyl operator is labeled by a phase space point $(p, q) = a \in V$. A function $f : V \to V$ is said to be linear if $f(\lambda a + b) = \lambda f(a) + f(b)$, for $\lambda \in \mathbb{Z}_d$, $a, b \in V$. A function $S : V \to V$ is called symplectic if it is linear and preserves the symplectic product, i.e. $[S\cdot, S\cdot] = [\cdot,\cdot]$. A transformation of the form $S \cdot +a$ where S is symplectic and $a \in V$ is called a symplectic affine transformation. Note that the symplectic functions form a group, and that the symplectic affine transformations also form a group.

As shown in Ref. [133], every Clifford superoperator is of the form $\mathcal{W}_a\mathcal{M}_S$, where $\mathcal{W}_a$ is a Weyl superoperator labelled by $a \in V$, $S : V \to V$ is a symplectic function, $\mathcal{M}$ is a unitary superoperator representation of the symplectic group (i.e. $\mathcal{M}_S\mathcal{M}_T = \mathcal{M}_{ST}$), and where $\mathcal{M}_S\mathcal{W}_v\mathcal{M}_S^{\dagger} = \mathcal{W}_{Sv}$ for any symplectic function S and for all $v \in V$.

Hence, each Clifford operation can be indexed by a phase space vector a and a symplectic map S, and so we will denote them by $\mathcal{C}_{a,S} := \mathcal{W}_a\mathcal{M}_S$. Clearly, a Weyl operator $\mathcal{W}_{p,q}$ is a Clifford unitary $\mathcal{C}_{a,S}$, where $a = (p, q)$ and $S = \mathbb{1}$. Furthermore, the mapping $S \cdot +a \mapsto \mathcal{W}_a\mathcal{M}_S$ is a representation of the group of symplectic affine transformations [133].

The Clifford unitary superoperators form a group, often termed the Clifford group, with composition rule

$$\mathcal{C}_{a,S}\mathcal{C}_{b,T} = \mathcal{C}_{Sb+a,ST}. \tag{6.268}$$

[26] If d is a prime power $d = p^k$, then this is moreover a finite vector space.

The inverse of a Clifford unitary superoperator is

$$\mathcal{C}_{a,S}^{-1} = \mathcal{C}_{a,S}^{\dagger} = \mathcal{C}_{-S^{-1}a,S^{-1}}. \tag{6.269}$$

It is therefore clear that the Clifford superoperator group in dimension d and the symplectic affine group for $\mathbb{Z}_d \times \mathbb{Z}_d$ are isomorphic groups.

For a fixed dimension, the Clifford group is generated by the superoperators associated to the generalized Hadamard gate H and the generalized phase gate P [102], defined respectively by

$$H\,|x\rangle = \frac{1}{\sqrt{d}} \sum_{k\in\mathbb{Z}_d} \omega^{xk}\,|k\rangle\,, \tag{6.270}$$

$$P\,|x\rangle = \omega^{\frac{1}{2}x(x+d)}\,|x\rangle\,. \tag{6.271}$$

The stabilizer subtheory for a single system in dimension d is defined as the set of processes which can be generated by sequential composition of: i) pure states uniquely identified by being the simultaneous eigenstates of a given set of Weyl operators, ii) projective measurements in the spectral decomposition of the Weyl operators[27], and iii) Clifford unitary superoperators on the associated Hilbert space, as well as convex mixtures of such processes.

This construction is easily generalized to allow for parallel composition, that is, for systems made up of n qu$d$its[28], by defining the multiparticle Weyl operators as tensor products of those defined above, and defining the multiparticle Clifford operators as unitary superoperators that preserve the multiparticle Weyl operators under conjugation; see Ref. [133] for more details. An important feature is that in general the stabilizer subtheory defined by parallel composition of n qu$d$its is not the same as the stabilizer subtheory defined by a single d^n dimensional system—for instance, the latter generally has far fewer states [133]. Therefore, for a given dimension D there may be multiple different stabilizer theories which could be associated to it, depending on whether one views it as a single monolithic system of dimension D (which Gross calls the single-particle view), or views it as some tensor product of multiple qu$d$its (which Gross calls a multi-particle view).

[27] Note that although the Weyl operators are not Hermitian operators, they *are* normal operators, and hence have a spectral decomposition, which implies one can carry out a projective measurement in the eigenbasis of each.

[28] To the authors' knowledge, parallel composition of systems of different dimensions is at best highly nontrivial, and has not been considered in the literature.

B.2 Useful Preliminaries

It is well-known that a basis of a vector space uniquely defines a dual basis in the dual vector space (i.e. the space of functionals on the vector space). We will leverage this fact, but in a slightly different form:

Lemma B.1. *Given any basis* $\{F_\lambda\}_\lambda$ *for a* d^2*-dimensional real vector space* $\mathsf{Herm}(\mathcal{H})$ *of Hermitian operators on a Hilbert space* $\mathcal{H}$*, there is a unique set* $\{D_\lambda\}_\lambda$ *of* d^2 *Hermitian operators satisfying*

$$\mathsf{tr}(D_{\lambda'} F_\lambda) = \delta_{\lambda,\lambda'}, \tag{6.272}$$

and $\{D_\lambda\}_\lambda$ *also forms a basis for* $\mathsf{Herm}(\mathcal{H})$*.*

Proof. Consider any basis $\{F_\lambda\}_\lambda$ of $\mathsf{Herm}(\mathcal{H})$. It uniquely specifies a basis $\{\mathcal{D}_\lambda\}_\lambda$ of the dual vector space $\mathsf{Herm}(\mathcal{H})^*$, where $\{\mathcal{D}_\lambda\}_\lambda$ are linear functionals satisfying $\mathcal{D}_{\lambda'}(F_\lambda) = \delta_{\lambda,\lambda'}$.[29] Now, in order to obtain again a set of Hermitian operators $\{D_\lambda\}_\lambda$, we use the Riesz representation theorem [238], which states that each of these functionals $\mathcal{D}_\lambda$ can be written as the Hilbert-Schmidt inner product with a unique Hermitian operator D_λ, namely

$$\mathcal{D}_\lambda(\cdot) = \mathsf{tr}[(\cdot) D_\lambda]. \tag{6.273}$$

This picks out a unique basis $\{D_\lambda\}_\lambda$ which satisfies Eq. (6.272). □

Note that the operators $\{F_\lambda\}_\lambda$ and $\{D_\lambda\}_\lambda$ are both in $\mathsf{Herm}(\mathcal{H})$. For a basis $\{F_\lambda\}_\lambda$, we refer to the set $\{D_\lambda\}_\lambda$ constructed using this lemma as the *dual basis*.

Another useful lemma we will require is the following.

Lemma B.2. *A nonnegative and diagram-preserving quasiprobabilistic representation of any unitary superoperator* $\mathcal{U}(\cdot) := U(\cdot)U^\dagger$ *is given by a permutation; that is, by a conditional probability distribution*

$$\xi_{\mathcal{U}}(\lambda'|\lambda) = \delta_{\sigma_U(\lambda'),\lambda} \tag{6.274}$$

for some permutation $\sigma_U : \Lambda \to \Lambda$*.*

Proof. By definition, a nonnegative quasiprobabilistic representation ξ represents every unitary superoperator $\mathcal{U}$ as a stochastic map from Λ to itself, so $\xi_{\mathcal{U}}$ and $\xi_{\mathcal{U}^\dagger}$ are stochastic

[29] To see that this is unique, consider a linear functional $\mathcal{D}'_{\lambda'}$ satisfying $\mathcal{D}'_{\lambda'}(F_\lambda) = \delta_{\lambda,\lambda'}$ for all λ. Since a linear functional is fully specified by its action on a basis, $\mathcal{D}'_{\lambda'}$ is the exact same functional as $\mathcal{D}_{\lambda'}$.

maps. By diagram preservation, it holds that $\xi_{\mathcal{U}\mathcal{U}^\dagger} = \xi_{\mathcal{U}} \circ \xi_{\mathcal{U}^\dagger}$. But $\mathcal{U}\mathcal{U}^\dagger = \mathbb{1}$, and hence $\xi_{\mathcal{U}\mathcal{U}^\dagger} = \xi_{\mathbb{1}}$, where (by diagram preservation) $\xi_{\mathbb{1}}$ must be the identity matrix. Therefore $\xi_{\mathcal{U}^\dagger} \circ \xi_{\mathcal{U}}$ is the identity matrix, so $\xi_{\mathcal{U}^\dagger}$ is the left inverse of $\xi_{\mathcal{U}}$, and so (by the fact that they are square matrices) $\xi_{\mathcal{U}}$ and $\xi_{\mathcal{U}^\dagger}$ are inverses. But the only (square) stochastic matrices whose inverses are stochastic are permutations. Hence $\xi_{\mathcal{U}}$ is a permutation for every unitary U. □

A final useful lemma is a well-known result from Ref. [281]:

Lemma B.3. *Projective measurements have an outcome-deterministic representation in any noncontextual ontological model. That is, representation of the projectors in a projective measurement are conditional probability distributions valued in* $\{0, 1\}$*. Furthermore, every ontic state is in the support of the representations of one and only one of eigenstates in any given projective measurement.*

This lemma was originally proven for full quantum theory, but it immediately generalizes to the stabilizer subtheory.

B.3 Proof of Main Theorem

Here we prove

Theorem 18.

(a) For any stabilizer subtheory (single- or multi-particle) in **odd** *dimensions, the* unique *nonnegative and diagram-preserving quasiprobability representation for it is Gross's representation.*

(b) For any stabilizer subtheory (single- or multi-particle) in **even** *dimensions, there is no nonnegative and diagram-preserving quasiprobability representation.*

We first give a one-paragraph intuitive proof sketch. Recall that the structure theorem of Ref. [260] gives an exact frame representation as discussed in the main text. Starting with the single-particle case, we leverage the fact that noncontextuality implies outcome determinism to find a privileged labeling of the ontic states as points in phase space. We show that this implies translational covariance: that is, Clifford covariance for all Weyl operators. Using this and the fact that Weyl operators form a basis of the linear operators, we then show that the representation is fixed by the outcomes of measurements of Weyl operators on the $\lambda = (0, 0)$ ontic state [30]. We give various conditions on these outcomes due to the Hermiticity of the phase point operator, the representation of the Hadamard (which we show to be covariant), and from considering measurements of commuting pairs of Weyl operators. In odd dimensions, we show that the unique solution to these conditions is that which gives Gross's phase point operators. In even dimensions, we show that there is no solution. The generalization to multi-particle stabilizer subtheories is then shown to follow immediately.

We now give the full proof.

We start from the assumption that we have *some* nonnegative and diagram-preserving quasiprobability representation of the stabilizer subtheory in some finite dimension d. Note that this subtheory is tomographically local, and has GPT dimension d^2. Hence, Corollary VI.2 of Ref. [260] implies that the number of elements in the sample space is exactly d^2. Since a nonnegative and diagram-preserving quasiprobability representation is equivalent to a noncontextual ontological representation, we will refer to the elements of the sample space as 'ontic states'.

The structure theorems in Ref. [260] (in particular, Corollary VI.2) imply that this representation is an *exact frame representation* [103] composed of a basis $\{F_\lambda\}_\lambda$ and its dual $\{D_\lambda\}_\lambda$ (in the sense of Lemma B.1,) such that the representation of a completely positive trace preserving map $\mathcal{E}$ is given by the conditional quasiprobability distribution

$$\xi_{\mathcal{E}}(\lambda'|\lambda) = \mathsf{tr}[D_{\lambda'}\mathcal{E}(F_\lambda)]. \tag{6.275}$$

Here, $\{F_\lambda\}_\lambda$ is a spanning and linearly independent set of d^2 Hermitian operators, as is $\{D_\lambda\}_\lambda$, where these satisfy

$$\mathsf{tr}[F_\lambda] = 1, \tag{6.276}$$

$$\sum_\lambda D_\lambda = \mathbb{1}, \tag{6.277}$$

and

$$\mathsf{tr}[D_{\lambda'}F_\lambda] = \delta_{\lambda\lambda'}. \tag{6.278}$$

(Note, however, that the elements of each basis need not be pairwise orthogonal.) Given an $\{F_\lambda\}$, the $\{D_\lambda\}$ satisfying these conditions are unique, so to specify a representation it suffices to determine the $\{F_\lambda\}$, as we will now do.

[30] We believe, but have not shown, that distinct GHW representations differ by exactly these choices of outcomes.

Consider in particular the two stabilizer measurements corresponding to the $X^\dagger$ and Z operators. If we label the outcome of $X^\dagger$ by $p \in \mathbb{Z}_d$ and the outcome of Z by $q \in \mathbb{Z}_d$, then by outcome determinism (Lemma B.3), each ontic state corresponds to an ordered pair (p, q). In fact, this correspondence is bijective, and hence we can choose a useful labelling of the ontic states, i.e. $\lambda \mapsto (p, q)$ (so that measurements of $X^\dagger$ reveal p and measurements of Z reveal q). To see that the correspondence is surjective, consider an eigenstate of X with eigenvalue ω^{-p_1}. The ontic states in the support of its representation must have $p = p_1$ so that the outcome of an $X^\dagger$ measurement is always p_1. Furthermore, a measurement of Z on this eigenstate gives a uniformly random outcome q, and so the ontic states in the support of its representation must include *every* ontic state of the form (p_1, q), for arbitrary $q \in \mathbb{Z}_d$. This holds for all d eigenstates of X, and thus for all $p_1 \in \mathbb{Z}_d$. So for every pair p, q, there exists some ontic state (in the support of one of the eigenstates of X) which has (p, q) as its label. This establishes surjectivity. Since the number (d^2) of ontic states is the same as the number of pairs (p, q), surjectivity implies bijectivity.

