Quantum and Classical Markovian Graphical Causal Models and Their Identification

Jonathan Barrett
University of Oxford, UK

Isaac Friend
University of Oxford, UK

Aleks Kissinger
University of Oxford, UK

Abstract

Markov categories allow formalization of probabilistic and causal reasoning in a general setting that applies uniformly to many different kinds of classical probabilistic processes. It has so far been challenging, however, to generalize these techniques to reasoning about quantum processes, as the quantum no-cloning theorem forbids "copy" maps of the sort that have been used to axiomatize conditional independence, and the related notions of complete common causes and Markovianity, in classical Bayesian networks. Here, we introduce a new categorical notion of Markovian causal model, according to which a distinguished subcategory of "common cause" maps plays a similar role to that of "copy" maps in the categorical formulation of Bayesian networks. Moreover, defining causal models as second-order processes yields a clean and flexible formulation of interventions. Our formalism is both rich enough to handle "complete common cause" assumptions and general enough to encompass not only standard classical causal identification scenarios, but also quantum causal scenarios and new kinds of classical causal identification based on imperfect observations. Furthermore, we show that one can reason uniformly across all of these cases using string-diagrammatic techniques.

2012 ACM Subject Classification Computing methodologies → Causal reasoning and diagnostics; Theory of computation → Categorical semantics; Theory of computation → Quantum computation theory

Keywords and phrases causal inference, Bayesian networks, quantum combs, process theories

Digital Object Identifier 10.4230/LIPIcs.CSL.2025.48

Funding *Isaac Friend*: This research was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development and by the Province of Ontario through the Ministry of Colleges and Universities.
Aleks Kissinger: This publication was made possible through the support of the ID# 62312 grant from the John Templeton Foundation, as part of the "The Quantum Information Structure of Spacetime" Project (QISS). The opinions expressed in this publication are those of the author(s) and do not necessarily reflect the views of the John Templeton Foundation.

Acknowledgements We thank Elie Wolfe and Rob Spekkens for insights regarding the classes of instruments that can be used for causal identification.

1 Introduction

Within the study of probabilistic reasoning, causal inference involves discerning from data the causal relationships responsible for generating them, and hence the effects of hypothetical interventions. Many researchers in quantum information and the foundations of quantum theory have tried to adapt concepts from causal inference to a setting in which, roughly, quantum systems replace random variables as causal relata, and quantum channels replace functions or stochastic maps as causal mechanisms [15, 17, 19, 1, 12]. A natural approach to studying quantum generalizations of probabilistic reasoning is to start from the literature on categorical probability theory (e.g., [14, 11]), and simply replace a category of probabilistic processes with a category of quantum processes, which one hopes satisfies enough axioms to support analogous calculations to the classical case. But we quickly run into a major obstacle: an important axiom used in categorical probability to define *Markov categories*, and slight variations such as CD and CDU categories, is the existence of "copy" maps for all objects. Such maps are the basis of abstract formulations of conditional independence, and hence the Markov condition for Bayesian networks. When Bayesian networks are understood as causal models as in the work of Pearl [16], Markovianity sometimes involves a variable being "copied" and the copies distributed to other parts of the network that depend on the variable. The Markov condition for causal Bayesian networks is central to classical causal inference.

Such "copy" maps are forbidden for quantum processes, however, as famously shown by the quantum no-cloning [20] and no-broadcasting [2] theorems. While it is possible to define complete common causes, and hence Markovianity, in the quantum setting in several equivalent ways [1, 3], the absence of an explicit representation of copying as a well-defined quantum process in its own right limits the translation of standard causal inference techniques to the quantum setting. The present work solves this problem by introducing an abstract, categorical notion of Markovian causal model that can be instantiated in either a quantum or a classical setting[1]. As in prior work on categorical causal inference [13] (which depended on "copy" structure and did not include quantum models), a causal model involves two symmetric monoidal categories: a *syntactic category*, whose morphisms encode an abstract causal structure as a formal composition of "black boxes," and a *semantic category*, in which the abstract causal structure is functorially interpreted as a particular data-generating process, e.g., by filling in the black boxes with concrete stochastic matrices. Generalizing [13], we will provide a notion of abstract causal structure that can be interpreted in either a classical or a quantum semantic category. In particular, this framework subsumes ordinary classical Bayesian networks.

After defining the basic framework, we will describe how our formalism handles interventions, and pose the *causal identification* problem to which we will apply our mathematical technology. Causal identification is a type of causal inference problem for which the effects of counterfactual interventions are to be inferred from a combination of qualitative hypotheses (represented by a graph) and observational data. Our formulation of causal models lets us treat the statistics from a very restricted class of interventions as "observational data" available for inference. The precise class of interventions we choose is treated as a parameter in our framework, so different classes of interventions yield different kinds of causal identification problems. This flexibility is needed in the quantum case, where there is no standard notion of "passive observation," but is also useful in the classical case, where we can now study inference tasks whose input data have been obtained via imperfect procedures.

In classical causal inference, the assumption that an unknown causal model is a Markovian model based on a known graph, amounting to the assumption that there are no latent variables influencing multiple observed variables (i.e., no latent confounders), greatly expands the class of causal queries that can be answered with observational data. In particular, with this assumption, one can identify from the graph and the observational data the response of the model to arbitrary interventions. We demonstrate that a certain Markovianity

[1] The resulting notion of quantum Markovian causal model, specifically $(\mathbf{CPM}, \mathbf{Unitary}_{\bullet})$-valued Markovian model, is closely related to proposals in [8] and [1, 3].

assumption formulated in our new framework is similarly powerful in the quantum setting, allowing the identification of the entire data-generating process from only very limited probing operations. We simultaneously demonstrate that for classical causal identification, noisy and disturbing observations can sometimes serve in place of ordinary perfect passive observations. The uniform handling of the classical and quantum cases is made possible by the string-diagrammatic calculus for (compact) symmetric monoidal categories.

2 Process theories

Throughout the paper, we will use *process theoretic* terminology, following, e.g., [5], to discuss morphisms in a symmetric monoidal category. Namely, we will refer to symmetric monoidal categories $(\mathcal{C}, \otimes, I)$ as *process theories* and the morphisms therein as *processes*. Because of their physical interpretation, we also introduce special terminology for morphisms into and out of the monoidal unit I. In a process theory, morphisms of the form $\rho : I \to A$ are called *states*, and morphisms of the form $\pi : A \to I$ are called *effects*. Morphisms of the form $\lambda : I \to I$ are called *numbers* or *scalars*.

We will focus on process theories equipped with distinguished families of *discarding* maps $d_A : A \to I$, one map for each object A, satisfying $d_{A\otimes B} = d_A \otimes d_B$ and $d_I = 1_I$. The main utility of discarding maps is allowing us to say when a process is *causal*, which in the classical and quantum settings imposes a normalization constraint.

▶ **Definition 1.** *A process $f : A \to B$ is called causal if $d_B \circ f = d_A$.*

▶ **Example 2.** The process theory $\mathbf{Mat}[\mathbb{R}_+]$ has as objects natural numbers and as processes $M : m \to n$ the $n \times m$ matrices whose entries are non-negative real numbers $\{M^i_j \mid 1 \leq i \leq n, 1 \leq j \leq m\}$. The monoidal product is given by tensor product of matrices (a.k.a. Kronecker product), whose unit is the 1×1 matrix $[1] : 1 \to 1$. Discarding maps $d_n : n \to 1$ are the $1 \times n$ matrices (i.e., row vectors) consisting of all 1s. Composing with d_n corresponds to summing over an output index (i.e., marginalization). Consequently, causal states are column vectors of positive numbers whose entries sum to 1 (i.e., probability distributions), and causal processes are matrices whose columns each sum to 1 (i.e., stochastic maps, equivalent to conditional probability distributions with $P(i|j) := M^i_j$).

▶ **Example 3.** The process theory **CPM** has as objects finite-dimensional Hilbert spaces $\mathcal{H}, \mathcal{K}, \ldots$ and as morphisms completely positive maps $\Phi : L(\mathcal{H}) \to L(\mathcal{K})$, where $L(\mathcal{H})$ is the algebra of operators $\mathcal{H} \to \mathcal{H}$. The monoidal product is again given by tensor product, whose unit is the identity map on $L(\mathbb{C}) \cong \mathbb{C}$. Discarding maps are trace maps. A state $\rho : \mathbb{C} \to L(\mathcal{H})$ is fixed by a single positive operator $\rho(1) \in L(\mathcal{H})$, and causal states correspond to trace-1 positive operators. More generally, causal processes are the trace-preserving completely positive maps.

Since both matrices of positive numbers and completely positive maps are closed under sums, both $\mathbf{Mat}[\mathbb{R}_+]$ and **CPM** are additively enriched. We will first use this fact in Definition 4 to define *instruments*.

The presentation will use *string diagram* notation, with processes depicted as boxes and objects as wires in diagrams read from bottom to top. A process theory's monoidal unit object I and the identity process $I \to I$ are both depicted by empty space, and other identity processes are depicted as wires. Discarding, which will later serve as a counit for an internal comonoid structure, is depicted with a black dot.

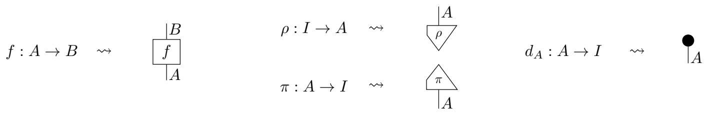

Diagrammatically, the causality condition for a process $f : A \to B$ from Definition 1 is

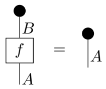

(1)

3 Causal Bayesian networks

The usual notion of a joint probability distribution being Markov compatible with a directed acyclic graph (DAG) is that it factorizes in such a way that each variable (labeling a node of the graph) is independent when conditioned on its parents. For example, a joint distribution $P(ABCDE)$ is Markov compatible with DAG

$G :=$ 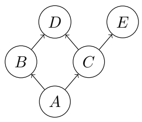

(2)

precisely when $P(ABCDE) = P(A)P(B|A)P(C|A)P(D|BC)P(E|C)$.

We now recall the string-diagrammatic formulation of Markov compatibility of a joint probability distribution with a DAG G, given, e.g., in [13]. First, we introduce for each object X in $\mathbf{Mat}[\mathbb{R}_+]$ a "copy" map $X \to X \otimes X$, whose composition with a point distribution $\psi : I \to X$ (i.e., a column vector with a single 1 entry and 0s elsewhere) is $\psi \otimes \psi$. With "copy" as comultiplication–depicted by a black dot with one input and two output wires–and discarding as counit, each object in $\mathbf{Mat}[\mathbb{R}_+]$ is given a cocommutative comonoid structure:

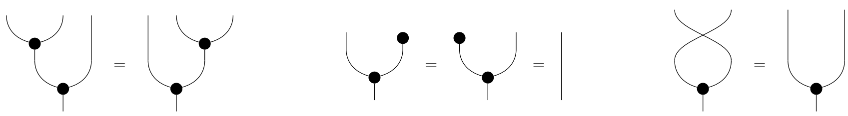
(3)

A symmetric monoidal category equipped with a compatible family of "copy" and discard maps for all objects is called a *CD category*. If in addition we impose the causality condition of Definition 1 on all maps, the category is called a *Markov category*. The copy and discard maps above endow $\mathbf{Mat}[\mathbb{R}_+]$ with the structure of a CD category; the subcategory $\mathbf{Stoch} \subseteq \mathbf{Mat}[\mathbb{R}_+]$ of stochastic (i.e., causal) maps is a Markov category.

We can form the string diagram associated with a DAG G by introducing a box $a : X_1 \otimes \ldots \otimes X_n \to A$ for every node A in G with parents $\{X_1, \ldots, X_n\}$. We compose these boxes by connecting each output A to the output of the overall diagram, as well as to the inputs of each of the children of A in G, introducing copy maps where necessary. A state being Markov with respect to G then means simply that it factorizes according to that diagram, for some choice of stochastic matrices $a, b, c, \ldots$ For example, a state ω is Markov with respect to the graph G from (2) when it can be decomposed as follows:

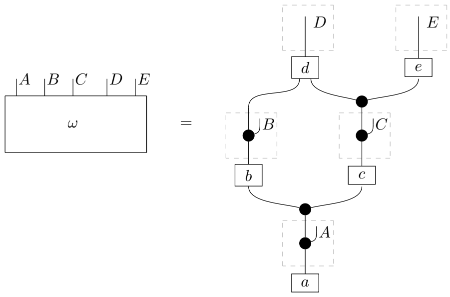

(4)

This diagrammatic condition corresponds precisely to the usual factorization of $P(ABCDE)$ given before, where $P(ABCDE)$ is the joint probability distribution given by the state ω.

The right-hand side of (4) may be said to represent ω as a *Bayesian network.* While Bayesian networks can in general just be seen as an efficient way to represent joint probability distributions, we can additionally provide them with a causal interpretation, using them to model how a certain scenario would respond to possible interventions. To model the effects of interventions using a Bayesian network, we interpret each of the boxes $a, b, c, \ldots$ as some actual (e.g., physical) mechanism that determines (stochastically) the value of its output, given any value of its input. One introduces a concept of local intervention whereby, for example, a change can be made at the input to box c while the rest of the network is left unchanged. In [13], local intervention is represented by endofunctors that "cut" diagrams like the one in Eq. 4. Such an intervention results in a different overall state from ω (the new state is sometimes called a "do-conditional" in the causal inference literature [16]). A Bayesian network representation of a state that is Markov compatible with DAG G, interpreted causally in this way, is called a *Markovian G-based causal model.*

In a causal inference problem, one is given a state like ω together with certain qualitative assumptions about how the state is generated, and the task is to determine further properties of the data-generating process and compute how ω would change under hypothetical interventions. One sort of qualitative assumption is that ω is generated by a Markovian causal model based on a certain DAG. Such an assumption turns out quite powerful for inference: with it, one can evaluate the results of essentially any hypothetical intervention. Discussions of "quantum causal modeling" naturally suggest the question of whether a "quantum Markovianity assumption" might provide similar inferential power in quantum causal scenarios. We therefore seek to formulate quantum Markovian causal models based on DAGs, and an associated inference problem.

Two obstacles arise. First, it is unclear what quantity constitutes the quantum analog of the state (ω in Eq. 4) that is an input in classical inference. That state carries "observational data," i.e., the probability distribution generated by the causal model when variables are merely observed rather than intervened on. In operational quantum theory, there is no standard notion of passive observation as distinct from more "active" intervention. Our solution makes no such distinction in principle, and instead simply allows any set of interventions to be declared the "accessible" ones whose outcome distributions will be available for inference. The second obstacle is the absence of copy maps in quantum theory. The assumption that observational data are generated by a process like the one on the right-hand side of Eq. 4 is useful for inference because copy maps guarantee, for example, that any randomness shared between the inputs to b and c is accounted for by variable A. (The Bayesian network representation in Eq. 4 also uses copy maps to produce an observed output for each variable while allowing the variable's value to be fed forward, undisturbed, to the rest of the network.) Our solution here allows any subtheory of maps in a process theory to take the role typically played by copy maps in distributing "information" from the output of one box to the inputs of other boxes. With such a general framework, there remains the question of which choices of parameters in the quantum setting give a specific notion of "Markovian causal model" that is especially useful for inference. The answer, in Section 7, depends on a theorem showing how unitary quantum maps (more generally, what are called "autonomous" quantum channels) mimic the "function-maps" in $\mathbf{Mat}[\mathbb{R}_+]$.

4 Generalized causal models

4.1 Combs and instruments

The interventional causal models studied in this paper will involve second-order processes, or *combs* [4], taking first-order processes as input and producing other first-order processes (usually numbers, i.e. processes $I \to I$). We will represent second-order processes as first-order ones by invoking the (self-dual) *compact structures* in the process theories $\mathbf{Mat}[\mathbb{R}_+]$ and **CPM**: every object A is equipped with a pair of maps $\cup_A : I \to A \otimes A$ and $\cap_A : A \otimes A \to I$, called "cups" and "caps" respectively, satisfying the so-called *yanking equations*, which are depicted in string diagram notation as follows:

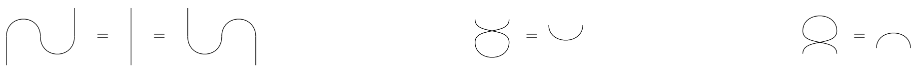

In $\mathbf{Mat}[\mathbb{R}_+]$, cups and caps are given by Kronecker delta matrices, with the two indices treated as either inputs or outputs: $\cup^{ij} = \cap_{ij} = \delta_{ij}$. In **CPM**, $\cup_{\mathcal{H}}$ is given by the unnormalized maximally entangled state $\sum_{ij} |ii\rangle\langle jj|$ and $\cap_{\mathcal{H}}$ is its associated effect, seen as a completely positive map from $L(\mathcal{H}) \otimes L(\mathcal{H})$ to $\mathbb{C}$.