Next, we show that the assumed labelling forces the representation to manifestly satisfy translational covariance: that is, the Weyl unitaries must be represented in a Clifford covariant manner, so that the unitary superoperator $\mathcal{W}_{p_1,q_1}(\cdot) := Z^{p_1} X^{q_1} (\cdot) \left(X^{q_1}\right)^\dagger \left(Z^{p_1}\right)^\dagger$ is represented by the permutation $(p, q) \to (p + p_1, q + q_1)$. To see this, first recall that the representation of a unitary superoperator is necessarily a permutation, as shown in Lemma B.2. Next, we determine the representation of the unitary superoperator $\mathcal{X}(\cdot) := X(\cdot)X^\dagger$. Consider an eigenstate of X with eigenvalue ω^{-p_1}. We argued above that the ontic states in the support of its representation must have $p = p_1$. Because the state is invariant under the unitary superoperator $\mathcal{X}$, the value of p must be unchanged by it. Similarly, consider an eigenstate of Z with eigenvalue ω^{q_1}. The ontic states in the support of its representation must have $q = q_1$. Applying the unitary superoperator $\mathcal{X}$ increments the Z eigenstate and corresponding eigenvalue by one, so that the value of q is transformed to $q_1 + 1$. Hence, we see that the representation of the unitary superoperator $\mathcal{X}$ takes $p \to p$ and $q \to q + 1$, which fully specifies its action as a permutation on the ontic states. (Note that this argument only holds for ontic states in the support of one of the X eigenstates and also in the support of one of the Z eigenstates. But by Lemma B.3, every ontic state is of this sort.) By a similar argument, the representation of the unitary superoperator $\mathcal{Z}(\cdot) := Z(\cdot)Z^\dagger$ takes $p \to p + 1$ and $q \to q$. Since all Weyl unitary superoperators can be generated by composing $\mathcal{X}$ and $\mathcal{Z}$, and since the representation is diagram-preserving, this fully specifies the permutations representing all of the Weyl unitary superoperators. In particular, the unitary superoperator $\mathcal{W}_{p_1,q_1}$ is indeed represented by the permutation $(p, q) \mapsto (p + p_1, q + q_1)$.

By a similar argument, we can deduce the representation of the Hadamard unitary superoperator $\mathcal{H}(\cdot) := H(\cdot)H^\dagger$, where H is defined in Eq. (6.270). In particular, if we start in the eigenstate of X with eigenvalue ω^{p_1}, then $p = -p_1$, and the Hadamard maps this to the eigenstate of Z with eigenvalue ω^{p_1}, for which $q = p_1$. So we see that the permutation representing the Hadamard superoperator results in a final value for q equal to the initial value of $-p$. Similarly, for the eigenstate of Z with eigenvalue ω^{q_1}, one has $q = q_1$, and this is mapped to the eigenstate of X with eigenvalue ω^{-q_1}, for which $p = q_1$. So we see that the permutation representing the Hadamard superoperator also results in a final value for p equal to the initial value of q. This fully specifies its action as a permutation on the ontic states, namely $(p, q) \mapsto (q, -p)$. In particular, $(0, 0) \mapsto (0, 0)$.

Define the Weyl operators as

$$W_{p,q} = Z^p X^q. \tag{6.279}$$

Now, since these operators (or their conjugates) are a basis for the complex vector space of linear operators on the Hilbert space, we can decompose the operator $F_{0,0}$ (namely, the element of the basis $\{F_\lambda\}_\lambda$ with $\lambda = (0, 0)$) as

$$F_{0,0} = \frac{1}{d} \sum_{p,q} f_{p,q} W^\dagger_{p,q} \tag{6.280}$$

Consider a measurement of a given Weyl operator W_{p_1,q_1} when the ontic state happens to be $(0, 0)$. By outcome determinism, we will always get a particular outcome, which we will label v_{p_1,q_1}. W_{p_1,q_1} has a spectral decomposition $\sum_\alpha \omega^\alpha \Pi^{p_1,q_1}_\alpha$ in terms of its eigenvalues ω^α for $\alpha \in \mathbb{R}$ and the projectors Π^{p_1,q_1}_α onto the corresponding eigenvectors. Computing the quantity $\mathsf{tr}[F_{0,0}W_{p_1,q_1}]$, we obtain

$$\mathsf{tr}[F_{0,0}W_{p_1,q_1}] = \sum_\alpha \omega^\alpha \mathsf{tr}[F_{0,0}\Pi^{p_1,q_1}_\alpha]. \tag{6.281}$$

But we know that $\mathsf{tr}[F_{0,0}\Pi^{p_1,q_1}_\alpha]$ is the probability of outcome α occurring in a measurement of W_{p_1,q_1} when the ontic state is $(0, 0)$, and we have already defined that the outcome that must occur in this case is that corresponding to eigenvalue $\omega^{v_{p_1,q_1}}$. It follows that $\mathsf{tr}[F_{0,0}W_{p_1,q_1}] = \omega^{v_{p_1,q_1}}$.

But a substitution of Eq. (6.280) into the left-hand side of Eq. (6.281) also allows us to compute this value as

$$\mathsf{tr}[F_{0,0}W_{p_1,q_1}] = \frac{1}{d} \sum_{p,q} f_{p,q} \mathsf{tr}[W^\dagger_{p,q} W_{p_1,q_1}] = f_{p_1,q_1}, \tag{6.282}$$

where the last equality follows from Eq. (6.265). Hence $f_{p,q} = \omega^{v_{p,q}}$, and so

$$F_{0,0} = \frac{1}{d} \sum_{p,q} \omega^{v_{p,q}} W_{p,q}^{\dagger}. \tag{6.283}$$

Using $X^q Z^p = \omega^{-pq} Z^p X^q$ we can calculate

$$W_{p,q}^{\dagger} = X^{-q} Z^{-p} = \omega^{-pq} Z^{-p} X^{-q} = \omega^{-pq} W_{-p,-q}, \tag{6.284}$$

so that

$$W_{p,q} = \omega^{pq} W_{-p,-q}^{\dagger}. \tag{6.285}$$

We require $F_{0,0}$ to be Hermitian, i.e.

$$\frac{1}{d} \sum_{p,q} \omega^{v_{p,q}} W_{p,q}^{\dagger} = F_{0,0} = F_{0,0}^{\dagger} = \frac{1}{d} \sum_{p,q} \omega^{-v_{p,q}} W_{p,q} \tag{6.286}$$

and so by using Eq. (6.285) equating the phases in front of $W_{p,q}^{\dagger}$ this becomes

$$v_{p,q} = -v_{-p,-q} + pq. \tag{6.287}$$

For the Hadamard we have

$$H W_{p,q}^{\dagger} H^{\dagger} = H X^{-q} Z^{-p} H^{\dagger} = Z^{-q} X^{p} = \omega^{-pq} X^{p} Z^{-q} = \omega^{-pq} W_{q,-p}^{\dagger}. \tag{6.288}$$

Hence covariance $F_{0,0} = H F_{0,0} H^{\dagger}$ becomes

$$v_{p,q} = v_{-q,p} + pq. \tag{6.289}$$

Applying this twice gives

$$v_{p,q} = v_{-p,-q}. \tag{6.290}$$

Summing Eqs. (6.287) and (6.290) gives

$$2v_{p,q} = pq. \tag{6.291}$$

Now consider a pair of commuting Weyl operators $W_{p,q}$ and $W_{p',q'}$, where the requirement that they commute can be expressed as $pq' - qp' = 0$. They are jointly measurable, and give outcomes $v_{p,q}$ and $v_{p',q'}$ on the $(0,0)$ ontic state. Their product $W_{p,q}W_{p',q'}$ is also jointly measurable with both. It is a general feature of quantum theory that if measurements of some commuting A and B give eigenvalues a and b, a measurement of their product AB gives eigenvalue ab. Here we have $a = \omega^{v_{p,q}}$ and $b = \omega^{v_{p',q'}}$ so the outcome of $W_{p,q}W_{p'q'}$ on the $(0,0)$ ontic state must also be $v_{p,q} + v_{p',q'}$. But

$$\begin{aligned} W_{p,q}W_{p'q'} &= Z^p X^q Z^{p'} X^{q'} \\ &= \omega^{-p'q} Z^{p+p'} X^{q+q'} = \omega^{-p'q} W_{p+p',q+q'}. \end{aligned} \tag{6.292}$$

Since the outcome of $W_{p+p',q+q'}$ on $(0,0)$ is $v_{p+p',q+q'}$, this gives the outcome of $W_{p,q}W_{p'q'}$ as $v_{p+p',q+q'} - p'q$. But we already established that this outcome must be $v_{p,q} + v_{p',q'}$, so that

$$v_{p+p',q+q'} = v_{p,q} + v_{p',q'} + p'q \tag{6.293}$$

for all such (p, q, p', q').

In the special case when $p' = p$ and $q' = q$ the commutation condition is clearly satisfied, and hence

$$v_{2p,2q} = 2v_{p,q} + pq. \tag{6.294}$$

Then we can apply Eq. (6.291) to obtain

$$v_{2p,2q} = 2pq. \tag{6.295}$$

We now consider three cases, depending on the dimension d.

Odd d

In odd d we have $W_{p,q}^d = (Z^p X^q)^d = \mathbb{1}$ [153], so that the eigenvalues of $W_{p,q}$ are d-th roots of unity. Hence the $v_{p,q} \in \mathbb{Z}_d$. In odd d, $\mathbb{Z}_d$ contains a unique inverse of 2 so we can multiply each side of Eq. (6.291) by 2^{-1} to obtain the unique solution

$$v_{p,q} = 2^{-1}pq \tag{6.296}$$

Hence $F_{0,0} = \sum_{p,q} \left(\omega^{-2^{-1}pq} W_{p,q}\right)^\dagger$ is Gross's phase point operator.

Furthermore, we already argued that our representation must satisfy translation covariance, which is satisfied if and only if $F_{p,q} = W_{p,q} F_{0,0} W_{p,q}^\dagger$; since Gross's representation also satisfies translation covariance its $F_{p,q}$ are likewise. Hence, the set of basis operators $\{F_\lambda\}_\lambda = \{F_{p,q}\}_{p,q}$ is exactly equal to the set of phase point operators in Gross's representation.

Hence, any nonnegative and diagram-preserving quasiprobability representation for the stabilizer subtheory in odd dimensions is equivalent to Gross's.

Even d, not a multiple of 4

In even d, there are values of p, q for which $(Z^p X^q)^d \neq \mathbb{1}$ [153], so the above argument for $v_{p,q} \in \mathbb{Z}_d$ is not applicable. Hence we allow arbitrary $v_{p,q} \in \mathbb{R}$ in the following.

If d is even but not a multiple of 4 then we can write $d = 2h$ where h is odd and $h = -h \mod d$. If we set $p = q = h$ then the Hadamard covariance condition in Eq. (6.289) becomes

$$v_{h,h} = v_{h,h} + h^2 \tag{6.297}$$

so that $h^2 = 0$. But we have $h = 1 \mod 2$ and so $h^2 = h \mod 2h$.

d a multiple of 4

The remaining case is that d is a multiple of 4, i.e. $d = 4r$ for some non-zero r. If we set $(p, q) = (0, 2)$ and $(p', q') = (2r, 2(r-1))$ then $pq' - qp' = -4r = -d = 0$ so that we can apply Eq. (6.293) to obtain

$$v_{2r,2r} = v_{0,2} + v_{2r,2(r-1)}. \tag{6.298}$$

Applying Eq. (6.295) to each term this becomes

$$2r^2 = 0 + 2r(r-1), \tag{6.299}$$

so that $2r = 0$. But $2r \neq 0$, so there is no valid model in this d either. Together with the previous case this establishes there are no valid models in any even dimension.

Multipartite cases

The multipartite generalization of these results follows immediately from Proposition VI.6 of Ref. [260], which implies that the frame representation for processes on a pair of systems is uniquely determined by the frame representation for processes on each component system. In the case that the component systems are odd-dimensional, they each have a unique representation, and hence, so too does the composite system. In the case that the component systems are even-dimensional, they do not admit of any noncontextual representation, and hence, neither does the composite system.

B.4 Alternative arguments for the necessity of generalized contextuality

Recall from the main text

Theorem 19. *Consider any state ρ which promotes the stabilizer subtheory to universal quantum computation. There is no generalized noncontextual model for the stabilizer subtheory together with ρ.*

One might expect that this result follows immediately from the fact that there is no nonnegative quasiprobability representation of full quantum theory, and that such a proof would hold in every model of quantum computation. However, the mere fact that a universal quantum computer can *simulate* every quantum circuit does not necessarily imply that one can *implement* every quantum circuit. (The loophole here follows from the distinction between computational universality and strict universality [10]. For example, the Toffoli and Hadamard gate together form a computationally universal gate set, and yet composition of these two gates cannot generate arbitrary unitary gates—only those with real matrix elements.) Hence, one cannot without further arguments conclude that a universal quantum computer is capable of implementing circuits with negativity (or contextuality)—one can only conclude that it can simulate such circuits.

However, we believe that Theorem 19 could be proven in by leveraging the previous necessity result for Kochen-Specker contextuality [154] together with the fact that Kochen-Specker contextuality implies generalized contextuality [169]. Such an argument would not be entirely trivial, as the latter implication requires bringing auxiliary operational processes into the argument, and one must establish that all of these additional processes are within the stabilizer subtheory. However, this does seem to be the case. First, one can establish outcome determinism for ontic states in the support of the maximally mixed state following the logic of Ref. [281], but using only stabilizer preparations. One can then establish that every ontic state in the support of the given nonstabilizer state (from the state-dependent proof of Ref. [154]) is also in the support of the maximally mixed state, using the fact that there always exists a decomposition of the maximally mixed state into the given nonstabilizer state together with *only* stabilizer states.

C Appendices for Chapter 6

C.1 Related work

A number of previous works either inspired parts of our work, or would be interesting to relate to our work.

The basic diagrammatic notation underpinning this work can be traced back to the work of, for example, [156, 159], which used string diagrams to represent particular types of categories. See [271] for a clear survey of these notations, and see [217] for a graphical representation of tensors. The two-directional diagrams which we used here were inspired by Hardy's duotensor notation [138]. A seemingly related notation has also appeared in the context of double categories [208], and it would be interesting to see if there is a formal connection between these. The work of [110] (which was itself based on [204]) first introduced us to the graphical representation of diagram-preserving maps which we used in this work.

Diagrammatic notation was first used in the context of quantum theory within the research program of categorical quantum mechanics, which began in [77], was axiomatised in [5], and is now the basis of the textbook [82]. This sparked the quantum picturalism revolution [78, 137], as well as use of similar notation for GPTs [138] and the operational probabilistic theories of the Pavia group [69, 70, 92]. Stronger connections between these notations have been developed in, for example, [127, 269, 318]. Moreover, more categorical approaches to generalized theories have been studied extensively using diagrammatic notation, in particular by Gogioso in Refs. [124, 127, 140], which also contains a formal treatment of the infinite dimensional case [125, 126].

There are many connections to the framework of operational probabilistic theories [69, 70, 92], which served as inspiration for multiple aspects of our framework. Developing a full understanding of the relations between the two is left for future work. Of particular note is the idea of a prediction map being used to define a notion of equivalence, with respect to which one can quotient. This notion of quotienting also appeared in [138] and [260]. Moreover, the causality axiom of Ref. [69] is closely related to our ignorability assumption, and both of these are closely related to the notion of terminality of Refs. [79, 160].

Moreover, the rough idea of structure preservation in ontological models has appeared in various forms (e.g. as a diagram-preserving map, or equivalently as a functor between categories) in Refs. [4, 54, 121, 158, 193, 260].

C.2 On the meaning of diagrams in our operational CI theories

In Section 6.6.1, we noted that diagram preservation is an immediate consequence of our choice to take diagrams in an operational CI theory to represent one's hypothesis about the fundamental causal and inferential structure in the given scenario. We now contrast this with the usual approach to operational theories, wherein one typically takes operational diagrams to be a representation of some kind of structure that is independent of one's interpretation.

For example, in quantum theory, any given scenario can be described as a circuit of completely-positive trace-preserving maps. The circuit assigned to a particular experiment (or to the idealized conception thereof) is essentially unique, and is a fact on which physicists of virtually all interpretational camps will agree upon. At a minimum, these camps agree on this circuit as the 'correct' one in the sense of having maximal pragmatic utility as a mathematical representation of one's experiment. In the usual approach to operational theories, it is this circuit depicting the calculational structure that is typically taken as the diagram representing one's scenario.

To provide a realist representation of one's scenario, however, requires one (in our view) to furthermore commit to an underlying causal structure. In general (depending on one's interpretational camp), this causal structure will not correspond to the calculational circuit just described. Hence, in such an approach, one's realist representation map would not be diagram-preserving, but must somehow map from the calculational circuit to one's hypothesized causal structure.

In contrast, in our framework, we do not represent the calculational circuit at all. Rather, we stipulate that the diagram one draws to describe a given scenario in an operational CI theory must be chosen to respect one's hypothesis about the fundamental causal-inferential structure. Hence, the classical realist representation map is diagram-preserving.

The only real novelty here is that in our framework, the term 'operational theory' no longer describes a description which is so bare-bones that all users of the framework will agree on it.

We now note a key consequence of our choice to take operational diagrams to represent one's hypothesis about the fundamental causal-inferential structure: namely, that diagram-preservation does *not* constitute a limitation on the scope of realist representations within our framework.

To demonstrate this, consider the classical realist representation of a pair of independent causal systems in our framework. Diagram preservation implies that these are represented by a pair of independent systems in the classical realist CI theory. Since system composition in a classical realist CI theory is given by the cartesian product of the corresponding ontic state spaces, it appears as though our framework commits one to represent every pair of operational systems by a cartesian product of the corresponding state spaces, an assumption sometimes termed *ontic separability* [139]. If one is committed to the idea that the two systems in question fundamentally exhibit some holistic properties, then this assumption (and hence our assumption of diagram preservation) might appear overly restrictive. Such an impression is mistaken, however. In our framework, to posit such holistic properties is to grant that the actual causal situation is one in which the relevant degrees of freedom fundamentally *cannot* be divided into two independent subsystems—even if they are represented by a tensor product in the calculational diagram. Rather, they fundamentally behave as a single monolithic causal system. With this causal hypothesis, then, our framework demands that one represent the operational scenario using a single system rather than a pair of systems, and the classical realist representation of this single system is thereby allowed to be an arbitrary ontic state space, not necessarily a Cartesian product of ontic state spaces of two components. So we see that our framework does not limit the scope of classical realist representations.

Of course, given a commitment to a *particular* causal-inferential hypothesis, the assumption of diagram-preservation provides strong constraints on the scope of possible classical realist representations. These constraints take the form of causal compatibility constraints, as discussed in Section 6.7.1. Indeed, one can subsume a number of assumptions made in deriving no-go theorems on ontological representations (including those needed to derive Bell's theorem, a version of the preparation-independence postulate [230], the Markovianity assumption used in Ref. [207], lambda-screening [280], and the assumptions used in Ref. [260]) under the assumption that the fundamental causal-inferential structure respects the standard (calculational) quantum circuit.

C.3 Useful results in SubStoch

We now list a number of useful equalities, some of which we will need for proofs in the next section. Each can be verified immediately by composing the partial functions defining the relevant processes.

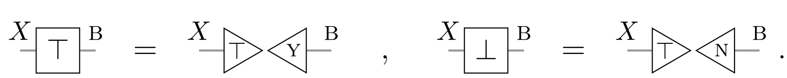

(6.300)

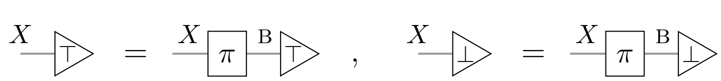

(6.301)

(6.302)

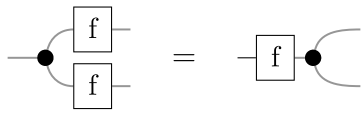

(6.303)

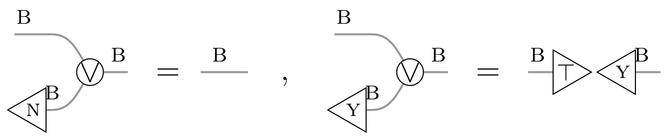

(6.304)

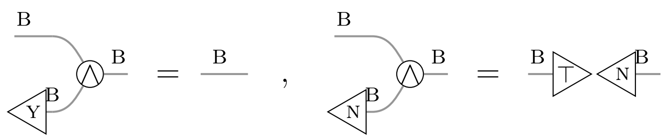

(6.305)

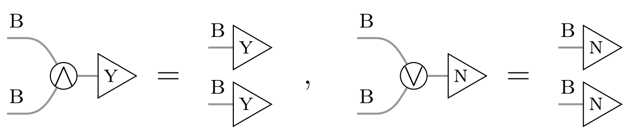

(6.306)

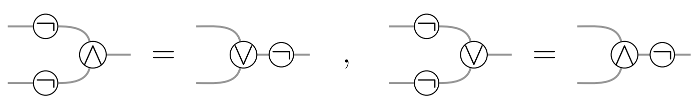

(6.307)

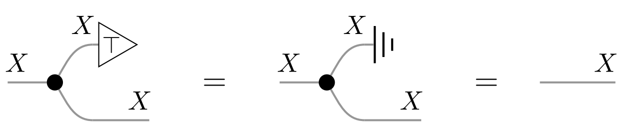

(6.308)

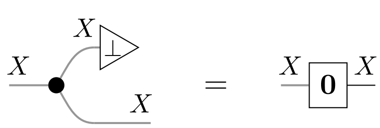

(6.309)

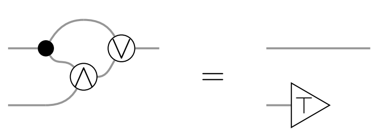

(6.310)

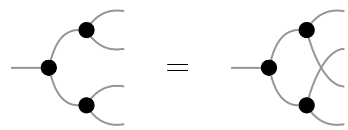

(6.311)

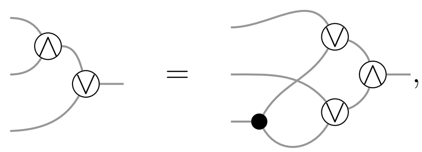

(6.312)

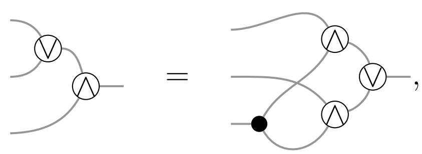

(6.313)

Properties of a Boolean algebra (proved diagrammatically)

A Boolean algebra satisfies the following properties (which are not all independent). For simplicity, we will here use α, β, and γ to denote propositions.

- associativity: $\alpha \vee (\beta \vee \gamma) = (\alpha \vee \beta) \vee \gamma$ and $\alpha \wedge (\beta \wedge \gamma) = (\alpha \wedge \beta) \wedge \gamma$
- commutativity: $\alpha \vee \beta = \beta \vee \alpha$ and $\alpha \wedge \beta = \beta \wedge \alpha$
- identity: $\alpha \vee \bot = \alpha$ and $\alpha \wedge \top = \alpha$
- complements: $\alpha \vee \neg\alpha = \top$ and $\alpha \wedge \neg\alpha = \bot$
- distributivity: $\alpha \vee (\beta \wedge \gamma) = (\alpha \vee \beta) \wedge (\alpha \vee \gamma)$ and $\alpha \wedge (\beta \vee \gamma) = (\alpha \wedge \beta) \vee (\alpha \wedge \gamma)$
- idempotence: $\alpha \vee \alpha = \alpha$ and $\alpha \wedge \alpha = \alpha$
- annihilation: $\alpha \vee \top = \top$ and $\alpha \wedge \bot = \bot$
- absorption: $\alpha \vee (\alpha \wedge \beta) = \alpha$ and $\alpha \wedge (\alpha \vee \beta) = \alpha$.

We now prove that each of these expressions holds in our diagrammatic representations. We prove only the first of each of these expressions; in each case, the proof of the second is similar. The associativity and commutativity axiom follow immediately by the symmetry of Eq. (6.37). The identity axiom follows simply from

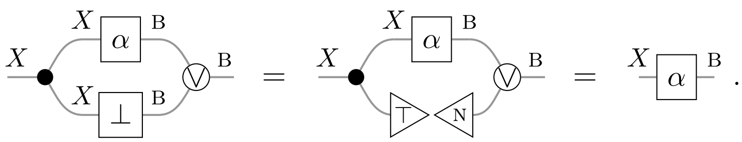

(6.314)

To prove the complements axiom, one has

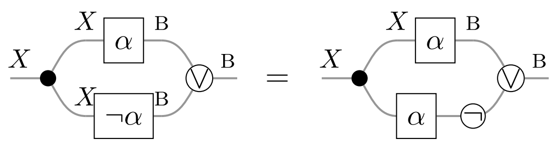

(6.315)

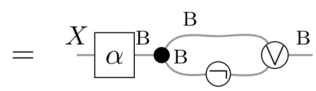

(6.316)

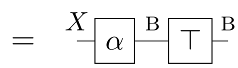

(6.317)

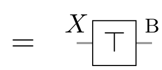

(6.318)

The proof of distributivity is as follows:

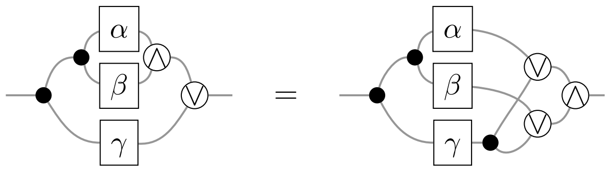

(6.319)

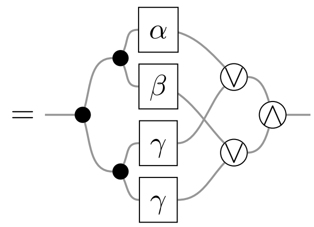

(6.320)

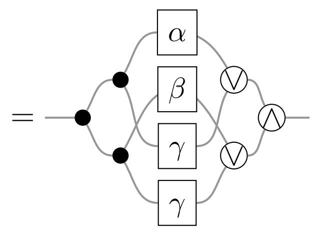

(6.321)

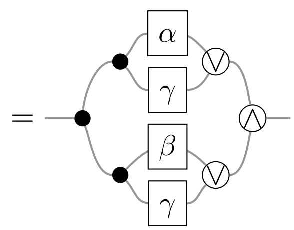

(6.322)

where Eq. (6.321) follows from Eq. (6.311) and Eq. (6.319) follows from Eq. (6.312).

The proof of idempotence is as follows:

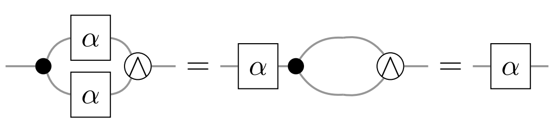

(6.323)

The proof of annihilation is as follows:

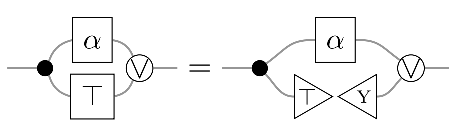

(6.324)

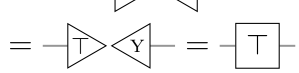

(6.325)

The proof of absorption is as follows:

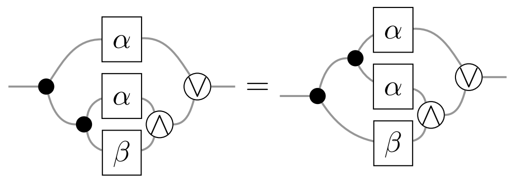

(6.326)

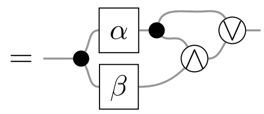

(6.327)

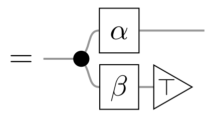

(6.328)

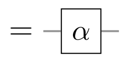

(6.329)

where Eq. (6.328) follows from Eq. (6.310).

Partial functions in Boole

Recall from Eq. (6.53) that any partial function $\hat{f}$ can be written as

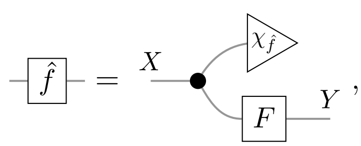

(6.330)

where $\chi_{\hat{f}}$ specifies the domain of $\hat{f}$ and F is a propositional map.