Using this structure, we can, for example, represent a process that takes processes of type $A \to A'$ and produces processes of type $B \to B'$ as a normal, first-order process $f : B \otimes A' \to A \otimes B'$. We then indicate its higher-order interpretation by drawing f as a box with a "hole" in it, often called a *comb*, and use cups and caps to define "plugging" another box into that hole:

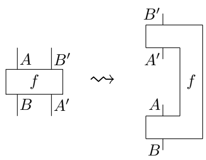

(5)

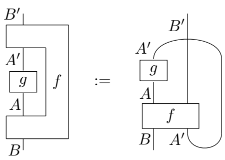

(6)

As in [9, 10], a classical or quantum causal model will involve a comb in $\mathbf{Mat}[\mathbb{R}_+]$ or **CPM**, respectively, encoding the stable mechanisms governing a repeated causal scenario. The "holes" in the comb will represent loci of intervention, where one can interact with the data-generating process in various ways, e.g., by implementing a causal map (a classical or quantum channel), or observing the value of a random variable and then feeding forward a certain state. An intervention procedure at a "hole" in a classical or quantum comb will be represented mathematically by an *instrument*.

▶ **Definition 4.** *An instrument of type $A \to A'$ valued in $\mathbf{Mat}[\mathbb{R}_+]$ or **CPM** is a finite set of maps $\{\phi_i : A \to A'\}_i$ whose sum $\sum_i \phi_i$ is a causal map. Each map ϕ_i is called a branch of the instrument.*

Branches correspond to possible outcomes of the intervention procedure. The probabilities of these outcomes are determined by the branches and by the process in which one is intervening.

▶ **Example 5.** The preparation of a causal state is represented by an instrument branch $\{\rho : I \to A\}$. A *demolition measurement* on a quantum system is represented by an instrument whose branches are effects $\{\phi_i : A \to I\}_i$. The only causal effect is the discarding map d_A, so the instrument condition says $\sum_i \phi_i = d_A$. For a causal state ρ, the probability of getting outcome i is $P(i|\rho) := \phi_i \circ \rho$. From the instrument condition, it follows that $\sum_i P(i|\rho) = \sum_i \phi_i \circ \rho = d_A \circ \rho = 1$.

If f in Eq. 5 is a causal process in $\mathbf{Mat}[\mathbb{R}_+]$ or **CPM**, and one selects instruments $I \to B$, $A \to A'$, and $B' \to I$, then f will map each possible triple of branches to a probability, understood as the probability of realizing this triple when probing f with the selected instruments. Causal inference in general consists in using such probabilities–imagined to have been learned experimentally over many trials–to compute properties of f, whose value is initially unknown, and thereby predict how f would respond to other combinations of instruments.

This paper focuses on a kind of causal inference problem called *causal identification*, for which certain properties of the comb, namely its "shape," are assumed in advance, and those assumptions used together with the probabilities just described to compute the further properties of interest. The assumption we will formalize and use for inference is *Markovianity*. We will now give a process-theoretic definition of Markovian causal model that can be instantiated in either $\mathbf{Mat}[\mathbb{R}_+]$–where we recover a standard definition of Markovian causal model–or **CPM**.

4.2 Abstract and concrete causal structures

We extend the recipe from [13] where a directed acyclic graph is used to generate a process theory whose morphisms are abstract causal structures encoding qualitative assumptions that will be used for inference. For a finite directed acyclic graph $G = (V_G, E_G)$ with vertex set V_G and edge set E_G, let $\mathbf{G}_0$ be a free symmetric monoidal category whose objects are generated by the set $V_G \uplus E_G$ and whose morphisms are generated by discarding maps for all objects and two additional kinds of maps:

$$x : e_1 \otimes \ldots \otimes e_j \to X \qquad\qquad A_X : X \to e'_1 \otimes \ldots \otimes e'_k \tag{7}$$

for each $X \in V_G$, where $\{e_1, \ldots, e_j\}$ are the in-edges of X and $\{e'_1, \ldots, e'_k\}$ are the out-edges of X.

From the free category $\mathbf{G}_0$, we form the "syntactic" process theory $\mathbf{G}$ by additionally imposing the causality equation (1) for every generating map. In particular, for Z a vertex with no out-edges, $A_Z = d_Z$.

We then associate to the graph G a process $c_G : X_1 \otimes \ldots \otimes X_n \to X_1 \otimes \ldots \otimes X_n$ in $\mathbf{G}$, called the *abstract causal structure* associated with G, by taking each of the generators from (7) and plugging each input wire labeled by an edge to the unique associated output wire. The inputs and outputs of c_G are all labeled by vertices, each vertex labeling exactly one input and one output wire. Each input/output pair is depicted as a hole in a wire and labeled by the corresponding vertex in V_G.

▶ **Example 6.** The directed graph G indicates the abstract causal structure c_G:

$G =$ 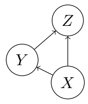

$c_G =$ 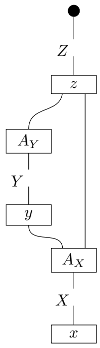

The graph is one of the inputs for the causal inference problem we will be studying. Just as in [13] being given the graph in (2) would let an agent assume that the observed probability distribution is generated by a process conforming to (4), being given the three-node graph in this example will let an agent assume that the unknown causal scenario is represented by a comb conforming to c_G. A class of *interventional causal models* respecting this abstract causal structure is defined relative to a pair of process theories $(\mathcal{C}, \mathcal{C}_{cc})$, where $\mathcal{C}_{cc}$ is a subtheory of $\mathcal{C}$ called the "common-cause" subtheory. (The common-cause subtheory is a parameter in the framework; specifying the common-cause subtheory is part of defining a causal identification problem. We will study the consequences of various choices of common-cause subtheory.)

▶ **Definition 7.** *A G-based, $(\mathcal{C}, \mathcal{C}_{cc})$-valued Markovian interventional causal model consists of a discarding-preserving functor of process theories (i.e., a discarding-preserving symmetric monoidal functor) $F : \boldsymbol{G} \to \mathcal{C}$ such that $F(A_X)$ is in $\mathcal{C}_{cc}$ for every $X \in V_G$.*

The process $F(c_G)$ is a concrete causal structure, i.e., it is a morphism in $\mathcal{C}$, such as a stochastic matrix (in the classical case) or a quantum channel, that assigns probabilities to outcomes of intervention procedures implemented at *intervention loci* (*loci* for short) represented by the input/output pairs that form the "holes" in the abstract and concrete causal structures. We will consider scenarios in which F is initially unknown, and we will try to compute elements in F's image from those probabilities.

In these scenarios, the process theories $\mathcal{C}$ and $\mathcal{C}_{cc}$, like the graph G, are given in advance. A process $F(A_X)$ in the common-cause subtheory $\mathcal{C}_{cc}$ distributes information from locus X toward the loci labeled by X's children in G. The common-cause subtheory determines what it means to assume (ultimately for the purpose of inference) that a locus is the *complete common cause* of the loci labeled by its children in G. There are no conditions on $\mathcal{C}_{cc}$ *a priori*, except that it should contain the family of discarding processes. In particular, we could have $\mathcal{C}_{cc} = \mathcal{C}$, in which case the notion of "complete common cause" is trivialized. However, for some classes of models below, $\mathcal{C}_{cc}$ will be a subtheory of processes that we think of as disallowing confounding between their outputs due to latent variables/systems. In this case, any observed correlations are thought of as arising entirely from the causal dependency of multiple output variables/systems on some input. We will formulate this concept for relevant subtheories of classical and quantum processes, where it will bestow significant inferential power. The basic idea is simple: if one wishes to infer the value of a process known to decompose according to a certain string diagram, then knowing that certain boxes are valued in a smaller subtheory will tend to make the task easier. We will show that certain classical and quantum subtheories are particularly useful in this regard, for mathematical reasons that are precisely analogous between the two settings. Nevertheless, it is important to understand that the term "Markovian" is used in Def. 7 in a new and abstract sense.