Consider now the action of $\hat{f}$ on an arbitrary propositional effect defined by $\pi \in \mathcal{B}(Y)$, namely

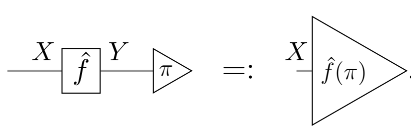

(6.331)

This defines a map from $\mathcal{B}(Y)$ to $\mathcal{B}(X)$, but it remains to see what structure of the Boolean algebra of propositional effects this map preserves. We now show that the action of partial functions on propositional effects preserves $\bot$, $\vee$ and $\wedge$, but not $\top$ and $\neg$.

First, note that we can reexpress $\hat{f}(\pi)$ in terms of the propositional effect $\chi_{\hat{f}} \in \mathcal{B}(X)$

and the total function F as follows:

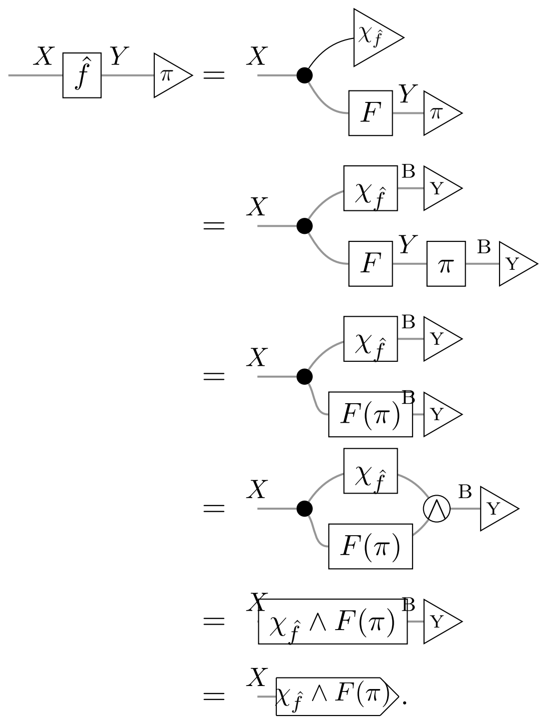

(6.332)

(6.333)

(6.334)

(6.335)

(6.336)

(6.337)

This is a very natural expression, stating that $\hat{f}(\pi)$ is equivalent to a propositional effect defined by the subset of X which is both in the domain of $\hat{f}$ and in the image of π under F.

At this point it is easy to verify that $\perp$ is preserved:

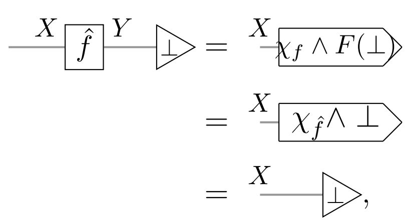

(6.338)

(6.339)

(6.340)

but that $\top$ is not preserved if $\chi_f \subsetneq X$:

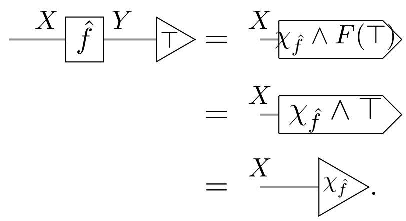

(6.341)

(6.342)

(6.343)

Hence, $\hat{f}$ does not define a Boolean algebra homomorphism (as these preserve $\top$).

However, it does preserve $\vee$ and $\wedge$, as we now show. Preservation of $\vee$ can be derived as

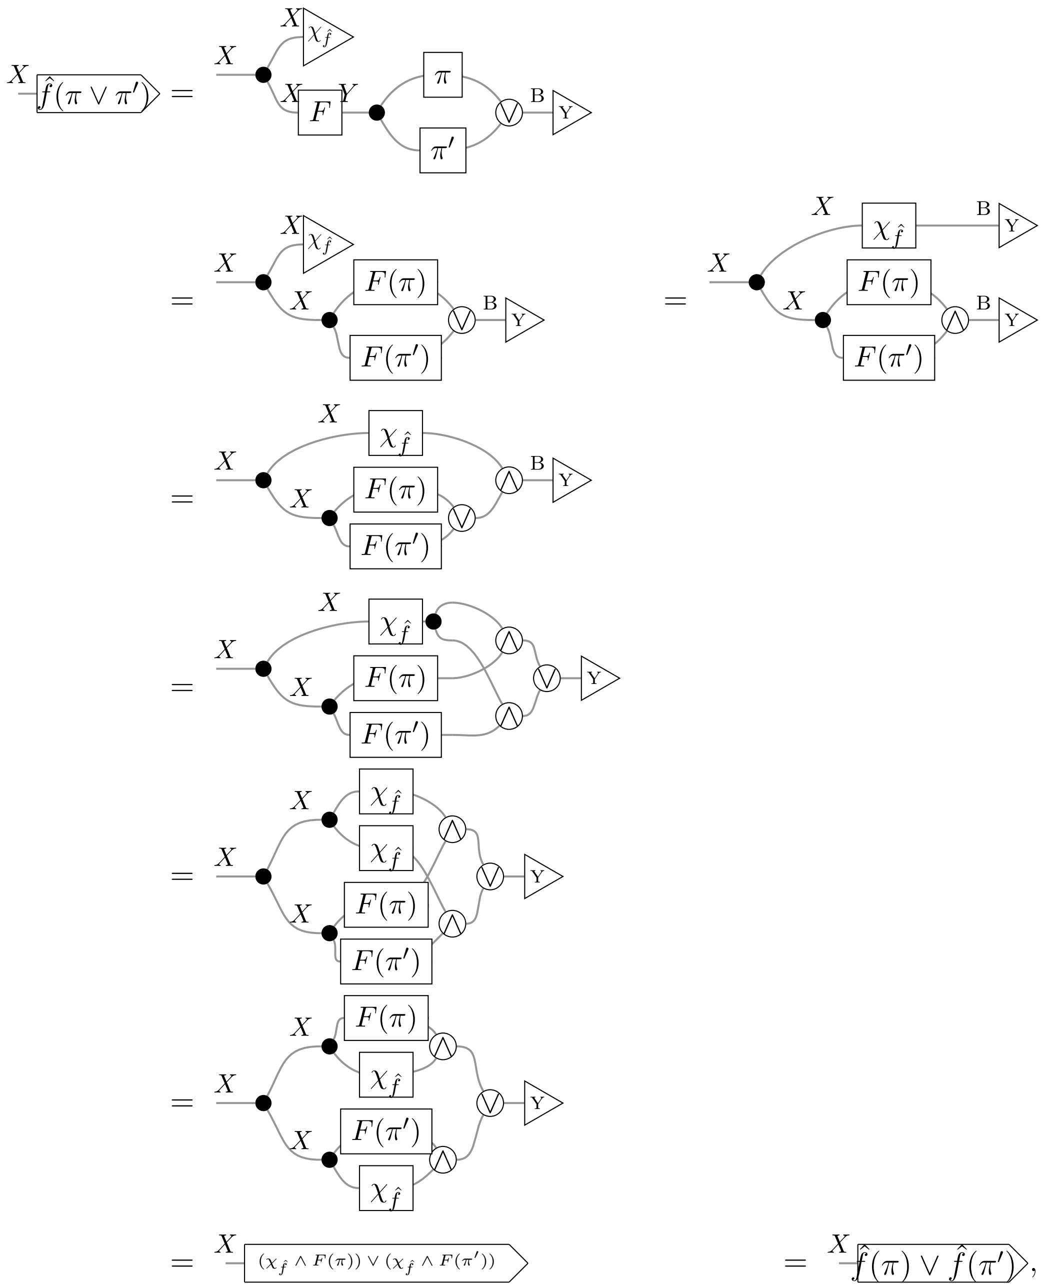

(6.344)

(6.345)

(6.346)

(6.347)

(6.348)

(6.349)

(6.350)

where Eq. (6.347) follows from Eq. (6.313) and Eq. (6.349) follows from Eq. (6.311).

Preservation of $\wedge$ can be derived as

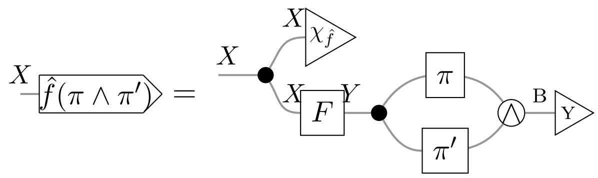

(6.351)

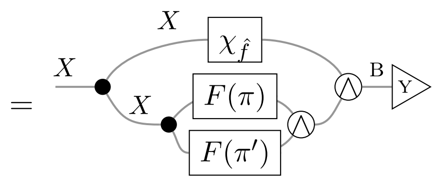

(6.352)

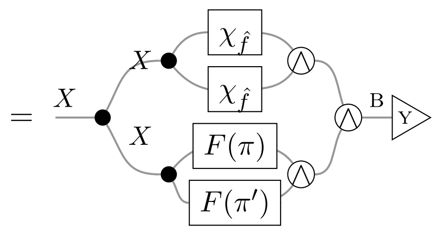

(6.353)

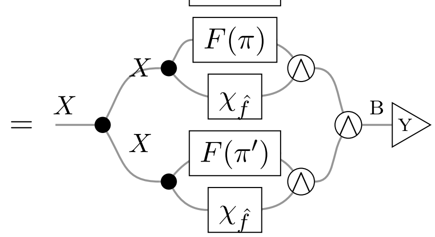

(6.354)

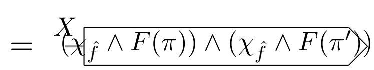

(6.355)

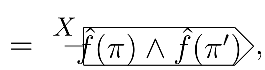

, (6.356)

where Eq. (6.353) follows from Eq. (6.319) and Eq. (6.354) follows from Eq. (6.311).

In summary, we see that $\hat{f}$ is a Boolean algebra homomorphism from $\mathcal{B}(Y)$ to $\mathcal{B}(\chi_{\hat{f}})$.

Propositions about composite systems

In Eq. (6.38), we noted that one can express propositional questions about composite systems, e.g. as

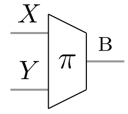

(6.357)

where $\pi \in \mathcal{B}(X \times Y)$. However, suppose that we have some propositional question π about X and some propositional question π' about Y; then, how should these be composed to give a propositional question about $X \times Y$? One's first guess might be to simply compose these in parallel within Boole; however, this would give

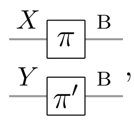

(6.358)

which is not a propositional question as its a function to B × B rather than simply B. The resolution comes from examining how we expect these to compose as Boolean Algebras. Note that the sets compose via the cartesian product, hence Boolean Algebras compose via the following rule: $\mathcal{B}(X) \otimes \mathcal{B}(Y) := \mathcal{B}(X \times Y)$ and moreover that $\pi \otimes \pi' \in \mathcal{B}(X) \otimes \mathcal{B}(Y)$ can be defined as a subset of $X \times Y$ by $\pi \times \pi'$ viewed as subsets of X and Y respectively. In our diagrammatic language this is represented by:

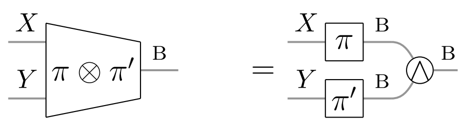

(6.359)

This composite therefore can be clearly interpreted as the situation in which we are interested in both π and π' being true about their respective systems.

Note, however, that we then clearly have other ways that we could compose these, for instance via:

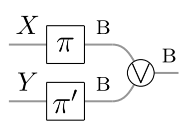

(6.360)

to see what this means in terms of the Boolean Algebra consider the following rewrites:

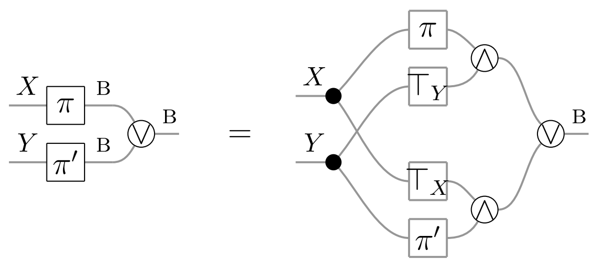

(6.361)

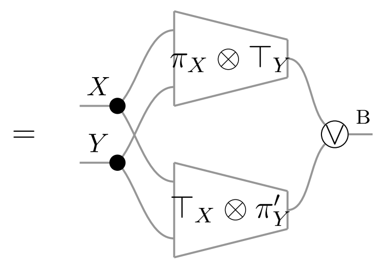

(6.362)

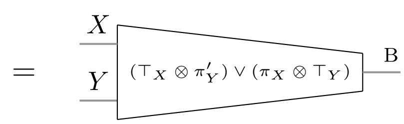

(6.363)

(Here we have included or omitted subscripts labeling the systems about which propositions are being made, as convenient.)

This corresponds, intuitively, to what we would mean for the logical disjunction of two propositions about distinct systems.

Useful relations between stochastic maps

There are various relationships between the stochastic maps which we have defined which are useful in the proofs in this paper. We list them here for reference.

(6.364)

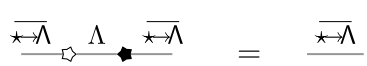

(6.365)

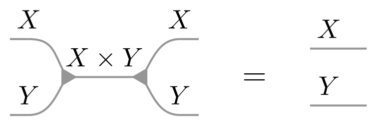

(6.366)

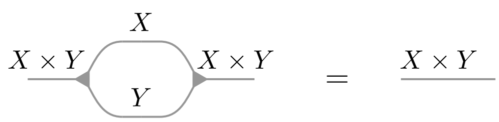

(6.367)

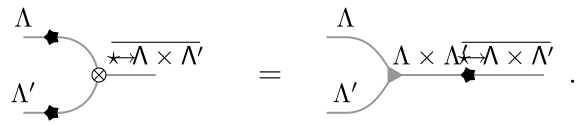

(6.368)

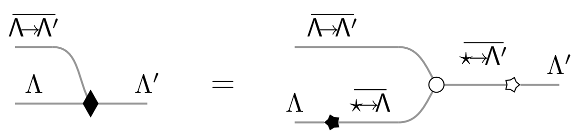

(6.369)

C.4 Useful results in FI

A useful lemma about (sub)stochastic maps

We now state and prove a lemma which was needed to justify Eq. (6.487).

Lemma C.1.

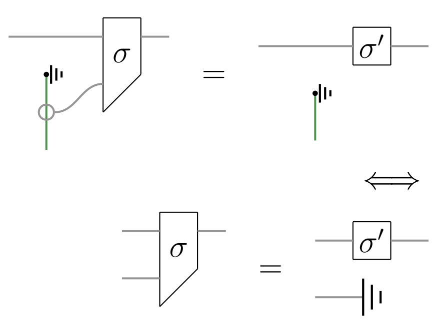

where σ is an arbitrary substochastic map. This result and proof also hold if one replaces all of the equalities with inferential equivalences.

Proof. The ⇐ direction trivially follows from Eq. (6.89). To prove the ⇒ direction, we begin by assuming that

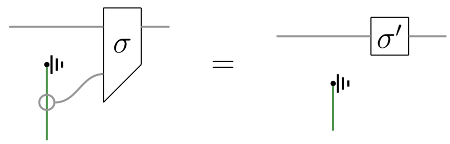

(6.370)

Composing this with a state preparation generator and the star isomorphism gives

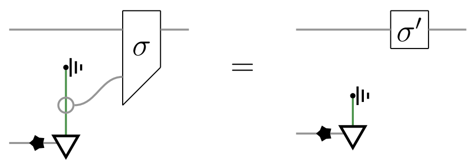

(6.371)

which can be rewritten using Eq. (6.142) to

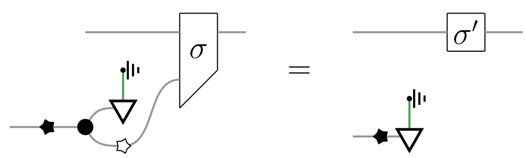

(6.372)

and then, using Eq. (6.189) and Eq. (6.308), to

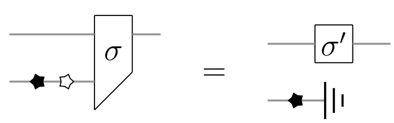

(6.373)

Finally, we use (on the LHS) the fact that the star is an isomorphism and (on the RHS) the fact that the star is stochastic to obtain the result:

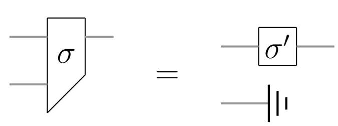

(6.374)

The proof proceeds in the same way if one replaces all equalities with inferential equivalences. If one assumes Eq. (6.370) but with the equality replaced by inferential equivalence, then Eq. (6.371) follows by the fact that inferential equivalence is preserved by composition. The remainder of the proof then follows by the same rewrite rules.

□

Copy function and complete common causes

First, let us define the *copy function* in FUNC, denoted

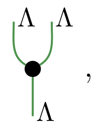

(6.375)

by $\bullet(\lambda) = (\lambda, \lambda)$ for all $\lambda \in \Lambda$.

We now show how some useful properties of the copy functioncan be lifted to define a corresponding process in F-S, via

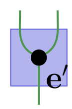

(6.376)

Specifically, that this acts as a suitable copy operation for F-S, that is, it is symmetric

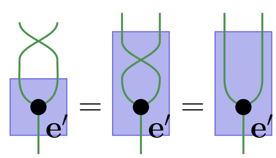

(6.377)

and associative,

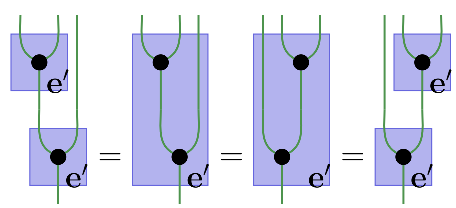

(6.378)

as follows immediately from diagram preservation of $\mathbf{e}'$ and the associativity of the underlying function. By Eq. (6.97), the embedding of the unique function u to the trivial system $\star$ is equal to the ignoring map:

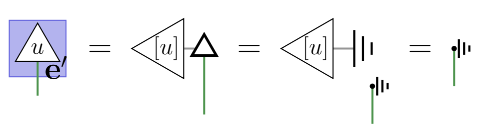

(6.379)

one can also see that it is a counit for the copy:

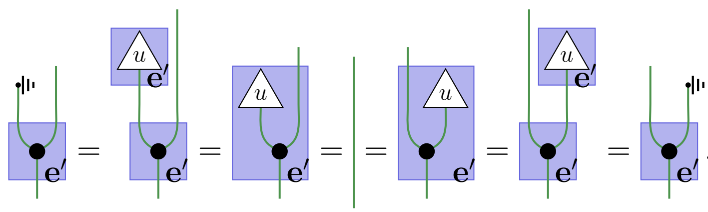

(6.380)

We now show that processes of the form

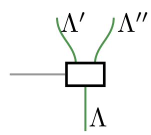

(6.381)

describe situations in which Λ is the *complete* common cause [12] of Λ′ and Λ″. This is a consequence of the fact that the outputs of a process in FUNC are (by construction) deterministic functions of the inputs, and hence in this diagram, Λ is the only possible cause of Λ′ and Λ″. (To represent a scenario in which the two have more than one common cause, one would represent these explicitly as inputs to the process. Note that our framework does not incorporate a diagrammatic distinction between latent and observed variables, although introducing such a distinction might be useful in future work.)

Lemma C.2.

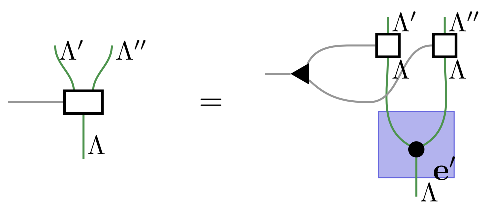

(6.382)

where the black triangle is a stochastic map defined by linearity and

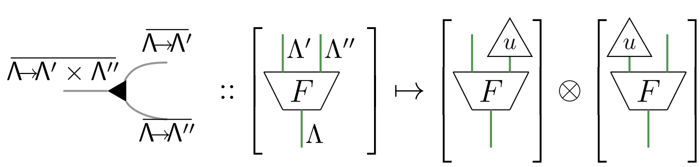

(6.383)

where u is the unique function to ⋆.

Proof. First, note that we have a similar decomposition in FUNC, that is for all functions F there exist f_l and f_r such that

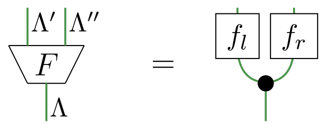

(6.384)

where

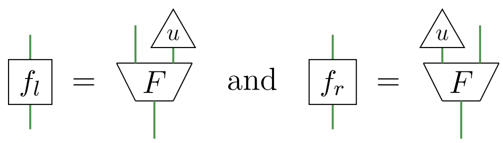

(6.385)

We will now show that this result can be lifted to F-S. Note that the definition of the black-triangle stochastic map, together with the above choices for f_r and f_l, imply that

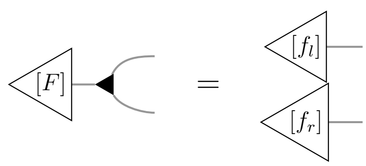

(6.386)

Now consider the following set of rewrites, where $[\bullet]$ is the state of certain knowledge that the copy operation, $\bullet : \Lambda \to \Lambda \times \Lambda$, of Eq. (6.375) has occurred.

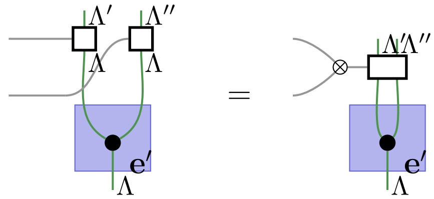

(6.387)

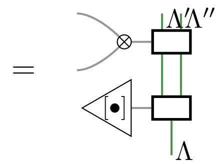

(6.388)

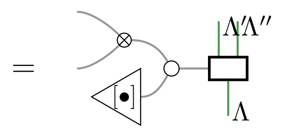

(6.389)

Next, note that

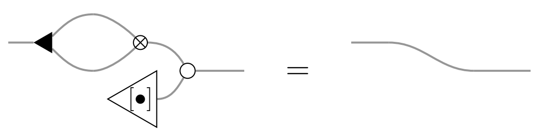

(6.390)

as can be verified by computing its action on an arbitrary delta function state of knowledge $[F]$, namely

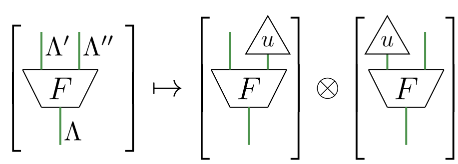

(6.391)

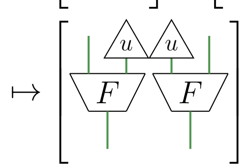

(6.392)

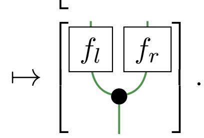

(6.393)

Hence, by Eq. (6.384), we see that

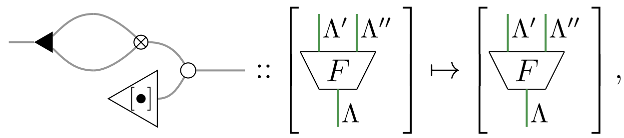

(6.394)

justifying Eq. (6.390). The conjunction of Eq. (6.390) and Eq. (6.389) immediately establishes the lemma. □

Next, we show that learning about an ontological system is the same as first copying that system and then learning about the copy:

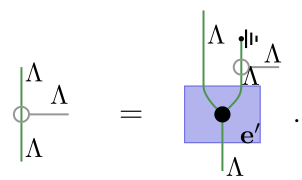

(6.395)

Proof. We start with the RHS and will rewrite it into the LHS. In the following equalities, Eq. (6.398) follows from Eq. (6.128), Eq. (6.399) follows from the fact that [•] is a point distribution, Eq. (6.401) follows from Eq. (6.308), and Eq. (6.402) follows from Eq. (6.379).

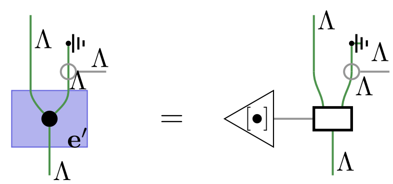

(6.396)

= 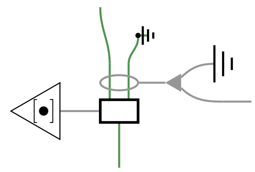 (6.397)

= 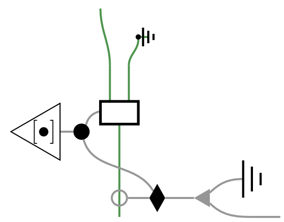 (6.398)

= 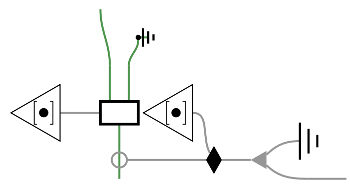 (6.399)

= 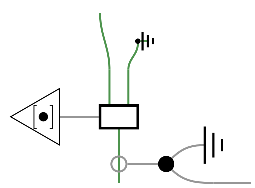 (6.400)

= 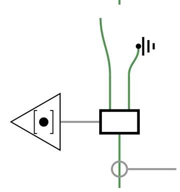 (6.401)

= 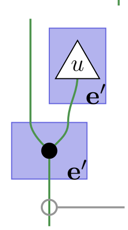

(6.402)

= 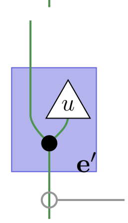

(6.403)

= 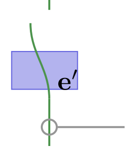

(6.404)

= 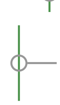 (6.405)

□

Other equalities

A special case of Eq. (6.128) is

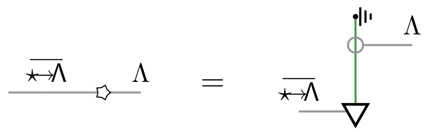

(6.406)

since

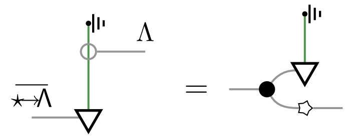

(6.407)

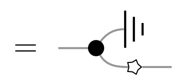

(6.408)

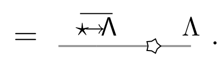

(6.409)

First, we show that one can always find at least one possible causal explanation in FI for any inference. A simple example of this is

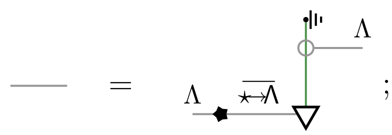

(6.410)

here, the inference described by the identity function is seen to have a possible causal explanation as the statement that a causal system has not evolved. As another simple example, inferences described by functions can always arise from a causal system evolving under that function as its dynamics, e.g. as

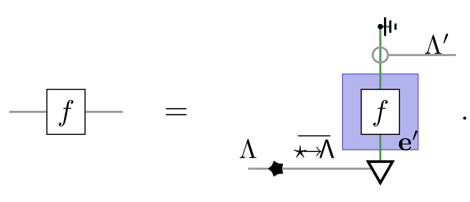

(6.411)

Most generally, an inference described by a general substochastic map can always be viewed as having partial knowledge about the input to some functional causal dynamics and considering a proposition about part of the output of the dynamics. That is, an arbitrary process $S \in$ SubStoch satisfies

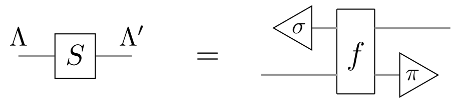

(6.412)

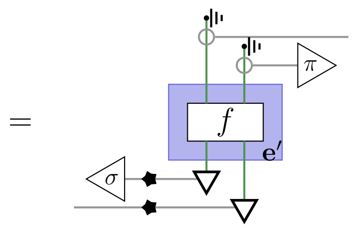

(6.413)

where σ is a probability distribution, f is a function, and π is a propositional effect.