The causal models studied in this paper are valued in the process theories $\mathbf{Mat}[\mathbb{R}_+]$ and **CPM**. Denote by **Func** the subtheory of $\mathbf{Mat}[\mathbb{R}_+]$ consisting of function-maps, i.e., stochastic maps whose columns each contain precisely one 1. This theory is equivalent to the theory of finite sets and functions: associate with each matrix M in **Func** the unique function f with $M^i_j := 1$ if and only if $f(j) = i$. Relevant subtheories of **CPM** include the theories **Unitary** and **Isom** of unitary and isometric quantum channels, i.e. completely positive maps of the form $\mathcal{U}(\rho) := U \circ \rho \circ U^\dagger$ for U a unitary or an isometry. For a subtheory $\mathcal{D}$ of a process theory $\mathcal{C}$ with discarding, the theory $\mathcal{D}_\bullet$ is formed by adjoining discarding maps for all objects. Thus we have, e.g., $\mathbf{Func}_\bullet = \mathbf{Func}$, and $\mathbf{Unitary}_\bullet$ is the theory of what are called *autonomous quantum channels* [18].

4.3 Recovering classical causal Bayesian networks

One reason for our use of the term "Markovian" in a manner specific to the new framework we are introducing is that this framework subsumes ordinary Markovian causal Bayesian networks. The key to establishing the relationship is the following property of function-maps in $\mathbf{Mat}[\mathbb{R}_+]$[2]:

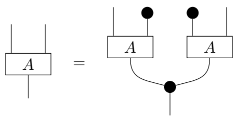

(8)

▶ **Proposition 8.** *For any* $(\mathbf{Mat}[\mathbb{R}_+], \mathbf{Func}_\bullet)$*-valued model of directed acyclic graph* G*, the state in* $\mathbf{Mat}[\mathbb{R}_+]$ *derived by plugging a copy map into each locus (as in the left-hand side of Eq. 9) is Markov compatible with* G *in the standard sense used in probabilistic graphical modeling. Conversely, any Bayesian network based on DAG* G *is derivable from some* G*-based,* $(\mathbf{Mat}[\mathbb{R}_+], \mathbf{Func}_\bullet)$*-valued causal model by this prescription.*

Proof. A state like the one in Eq. 4 is a G-based, $(\mathbf{Mat}[\mathbb{R}_+], \mathbf{Func}_\bullet)$-valued model composed with copy maps at all loci. The common-cause functions following the loci happen also to be copy maps, which are indeed morphisms in $\mathbf{Func}_\bullet$. On the other hand, starting from a $(\mathbf{Mat}[\mathbb{R}_+], \mathbf{Func}_\bullet)$-valued model with generic common-cause functions also yields a state of this form, thanks to Eq. 8 for functions:

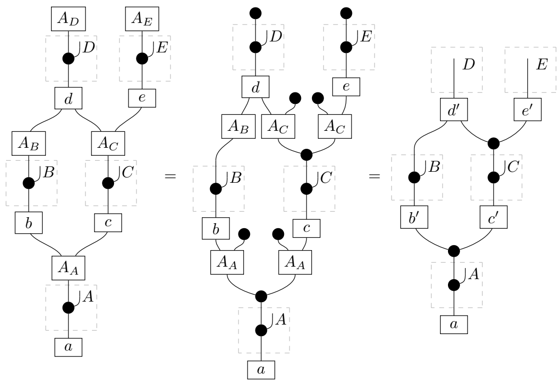

(9)

◀

[2] In synthetic probability based on Markov categories [11], Eq. 8 defines conditional independence of a map's outputs (conditioned on its input).

We have now shown using Eq. 8 how some of the copy maps in classical Bayesian networks emerge in our framework when $\mathbf{Func}_\bullet$ is the common-cause subtheory. In the next section, after posing the general identification problem, we will show how standard "observational data" can be extracted from our classical interventional models via what we call *perfect passive observation* instruments, which do not involve copy maps. We will then have recovered the standard notion of classical causal identification as just one case on the same footing as quantum and new classical problems.

5 Interventions and the identification problem

▶ **Definition 9.** *For a given abstract causal structure, a semantic process theory $\mathcal{C}$, and an object A in $\mathcal{C}$ for each locus in the abstract causal structure, a local intervention regime assigns to each A an instrument in $\mathcal{C}$ of type $A \to A$.*

For a fixed local intervention regime and a fixed classical or quantum model of the abstract causal structure, "implementing" the intervention regime for one iteration of the causal scenario results in the joint realization of a combination of maps at all the loci: at each locus, one branch of the instrument assigned to that locus is realized. The joint probability of this combination of local outcomes is the number resulting from plugging the maps into their loci. The problem of causal identification is to use such probabilities from a limited set of local intervention regimes, together with the shape of the abstract causal structure (equivalently, the graph), to infer probabilities of outcome combinations under other local intervention regimes.

The set of local intervention regimes whose outcome statistics are to be used for inference is constructed as follows: for each locus A, an *accessible set* $\mathcal{I}_A$ of instruments $A \to A$ is given. These accessible sets of instruments define a set of *accessible* local intervention regimes.

▶ **Definition 10.** *Given an accessible set $\mathcal{I}_A$ of instruments for each locus A in an abstract causal structure, an accessible local intervention regime assigns to each A an instrument from $\mathcal{I}_A$.*

The probabilities available for inference are the probabilities that can be "learned" from accessible local intervention regimes. That is, for each accessible local intervention regime, the joint probability of each combination of branches will be considered known.

Note that we will always assume every accessible set $\mathcal{I}_A$ contains the identity instrument of the appropriate type. An identity instrument has one branch, an identity process, depicted by a wire. Assuming the universal availability of identity instruments is a way of assuming that given any allowed local intervention regime, one can also choose to "do nothing" at one of the loci, keeping the same instruments at all other loci[3].

In causal identification, what one is trying to identify is the image of some map in the syntactic process theory $\mathbf{G}$ under F. Knowing such an image might allow one to determine how the model would respond to certain local intervention regimes. For instance, if one were confronted with an unknown model of the abstract causal structure in Example 6, inferring the value of $F(y \circ A_X)$ would allow one to predict the outcome probabilities for a

[3] The reason identity interventions are usually not discussed in classical causal inference literature is that they can be simulated from perfect passive observational data, by "marginalization." This is no longer the case, however, when we move beyond classical perfect passive observation. For example, performing a quantum measurement and then marginalizing over the outcome will *not* in general lead to the same statistics on the remaining loci as not doing the measurement at all.

local intervention regime consisting of an identity instrument at Z and arbitrary instruments at X and Y. However, one is initially given only limited access to the functor F, namely the probabilities associated with accessible local intervention regimes applied to the whole data-generating process $F(c_G)$. For simplicity, we will focus on the problem of inferring $F(c_G)$ itself, from which one can compute the outcome probabilities for arbitrary local intervention regimes. In general, however, we might only be interested in predicting the results of interventions at certain loci, in which case we might be able to focus on a simpler problem.

In full, the causal identification problem we will study is defined by the following in each instance: a directed acyclic graph G, specifying an abstract causal structure $c_G \in \mathbf{G}$; a semantic process theory $\mathcal{C}$ (either $\mathbf{Mat}[\mathbb{R}_+]$ or $\mathbf{CPM}$) and a subtheory $\mathcal{C}_{cc}$; a G-based, $(\mathcal{C}, \mathcal{C}_{cc})$-valued Markovian interventional causal model F; and an accessible set $\mathcal{I}_A$ of instruments for each locus A.

The inputs for the identification task are the (labeled) graph G, the pair of concrete process theories $(\mathcal{C}, \mathcal{C}_{cc})$, the accessible set of instruments for each locus A, and the data generated by $F(c_G)$ under each accessible local intervention regime (i.e., the joint probabilities of realizing combinations of branches). The task is to compute $F(c_G)$. If this task is possible, we will say G-based $(\mathcal{C}, \mathcal{C}_{cc})$-valued models are *identifiable* from the accessible sets of instruments.