Proof. The proof is as follows, where Eq. (6.420) follows from Eq. (6.368) and Eq. (6.421) follows from Eq. (6.369) (and the remaining equalities follow from the rewrite rules in F-S that we have introduced):

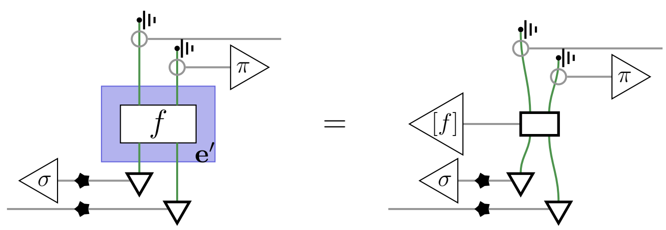

(6.414)

= 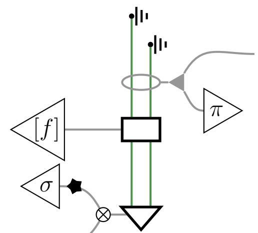

(6.415)

= 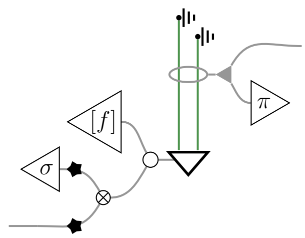

(6.416)

= 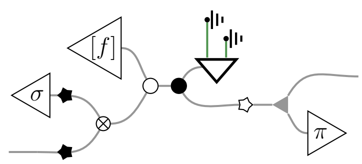

(6.417)

= 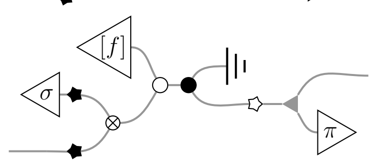

(6.418)

= 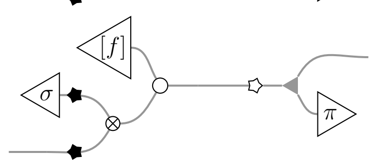

(6.419)

= 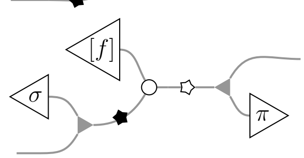

(6.420)

= 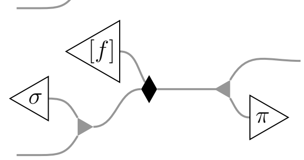

(6.421)

= 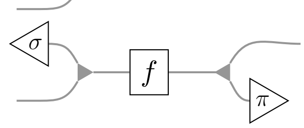

(6.422)

= 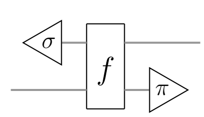

(6.423)

□

Next, we show that one can always replace what we know about a transformation with what we know about a variable that controls the transformation. First let us describe a "universal control" function ◆ ∈ FUNC, as follows:

$\forall f \in \overline{\Lambda \to \Lambda'}$

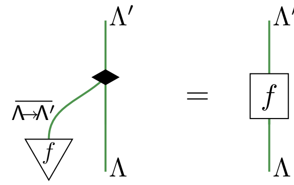

(6.424)

where the black diamond represents the universal control transformation ◆, and where we have introduced a causal system which ranges over the (finite) set of functions from Λ to Λ'. Then, one can show

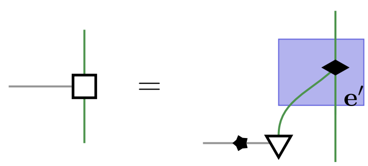

(6.425)

Proof. One can rewrite the RHS into the LHS, as follows:

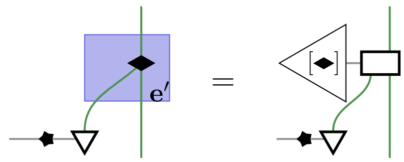

(6.426)

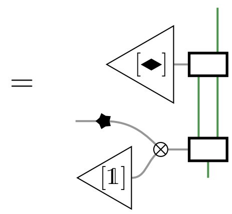

(6.427)

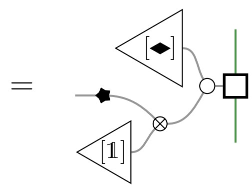

(6.428)

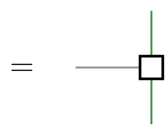

(6.429)

(6.430)

where the last step follows from the fact that

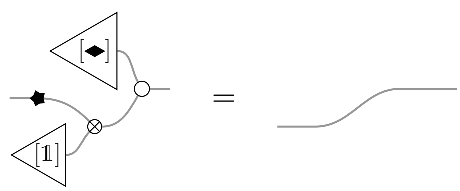

(6.431)

as can be verified by its action on an arbitrary delta function state of knowledge $[f]$. Namely,

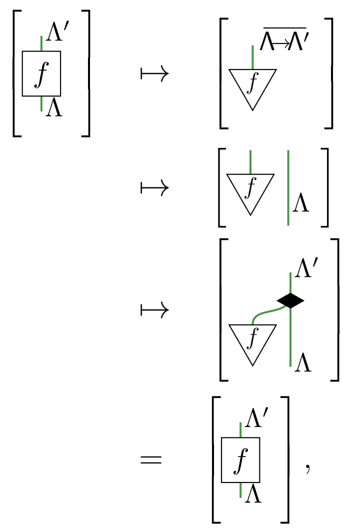

(6.432)

(6.433)

(6.434)

(6.435)

where the final equality is given by Eq. (6.425). □

Proof of normal form for F-S

We now prove Theorem 6.4.2; namely, the normal form

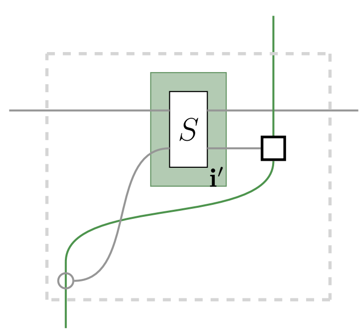

(6.436)

for F-S, where S is a substochastic map.

Proof. We will prove this by induction. First, we show (Step i) that every generator can be written into normal form. Second, we prove (Step ii) that the composite of two normal form diagrams can be rewritten into normal form. Given these, it is clear that one can write any diagram into normal form by first rewriting all of the generators involved into normal form using Step i, composing these according to Step ii, and iterating until the entire diagram is in normal form.

Step i—The fact that each generator is in normal form can be seen by inspection. For example, stochastic maps are generators in our theory, and are already in normal form—namely, the special case that arises when one takes all the causal systems in the normal form to be trivial. The other three generators (describing interactions between the causal and inferential systems) can be written in normal form as follows:

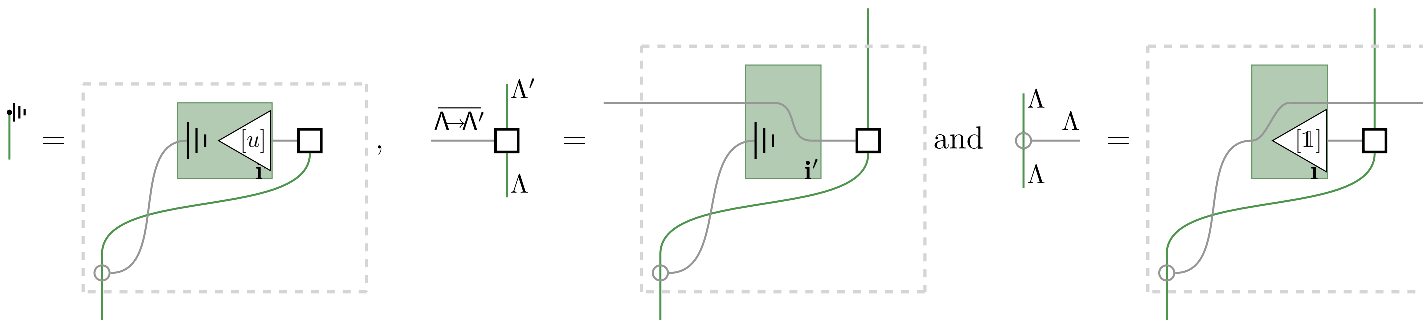

(6.437)

Step ii—First, we write down the most general way to compose two diagrams and then expand each of these in terms of the conjectured normal form:

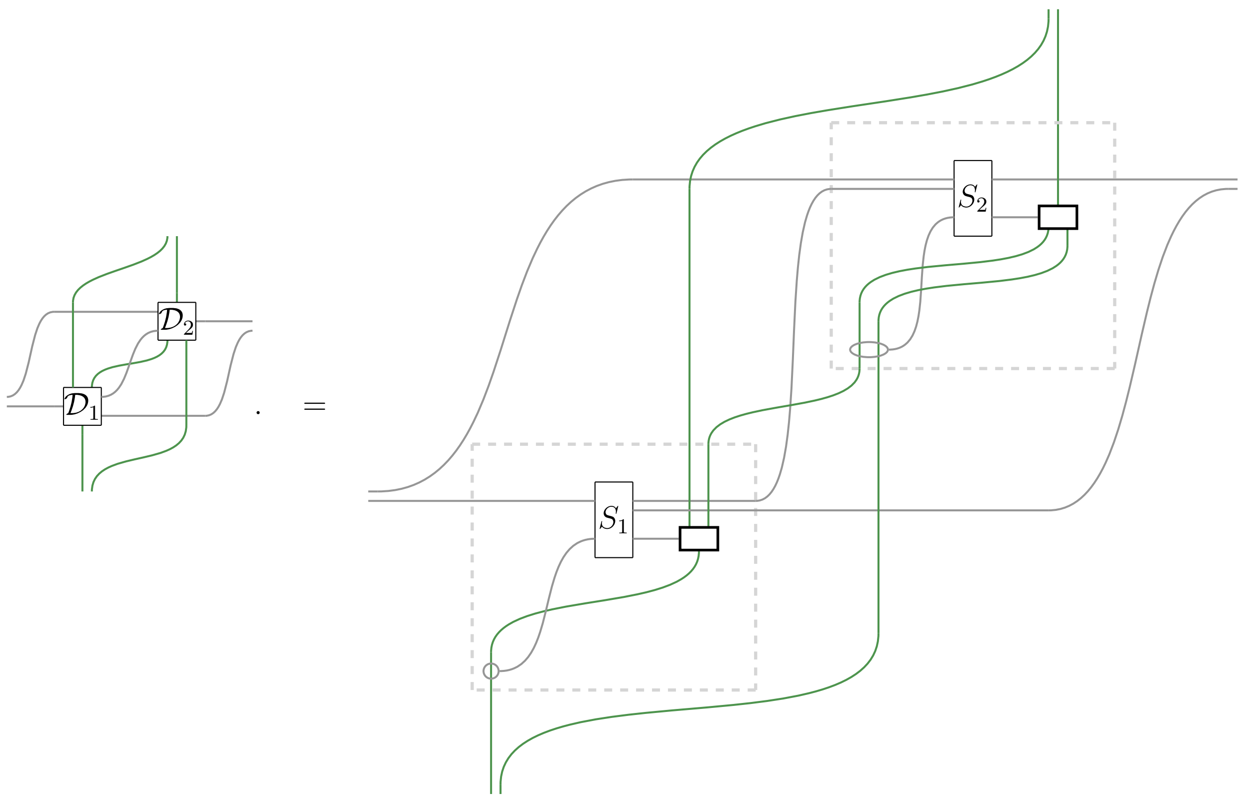

(6.438)

Removing the dashed gray lines and simply moving the wires around gives

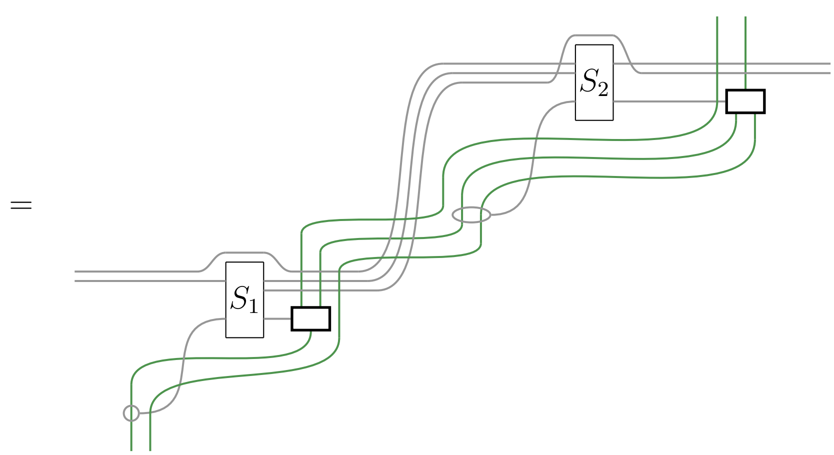

(6.439)

Next, we use Eq. (6.89) and Eq. (6.79) to add in extra processes to obtain

= 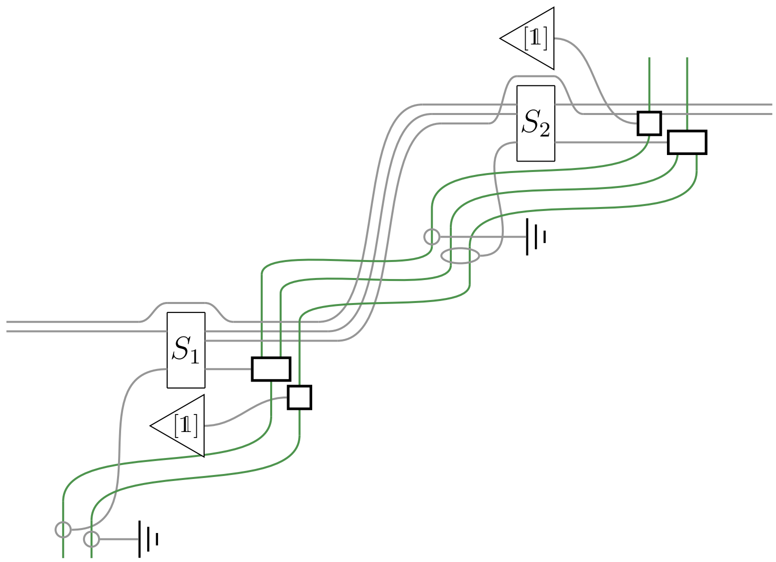 (6.440)