For the kinds of classical and quantum causal scenarios we are studying, there always exist finite sets of local instruments that, if declared accessible, suffice for identification regardless of the common-cause subtheory[4]. In contrast, we will consider how accessible sets that do not suffice for identification of models with one common-cause subtheory become sufficient when the common-cause subtheory is further restricted.

A typical example in the case of $\mathcal{C} = \mathbf{Mat}[\mathbb{R}_+]$ is for the accessible set of instruments at each locus to consist of the "perfect passive observations." Note that when reasoning in both the classical and quantum process theories simultaneously, we draw generic states and effects as asymmetric triangles, reserving symmetric triangles for $\mathbf{Mat}[\mathbb{R}_+]$ alone.

▶ **Definition 11.** *For object A in $\mathbf{Mat}[\mathbb{R}_+]$, the perfect passive observation instrument of type $A \to A$ has branches*

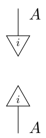

where the state labeled i is given by the column vector with 1 in row i and all other entries 0, and the effect labeled i is the matrix transpose of that column vector.

Knowing which branch of a perfect passive observation instrument has been implemented means being certain of the value a random variable has taken, and certain that the variable retains that value as it is input to subsequent causal mechanisms.

In the previous section, we saw that any $(\mathbf{Mat}[\mathbb{R}_+], \mathbf{Func}_\bullet)$-valued model can be translated diagrammatically into a Bayesian network by plugging "copy" maps into all loci. This procedure is equivalent to considering such a model with perfect passive observations at each locus. We can see this by noting that the probability of any particular joint outcome associated with a joint state ω such as the one in (4) can be obtained by plugging in the effect associated to that outcome, e.g.:

[4] See p. 107 of [9] for a way to construct such sets, and explanation of how the construction represents the idea of controlled experiments, from which we expect to be able to deduce any data-generating process.

$P(A = a, B = b, C = c, D = d, E = e) \;=$ 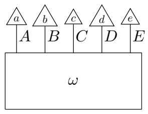

Then, we can apply the following equation satisfied by the copy and any effect associated with a unit vector:

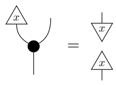

to obtain a perfect passive observation at every locus; e.g.,

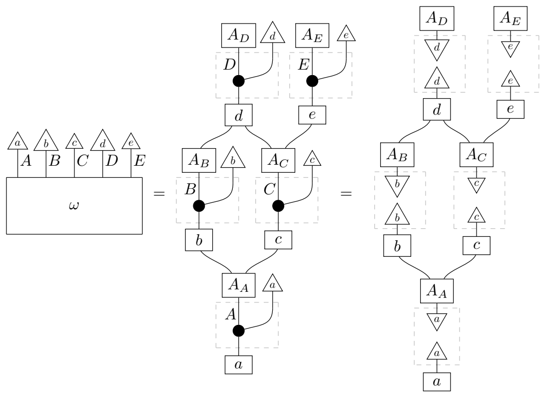

Hence, knowing the state ω is the same as knowing the probabilities associated with perfect passive observations.

We now study what kinds of classical models can be identified from perfect passive observations.

▶ **Proposition 12.** *Perfect passive observation instruments do not suffice for identifying* $(\boldsymbol{Mat}[\mathbb{R}_+], \boldsymbol{Mat}[\mathbb{R}_+])$*-valued Markovian models.*

The proposition means that for some graphs G, there are multiple G-based, $(\mathbf{Mat}[\mathbb{R}_+], \mathbf{Mat}[\mathbb{R}_+])$-valued Markovian models that behave identically under all local intervention regimes involving only perfect passive observation instruments, but differently under other local intervention regimes. It should not surprise readers familiar with causal inference; if the common-cause maps can be arbitrary stochastic matrices, they can essentially introduce confounding. The modifier "Markovian" would not ordinarily be applied to generic instances of what we are calling $(\mathbf{Mat}[\mathbb{R}_+], \mathbf{Mat}[\mathbb{R}_+])$-valued Markovian models. It would, however, describe what we call $(\mathbf{Mat}[\mathbb{R}_+], \mathbf{Func}_\bullet)$-valued Markovian models. For these models, where common-cause maps are restricted to functions, there are no hidden confounders and identification from perfect passive observation is always possible.

▶ **Proposition 13.** $(\boldsymbol{Mat}[\mathbb{R}_+], \boldsymbol{Func}_\bullet)$*-valued Markovian models are identifiable from perfect passive observation instruments.*

In the proof in Appendix A, the last rewriting step uses the fact that the classical "copy" map literally copies the state that leaves a locus after a perfect passive observation. In quantum inference, and in classical inference with generalized observation, this calculation will be unavailable, and a new technique will be introduced to take its place.

5.1 Quantum and generalized classical observation

In some classical causal scenarios, observational data are noisy[5]. After one learns the result of a test, one's credences about the possible values of the variable are given by a probability distribution. Furthermore, one's credences about the possible values being fed forward may be different–one may understand that the procedure whereby one learns about the variable's value tends to change the value. We now study causal identification in such situations and in quantum scenarios by process-theoretically generalizing perfect passive observation instruments to kinds of classical and quantum instruments that may be "noisy" rather than "perfect," and "disturbing" rather than "passive," but are consistent with the idea of observational instruments as those for which the state leaving a locus is determined by the effect that has been realized there, and not by the experimenter's further choice. Quantum causal identification with, e.g., projective measurement instruments turns out similar to classical causal identification with noisy and disturbing instruments.

The generalized observations that will constitute the accessible sets of instruments are such that when one learns an outcome, one models the observation as having implemented an effect followed by a state.

▶ **Definition 14.** *A process of type $A \to A'$ is called $\circ$-separable if it consists of an effect $A \to I$ followed by a state $I \to A'$.*

▶ **Definition 15.** *An instrument of type $A \to A'$ is $\circ$-separable if each of its branches is a $\circ$-separable process.*

▶ **Example 16.** A surgical intervention, composed of a discarding map followed by a causal state preparation, corresponds to a $\circ$-separable instrument with one branch.

▶ **Example 17.** A classical perfect passive observation instrument is a $\circ$-separable instrument.

▶ **Example 18.** Any orthonormal basis for a finite-dimensional Hilbert space induces a $\circ$-separable **CPM**-valued instrument called an ONB measurement (a.k.a. a non-degenerate von Neumann measurement).

The entities that will serve as well as perfect passive observation instruments for causal identification are "complete sets of $\circ$-separable instruments," whose definition invokes the concept of "informational completeness."

▶ **Definition 19.** *A set of effects $\{\pi_j : A \to I\}$ is called informationally complete for A if any state $\rho : I \to A$ is uniquely determined by the set of numbers $\pi_j \circ \rho$. Similarly, a set of states $\{\rho_j : I \to A\}$ is called informationally complete for A if any effect $\pi : A \to I$ is uniquely determined by the set of numbers $\pi \circ \rho_j$.*

In $\mathbf{Mat}[\mathbb{R}_+]$ and **CPM**, a set is informationally complete if and only if it spans the relevant vector space, where the vector space associated with an object A in **CPM** is the space $L(A)$ of linear operators on Hilbert space A.

▶ **Definition 20.** *A set of $\circ$-separable instruments of type $A \to A$ will be called complete if (i) the set of all states appearing in the branches of the instruments is informationally complete for A, and (ii) the set of all effects appearing in the branches of the instruments is informationally complete for A.*

[5] Hidden Markov models, the standard means of modeling this phenomenon, graphically represent variables quite differently from causal Bayesian networks, and moreover are not meant for studying interventions.

▶ **Example 21.** In $\mathbf{Mat}[\mathbb{R}_+]$, a perfect passive observation instrument by itself constitutes a complete set of $\circ$-separable instruments.

▶ **Example 22.** For any object in **CPM**, there is a finite set of ONB measurements forming a complete set of $\circ$-separable instruments. One example is the set of ONB measurements corresponding to the bases of eigenvectors for the d^2 generalized Pauli matrices on a Hilbert space of dimension d.

▶ **Example 23.** In $\mathbf{Mat}[\mathbb{R}_+]$, let

$$\phi = \begin{bmatrix} .8 & .9 \end{bmatrix} \qquad \phi' = \begin{bmatrix} .2 & .1 \end{bmatrix} \qquad \psi = \begin{bmatrix} .5 \\ .5 \end{bmatrix} \qquad \psi' = \begin{bmatrix} .9 \\ .1 \end{bmatrix}$$

The instrument with branches $\psi \circ \phi$ and $\psi' \circ \phi'$ constitutes a single-instrument marginally informationally complete set of $\circ$-separable instruments of type $2 \to 2$. When a locus representing a binary random variable (with values denoted 1 and 2) is probed with this instrument, if the variable's true value is 1, $\psi \circ \phi$ is realized with probability .8 and $\psi' \circ \phi'$ with probability .2. If the true value of the variable is 2, $\psi \circ \phi$ is realized with probability .9 and $\psi' \circ \phi'$ with probability .2. If branch $\psi \circ \phi$ is realized, the value of the variable fed forward after the probing is totally randomized. If branch $\psi' \circ \phi'$ is realized, the value fed forward is 1 with probability .9 and 2 with probability .1. Instruments that are "biased" toward the realization of certain branches can be thought of as modeling certain kinds of selection effects, which are an important topic of study both in statistics in general and in causal identification research [7].

We will show how classical and quantum Markovian causal models with appropriately restricted common-cause subtheories can be identified when the accessible set of instruments at each locus is a complete set of $\circ$-separable instruments.

6 Quantum common causes and convolution of maps

In this section, we will establish that autonomous quantum channels satisfy a quantum version of Eq. 8, where the quantum meaning of the "copy" dot will be given. Ultimately Eq. 8 will be exploited for identification of both classical and quantum Markovian causal models.

We pass from **CPM** to the larger process theory **FVect** of all linear maps to introduce super-operators that are not completely positive but will be used for diagrammatic quantum causal inference. First, we define super-operators

$$\mu : L(\mathcal{H}_A) \otimes L(\mathcal{H}_A) \to L(\mathcal{H}_A) \qquad \delta : L(\mathcal{H}_A) \to L(\mathcal{H}_A) \otimes L(\mathcal{H}_A)$$

where μ corresponds to matrix multiplication, i.e., $\mu(\rho \otimes \sigma) := \rho\sigma$, and δ is its adjoint with respect to the Hilbert-Schmidt inner product. The latter is easiest to describe concretely by its action on basis elements written in Dirac's "bra-ket" notation (see Appendix B):

$$\delta(|i\rangle\langle j|) := \sum_k |i\rangle\langle k| \otimes |k\rangle\langle j|$$

Now we introduce diagrammatic notation like that used in the classical case. We'll write:

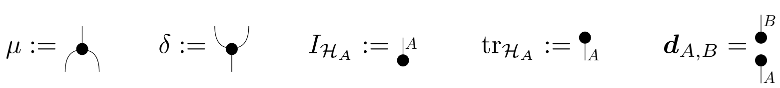

It is straightforward to show a few basic identities using these generators, e.g.,

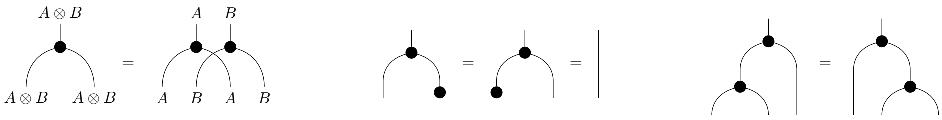

and their (vertical) mirror-images, which imply that δ and $\mathrm{tr}_{\mathcal{H}_A}$ give every object of the form $L(\mathcal{H}_A)$ in **FVect** a comonoid structure. Here, in contrast to the classical case, the comonoid structure is non-cocommutative, i.e., Eq. 3 does not hold. This structure has been described in [6].

The classical maps depicted by 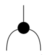 and  are the matrix transposes of those depicted by the already-defined vertical mirror images of the respective diagrams.

The following two definitions are the diagrammatic equivalents of those for the same terms in Appendix B.

▶ **Definition 24.** *The convolution* $\Phi_1 * \Phi_2$ *of two quantum or classical maps* $\Phi_1, \Phi_2 : A \to B$ *is* $\mu \circ (\Phi_1 \otimes \Phi_2) \circ \delta$*.*

▶ **Definition 25.** *For both* ***FVect*** *and* ***Mat***$[\mathbb{R}_+]$*, the convolution inverse of a map, indicated by that map's diagram inside* $\{-\}^{-1}$*, satisfies*

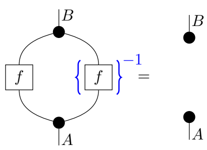

When this notation is used in calculations for Section 7, one or both of the objects A and B will be the unit object, i.e., f will be an effect or a number. This notation is consistent with the use of $\{-\}^{-1}$ for inverses of positive real numbers in the proof of Prop. 13.

With the quantum semantics for the black dot, the essential similarity between autonomous quantum channels and function-maps can be stated as follows:

▶ **Theorem 26.** *Any autonomous quantum channel A satisfies Eq. 8.*

The theorem follows from Propositions 31 and 36 in Appendix B. It implies that the common-cause maps in $(\mathbf{CPM}, \mathbf{Unitary}_\bullet)$-valued Markovian models can be rewritten just as the classical common-cause function-maps are, e.g., in the first equality of (9). This rewriting, together with convolution inverses, will yield an identification technique for $(\mathbf{CPM}, \mathbf{Unitary}_\bullet)$ models.

7 Identification

We will study the causal identification problem as described in Section 5, focusing on classical and quantum cases in which the accessible set of instruments at each locus is a complete set of $\circ$-separable instruments. Specifically, we will show, for the smallest graph in which two arrows leave a single vertex, that in both the classical and quantum settings, if the common-cause subtheory $\mathcal{C}_{cc}$ is appropriately restrictive, any complete sets of $\circ$-separable instruments at all loci suffice for identification; otherwise, generic such sets do not suffice.

▶ **Example 27.** Strictly positive Markovian models based on the graph G of Example 6 valued in $(\mathbf{Mat}[\mathbb{R}_+], \mathbf{Mat}[\mathbb{R}_+])$ or $(\mathbf{CPM}, \mathbf{Isom}_\bullet)$ are not identifiable from arbitrary complete sets of $\circ$-separable instruments.

For $(\mathbf{Mat}[\mathbb{R}_+], \mathbf{Mat}[\mathbb{R}_+])$, the statement follows from Proposition 12. For $(\mathbf{CPM}, \mathbf{Isom}_\bullet)$, an example of a pair of models that respond differently to some interventions but are indistinguishable under local projective measurement can be constructed from the example at the end of Section 3 of [9]. (The labels X and Z are swapped relative to the labeling in Example 6; the common-cause map $\boldsymbol{A_Z}$ is the parallel composite of an identity map and the state $\boldsymbol{u}$.) Here common-cause maps from $\mathbf{Isom}_\bullet$ can be thought of as potentially introducing unseen auxiliary systems that correlate outcomes at multiple loci.

If the common-cause subtheory is restricted to $\mathbf{Func}_\bullet$ in the classical case or $\mathbf{Unitary}_\bullet$ in the quantum case, complete sets of ∘-separable instruments at all loci become sufficient for identification.

▶ **Proposition 28.** *Markovian models based on the graph G of Example 6 valued in $(\boldsymbol{Mat}[\mathbb{R}_+], \boldsymbol{Func}_\bullet)$ or $(\boldsymbol{CPM}, \boldsymbol{Unitary}_\bullet)$ are identifiable whenever the accessible set of instruments at each locus is a complete set of ∘-separable instruments.*

The proof, in Appendix A, applies Theorem 26 and convolution inverses of quantum and classical maps.

8 Conclusion

This paper has addressed the problem of defining quantum Markovian graphical causal models by formulating causal models as combs and replacing the copy maps of categorical probability by "common cause" maps describing how information may be shared among intervention loci. The framework allows the formulation of causal inference problems parametrized by the process theory, common-cause subtheory, and available sets of instruments. When the quantum common-cause subtheory consists of autonomous quantum channels, Theorem 26 gives the quantum causal models a structure that can be used to identify the models even when one is given access to only highly restricted probing operations. Meanwhile, the formalism permits the study of new kinds of classical causal identification problems solvable when function-maps are taken as the classical common-cause subtheory.