Merging some of these together, one obtains

= 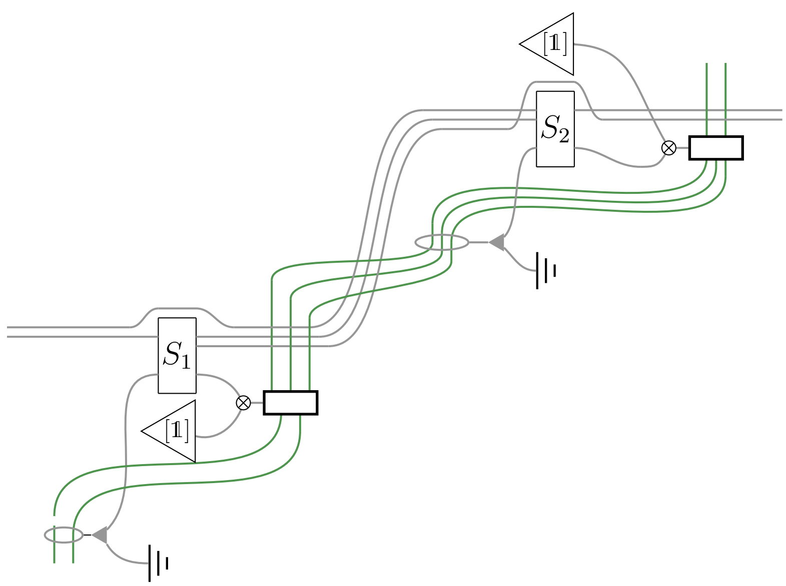 (6.441)

Next, simply moving wires around yields

$$= \quad$$ 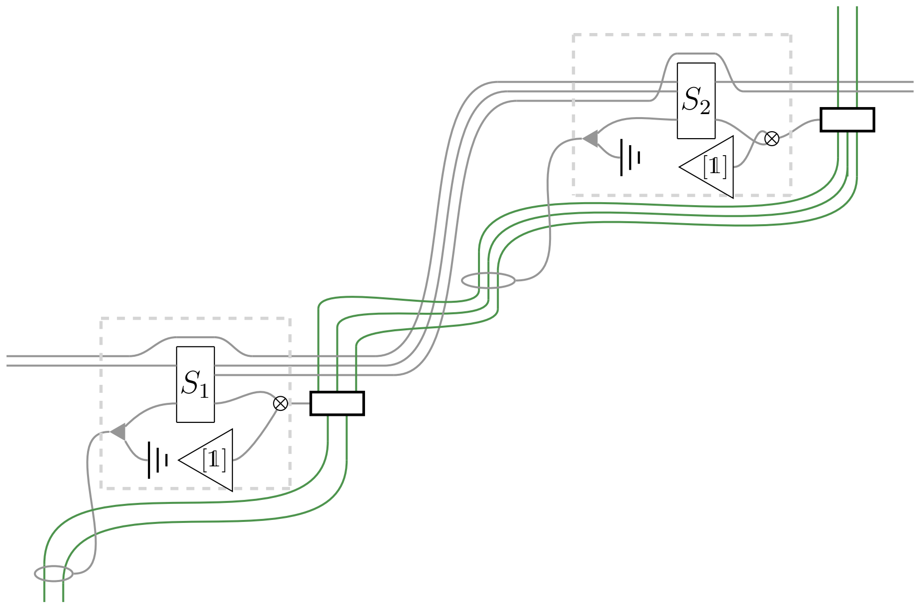 , (6.442)

at which point one can identify the two gray dashed boxes as stochastic maps (since these contain only normalized inferential processes). Denoting these S_1' and S_2' one obtains

$$=: \quad$$ 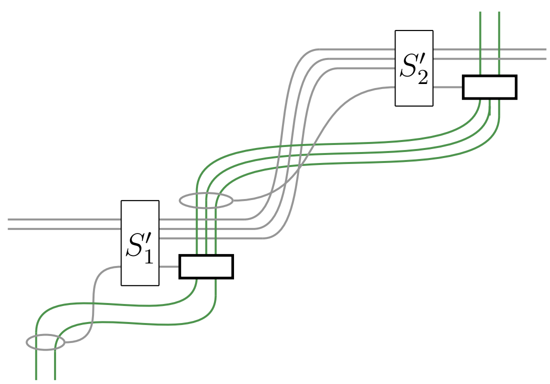 , (6.443)

We can use Eq. (6.128) to rewrite this as

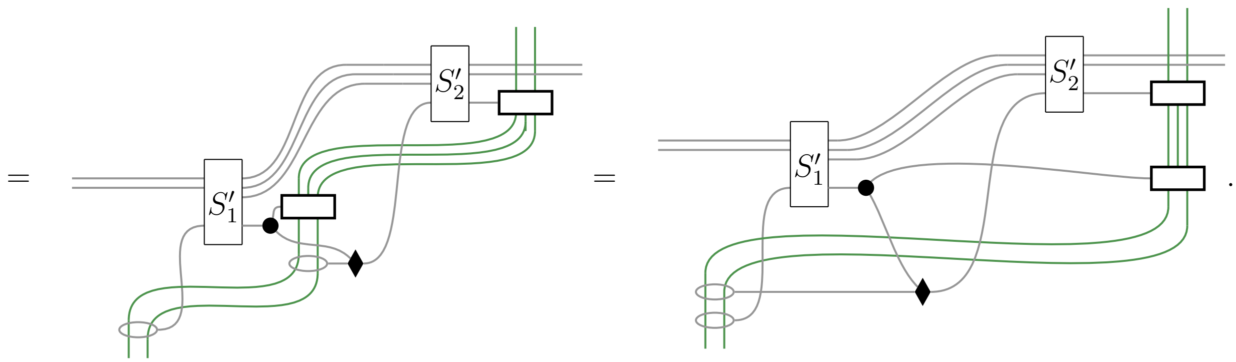

(6.444)

Rewriting to express compositional structure within one's inferences, one gets

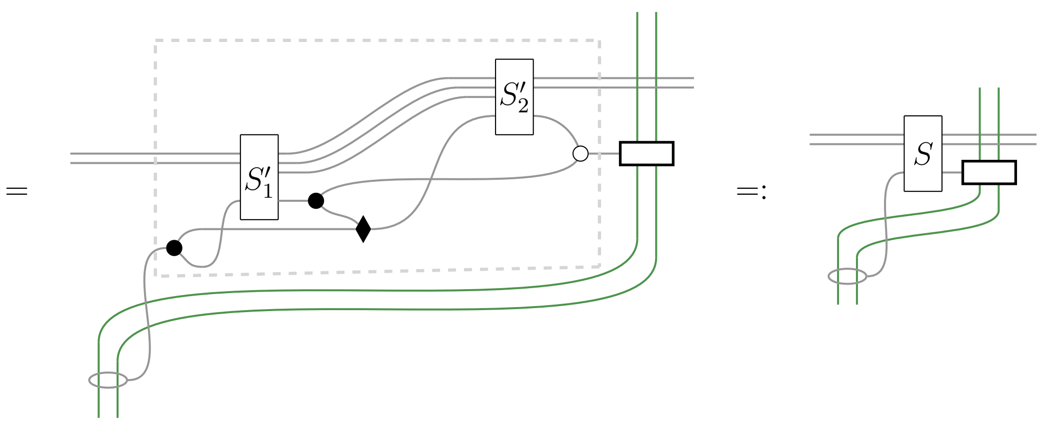

(6.445)

where one has identified the process in the dashed box as a stochastic map S. Noting then that each pair of wires can be considered as a single composite wire and that S is in the image of $\mathbf{i}'$, this is indeed seen to be in the claimed normal form. □

C.5 Useful results in $\widetilde{\mathbf{FS}}$

Statement and Proof of Lemma C.3

We now state and prove a lemma used in the main text.

Lemma C.3. *A causal identity is inferentially equivalent to a process which factors through* *an inferential system as*

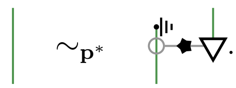

(6.446)

Proof. By Lemma 6.5.1, we can establish the inferential equivalence by showing the following:

$\forall \tau$

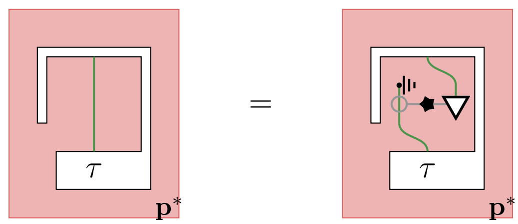

(6.447)

Now, as follows from Section 6.1.1, these testers are shorthand notation for a diagram of the form:

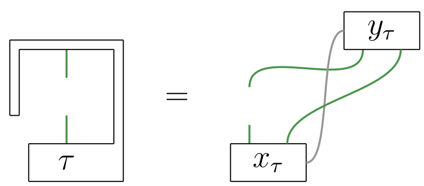

(6.448)

By applying the normal form of Theorem 6.4.2 to the special case of processes of the form of x_τ and of the form of y_τ, we can write this tester explicitly as

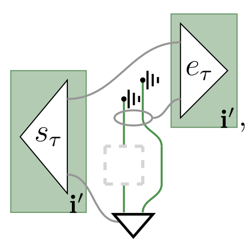

(6.449)

and hence our goal is to prove the following equality:

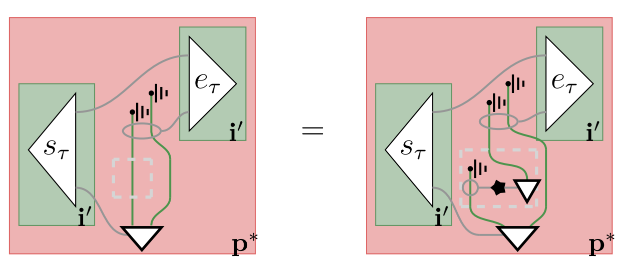

(6.450)

This equality follows immediately from the following set of rewrites (where the first and last equality follow from the special case of Lemma C.2 where Λ is trivial):

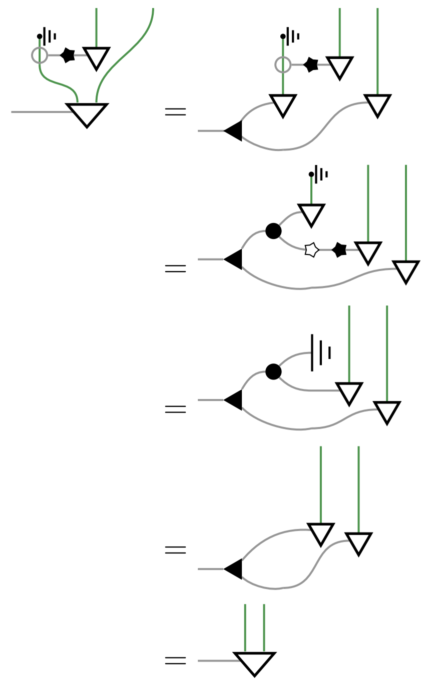

(6.451)

(6.452)

(6.453)

(6.454)

(6.455)

□

Proof of Lemma 6.5.6

We now prove Lemma 6.5.6, that two processes in $\widetilde{\mathrm{FS}}$ are inferentially equivalent if and only if they are associated with the same substochastic map.

Proof. The $\implies$ direction follows immediately from the definition of inferential equivalence and the fact that the following diagram is a tester:

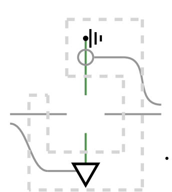

. (6.456)

To prove the $\Leftarrow$ direction, one can apply the fact that $\mathbf{p}^* = \tilde{\mathbf{p}}^* \circ \sim_{\mathbf{p}^*}$ and then apply

Lemma C.3 to show that

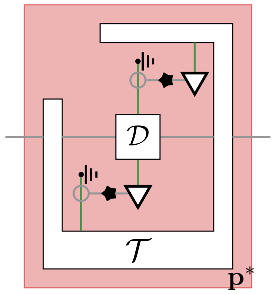

=

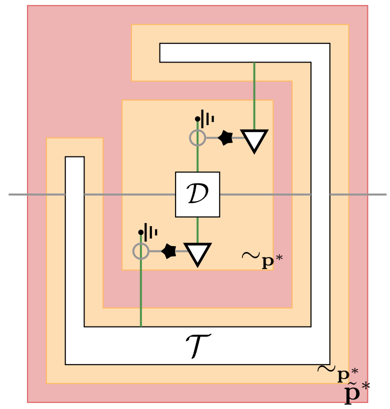

(6.457)

=

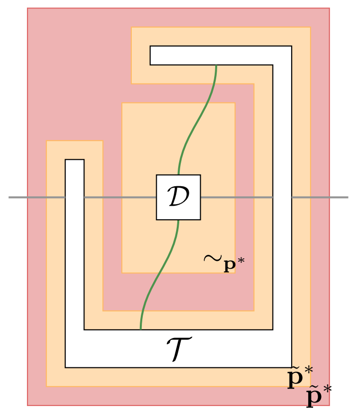

(6.458)

=

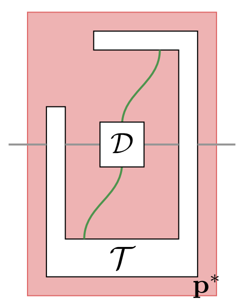

(6.459)

and similarly for $\mathcal{E}$.