References

1 John-Mark A. Allen, Jonathan Barrett, Dominic C. Horsman, Ciarán M. Lee, and Robert W. Spekkens. Quantum Common Causes and Quantum Causal Models. *Phys. Rev. X*, 7(3):031021, July 2017. `doi:10.1103/PhysRevX.7.031021`.

2 Howard Barnum, Carlton M. Caves, Christopher A. Fuchs, Richard Jozsa, and Benjamin Schumacher. Noncommuting Mixed States Cannot Be Broadcast. *Phys. Rev. Lett.*, 76(15):2818–2821, April 1996. Publisher: American Physical Society. `doi:10.1103/PhysRevLett.76.2818`.

3 Jonathan Barrett, Robin Lorenz, and Ognyan Oreshkov. Quantum Causal Models, 2019. `doi:10.48550/arXiv.1906.10726`.

4 G. Chiribella, G. M. D'Ariano, and P. Perinotti. Quantum Circuit Architecture. *Phys. Rev. Lett.*, 101(6):060401, August 2008. Publisher: American Physical Society. `doi:10.1103/PhysRevLett.101.060401`.

5 Bob Coecke and Aleks Kissinger. *Picturing Quantum Processes: A First Course in Quantum Theory and Diagrammatic Reasoning.* Cambridge University Press, 2017. `doi:10.1017/9781316219317`.

6 Bob Coecke and Robert W. Spekkens. Picturing classical and quantum Bayesian inference. *Synthese*, 186(3):651–696, June 2012. `doi:10.1007/s11229-011-9917-5`.

7 Juan D. Correa, Jin Tian, and Elias Bareinboim. Identification of Causal Effects in the Presence of Selection Bias. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):2744–2751, July 2019. Section: AAAI Technical Track: Knowledge Representation and Reasoning. `doi:10.1609/aaai.v33i01.33012744`.

8 Fabio Costa and Sally Shrapnel. Quantum causal modelling. *New Journal of Physics*, 18(6):063032, June 2016. Publisher: IOP Publishing. `doi:10.1088/1367-2630/18/6/063032`.

9 Isaac Friend and Aleks Kissinger. Identification of Causal Influences in Quantum Processes. In Stefano Gogioso and Matty Hoban, editors, *Proceedings 19th International Conference on Quantum Physics and Logic, Wolfson College, Oxford, UK, 27 June - 1 July 2022*, volume 394 of *Electronic Proceedings in Theoretical Computer Science*, pages 101–115. Open Publishing Association, 2023. `doi:10.4204/EPTCS.394.7`.

10 Isaac Friend and Aleks Kissinger. Identification of causal influences in quantum processes. *Physical Review A*, 109(4):042214, April 2024. Publisher: American Physical Society. `doi:10.1103/PhysRevA.109.042214`.

11 Tobias Fritz. A synthetic approach to Markov kernels, conditional independence and theorems on sufficient statistics. *Advances in Mathematics*, 370:107239, August 2020. `doi:10.1016/j.aim.2020.107239`.

12 Christina Giarmatzi and Fabio Costa. A quantum causal discovery algorithm. *npj Quantum Information*, 4(1):17, March 2018. `doi:10.1038/s41534-018-0062-6`.

13 Bart Jacobs, Aleks Kissinger, and Fabio Zanasi. Causal Inference by String Diagram Surgery. In Mikołaj Bojańczyk and Alex Simpson, editors, *Foundations of Software Science and Computation Structures*, pages 313–329, Cham, 2019. Springer International Publishing. `doi:10.1007/978-3-030-17127-8_18`.

14 Bart Jacobs and Fabio Zanasi. The Logical Essentials of Bayesian Reasoning. *CoRR*, abs/1804.01193, 2018. arXiv: 1804.01193. `arXiv:1804.01193`.

15 M. S. Leifer and Robert W. Spekkens. Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference. *Phys. Rev. A*, 88(5):052130, November 2013. Publisher: American Physical Society. `doi:10.1103/PhysRevA.88.052130`.

16 Judea Pearl. *Causality*. Cambridge University Press, 2 edition, 2009. `doi:10.1017/CBO9780511803161`.

17 Katja Ried, Megan Agnew, Lydia Vermeyden, Dominik Janzing, Robert W. Spekkens, and Kevin J. Resch. A quantum advantage for inferring causal structure. *Nature Physics*, 11(5):414–420, May 2015. `doi:10.1038/nphys3266`.

18 Benjamin Schumacher and Michael D. Westmoreland. Locality and Information Transfer in Quantum Operations. *Quantum Information Processing*, 4(1):13–34, February 2005. `doi:10.1007/s11128-004-3193-y`.

19 Sally Shrapnel, Fabio Costa, and Gerard Milburn. Quantum Markovianity as a supervised learning task. *International Journal of Quantum Information*, 16(08):1840010, 2018. `doi:10.1142/S0219749918400105`.

20 W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, October 1982. `doi:10.1038/299802a0`.

A Identifiability proofs

In our study of identifiability conditions, we always assume that all models, whether classical or quantum, are *strictly positive* in the following sense:

▶ **Definition 29.** *An interventional causal model based on graph G is called strictly positive if for each generating map of the form $x : e_1 \otimes \ldots \otimes e_j \to X$ in $\boldsymbol{G}$, in the case $\mathcal{C} = \boldsymbol{Mat}[\mathbb{R}_+]$ the stochastic matrix $F(x)$ has only strictly positive entries, or in the case $\mathcal{C} = \boldsymbol{CPM}$ the quantum channel $F(x)$ has full Choi rank.*

The Choi rank of a quantum channel is defined in terms of the channel's Choi matrix, discussed in Appendix B. When a strictly positive classical or quantum model is composed with any non-zero state and any non-zero effect, the result is a strictly positive real number. Our strict positivity assumption serves a similar purpose to that of standard requirements of strictly positive distributions in causal inference, which guarantee that relevant conditional probabilities are defined; here the guarantee is that scalars inverted in our identification protocols are in fact non-zero, and more generally that effects have "convolution inverses."

One common feature of $\mathbf{Mat}[\mathbb{R}_+]$ and $\mathbf{CPM}$ that will help with causal identification is *local process tomography.*

▶ **Proposition 30.** *The theories $\boldsymbol{Mat}[\mathbb{R}_+]$ and $\boldsymbol{CPM}$ have local process tomography: any process $f : A \otimes B \to C \otimes D$ is determined by numbers*

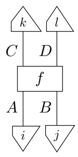

(10)

where i, j, k, and l index any informationally complete sets of states or effects for the appropriate objects.

We will often leave the interpretation functor F from the syntactic process theory $\mathbf{G}$ into the semantic process theory implicit and use boldface to distinguish abstract processes in $\mathbf{G}$ from their images under F, writing, e.g., $\boldsymbol{x} := F(x)$. Labels for objects/intervention loci will be identical between the syntactic and semantic process theories, since the distinction will already be clear from the labels for processes.

Proof of Proposition 13. The proposition can be proven with techniques from [13], via the equivalence we have discussed between our $(\mathbf{Mat}[\mathbb{R}_+], \mathbf{Func}_\bullet)$-valued Markovian models and the causal Bayesian networks formulated in that article. Here, however, we prove the proposition only for one graph, so that the procedure can be compared directly with the one in Section 7 for quantum and classical identification from generalized observation.

An unknown $(\mathbf{Mat}[\mathbb{R}_+], \mathbf{Func}_\bullet)$-valued Markovian model based on the graph G in Example 6 is a functorial interpretation of the abstract causal structure c_G in that example. Since the common-cause subtheory is $\mathbf{Func}_\bullet$, Eq. (8) applies to the common-cause maps, which can then be absorbed into larger processes $\boldsymbol{y'}$ and $\boldsymbol{z'}$, resulting in a new representation for the unknown model:

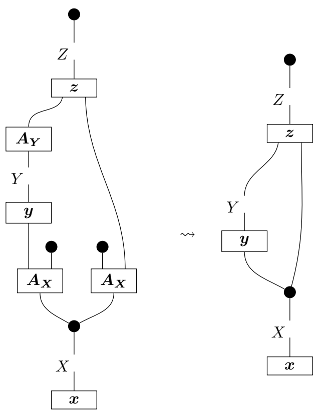

(11)

Inferring the values of the processes $\boldsymbol{x}$, $\boldsymbol{y'}$, and $\boldsymbol{z'}$ on the right-hand side–which one does by inferring all their matrix entries–is equivalent to inferring the process $F(c_G)$.