Now, starting with the RHS of the implication in Eq. (6.177),

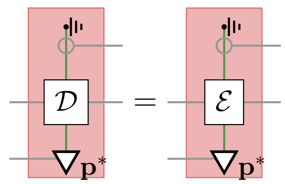

(6.460)

we will derive the LHS using the result we just proved. First, note that this equality implies that

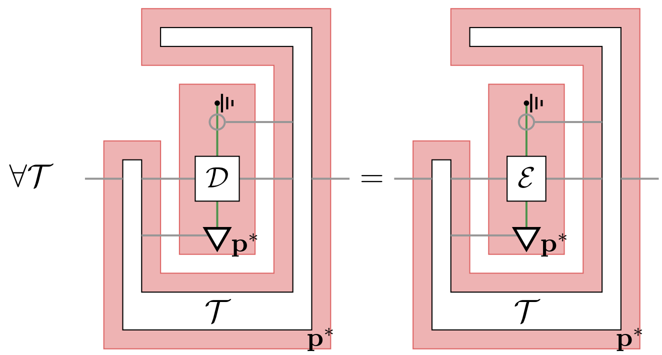

(6.461)

Diagram preservation then allows us to write this as:

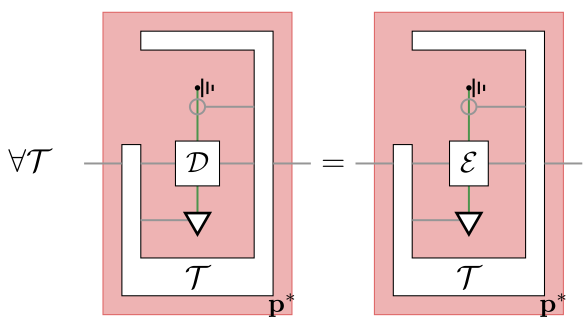

(6.462)

Now, consider a special class of testers, namely, those of the form:

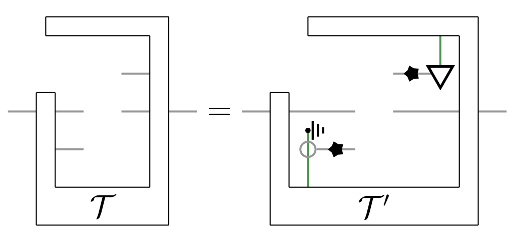

(6.463)

for any $\mathcal{T}'$. Condition 6.462 therefore implies that

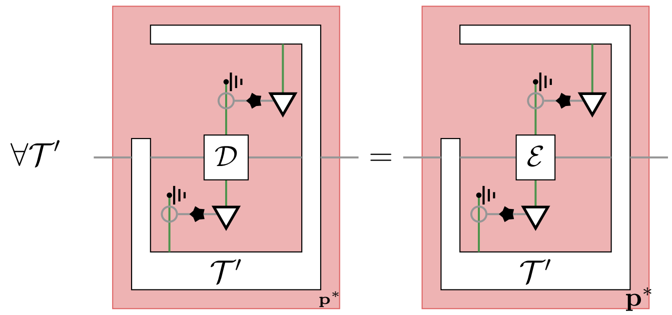

(6.464)

Using Eq. (6.459), this is equivalent to

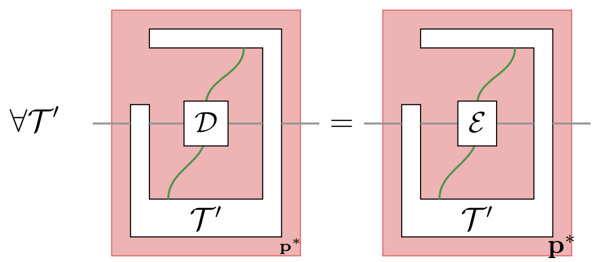

(6.465)

But this is just the definition of inferential equivalence:

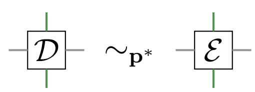

(6.466)

□

Proof of Theorem 6.5.7

We now prove Theorem 6.5.7, which immediately led to the normal form for $\widetilde{\text{F-S}}$ given in Corollary 6.5.7.1. In the equalities that follow, Eq. (6.468) follows from Eq. (6.413), Eq. (6.469) follows from Eq. (6.395) and Eq. (6.425), Eq. (6.471) follows from two applications of Eq. (6.446), Eq. (6.477) follows from Eq. (6.128), Eq. (6.478) follows from Lemma (C.2).

Proof.

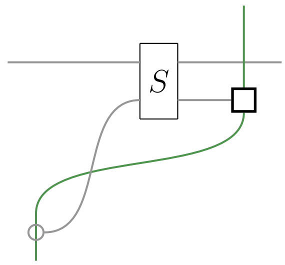

$=$

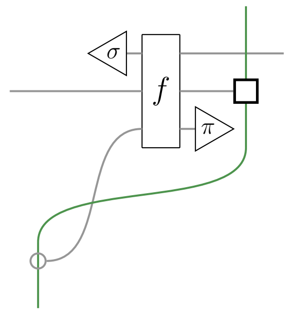

(6.467)

$=$

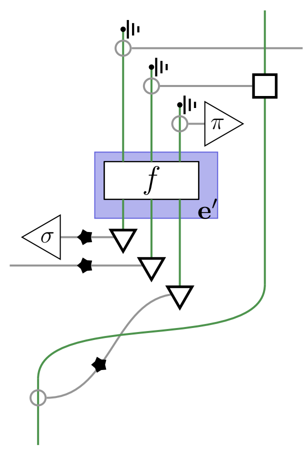

(6.468)

$=$

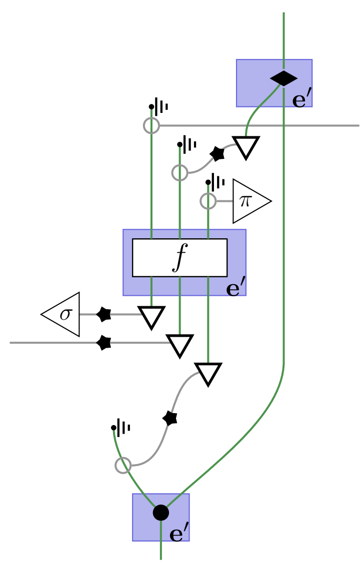

(6.469)

(6.470)

$\sim_{\mathbf{p}^*}$ 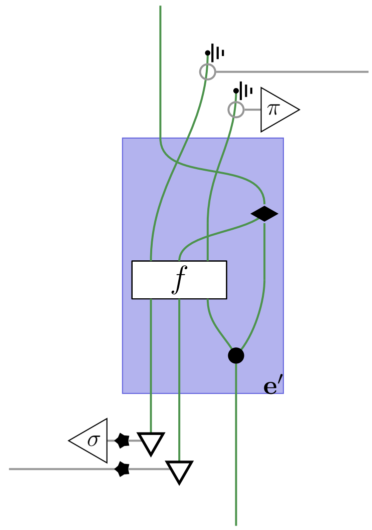 (6.471)

$=$ 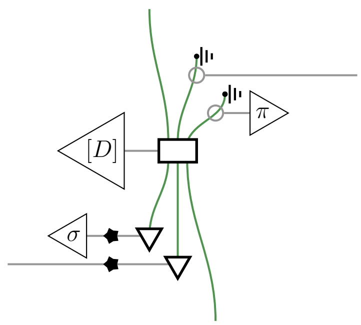 (6.472)

where

$D \quad =$ 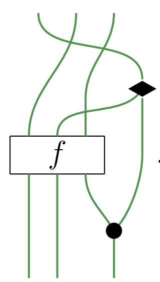 . (6.473)

We can then further rewrite this as:

= 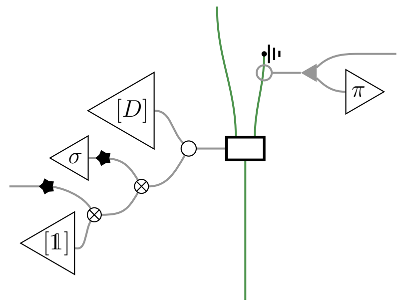

(6.474)

= 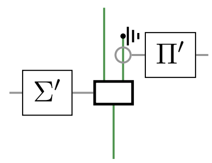

(6.475)

= 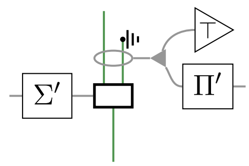

(6.476)

= 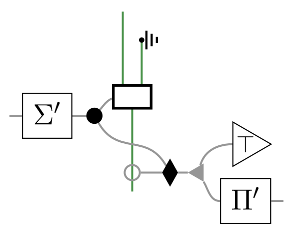

(6.477)

= 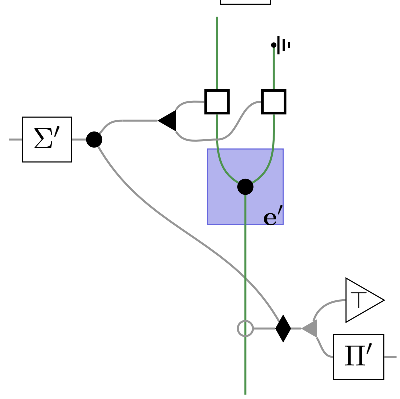

(6.478)

$$= \quad$$ 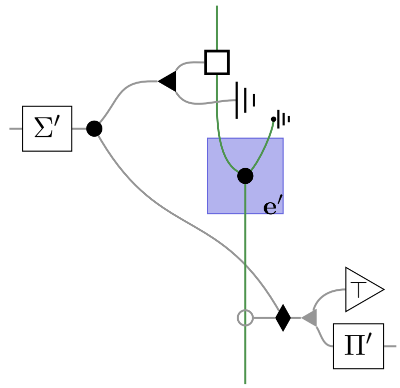 $$\qquad (6.479)$$

$$= \quad$$ 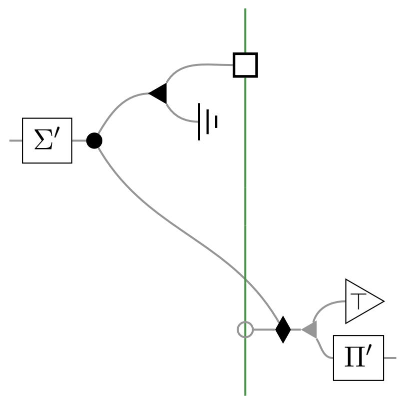 $$\qquad (6.480)$$

$$= \quad$$ 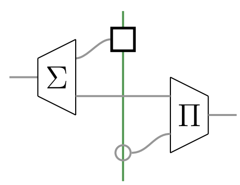 $$\qquad (6.481)$$

By their construction, one can see that Σ is a stochastic map and Π is a propositional map. □

C.6 Useful results for classical realist representations

Proof of Theorem 6.6.1

We now prove Theorem 6.6.1.

Proof. First, note that Eq. (6.189), diagram preservation of ξ, and the constraint of ignorability, Eq. (6.96), imply that

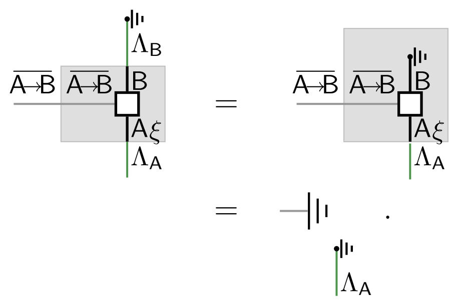

(6.482)

(6.483)

Now, Theorem 6.5.7 gives that

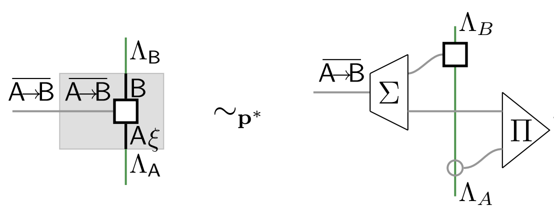

(6.484)

for some substochastic map Σ and some propositional effect Π. Applying this to decompose the process on the LHS of Eq. (6.482), one gets

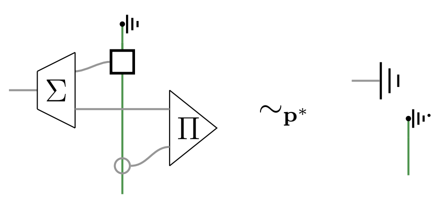

(6.485)

Rewriting the LHS of this we obtain

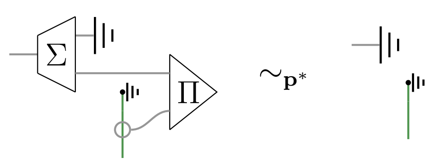

(6.486)

Using Lemma C.1 (stated and proved in Appendix C.4), we obtain

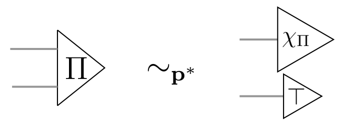

(6.487)

Substituting this in, we obtain

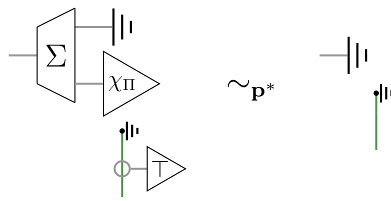

(6.488)

and so

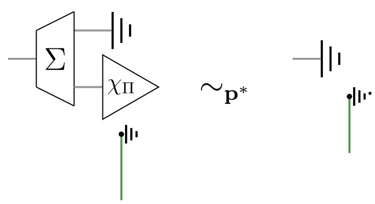

(6.489)

Hence, it must be that

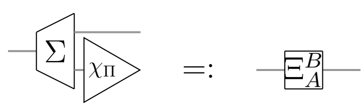

(6.490)

is a stochastic map.

Finally, substituting the decomposition of Π into Eq. (6.484) and then using the definition of Ξ_A^B, one obtains

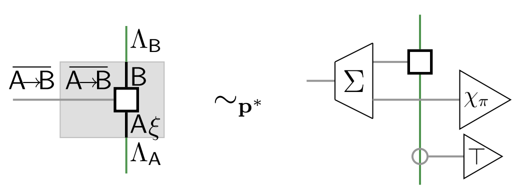

(6.491)

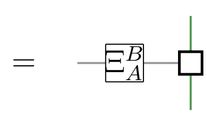

(6.492)

That is, every classical realist representation is inferentially equivalent to updating one's knowledge about the operational procedure to knowledge about functional dynamics. □