One first computes the processes $\boldsymbol{x}$ and $\boldsymbol{y'}$ by determining all their matrix entries. For $\boldsymbol{x}$, the numbers

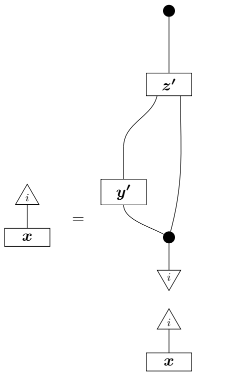

are obtained via a local intervention regime consisting of perfect passive observation at locus X and identity interventions at Y and Z. The process $\boldsymbol{y'}$ is tomographically determined as follows:

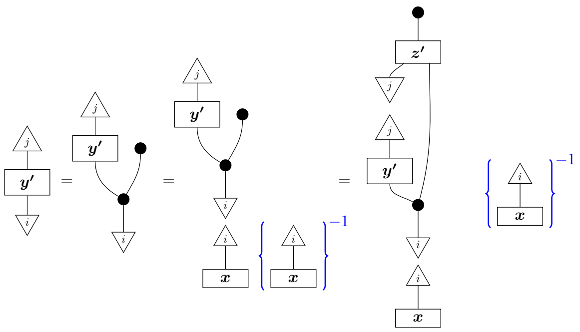

The desired quantity has been rewritten as the product of a probability obtained via perfect passive observation at X and Y (and identity intervention at Z) and a number obtainable from the known value of $\boldsymbol{x}$.

Finally, $\boldsymbol{z'}$ is computed from probabilities obtained via a local intervention regime consisting of perfect passive observation at all three loci, and from inverses of numbers already determined in previous steps.

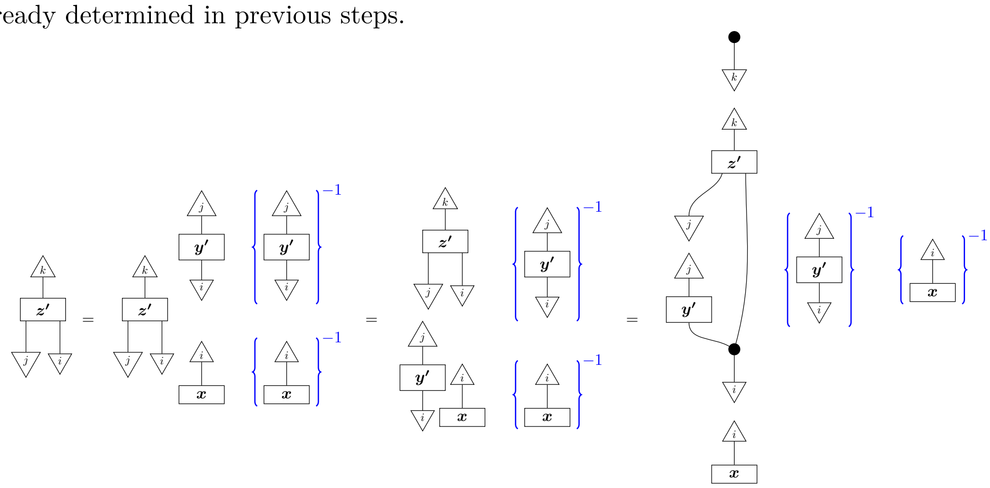

◀

Proof of Proposition 28. As in the demonstration of Prop. 13, but now for both the classical and quantum cases, Eq. 8 and defining new unknown processes $\boldsymbol{y'}$ and $\boldsymbol{z'}$ lead to a simplification of the unknown model as shown in expression (11). To identify the model, one proceeds as before to compute the processes $\boldsymbol{x}$ and $\boldsymbol{y'}$ by determining the probabilities given by their composition with appropriate informationally complete sets of states and effects.

The union of the accessible set of instruments at a locus, say X, is a set of maps, indexed by, say, ϕ. Each map ϕ is composed of an effect $f_X(\phi)$ and a state $g_X(\phi)$, where f_X and g_X are functions associated with locus X. Marginal informational completeness of the set of instruments means that the set $\{f_X(\phi)\}_\phi$ of effects and the set $\{g_X(\phi)\}_\phi$ of states are each informationally complete for system-type X.

One determines $\boldsymbol{x}$ by learning for the informationally complete set of effects $f_X(\phi)$ the probability

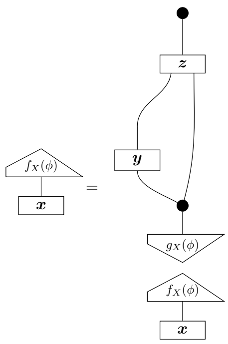

The right-hand diagram is a probability learned from probing with ∘-separable instruments at locus X and identity instruments elsewhere. In contrast to the case of classical perfect passive observation, inferring the value of $\boldsymbol{x}$ now might involve more than one local intervention regime, so that X can be probed with multiple instruments.

Next, one proceeds to determine $\boldsymbol{y'}$ tomographically as in the proof of Prop. 13, but in general collating data from multiple local intervention regimes:

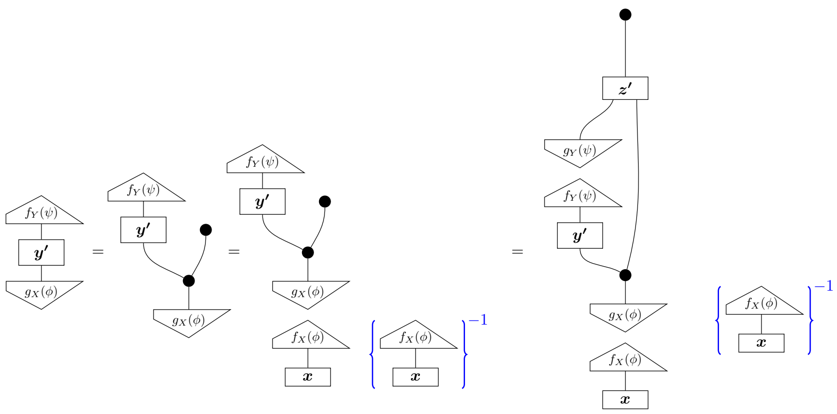

At this point, in the case of classical perfect passive observation, the computation of $\boldsymbol{z}$ in the proof of Prop. 13 uses the fact that the classical "copy" map literally copies the pure states that leave a locus after a perfect observation. For quantum measurements, even maximally informative projective measurements, and for generalized classical observation, the state leaving the first locus is not copiable, and hence this calculation isn't available.

In this general case, once one has computed the values of $\boldsymbol{x}$ and $\boldsymbol{y'}$, one can tomographically determine the value of the process

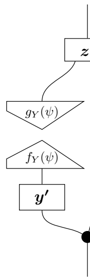

(12)

for each map ψ in a marginally informationally complete set, from the numbers

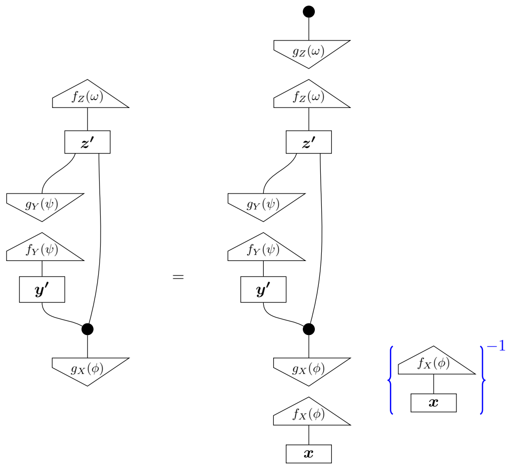

which one can learn for informationally complete sets of states $g_X(\phi)$ and effects $f_Z(\omega)$.

Knowing also the value of $\boldsymbol{y'}$, one can compute for each ψ the convolution inverse of $f_Y(\psi) \circ \boldsymbol{y'}$, and compose it with the process (12) as follows:

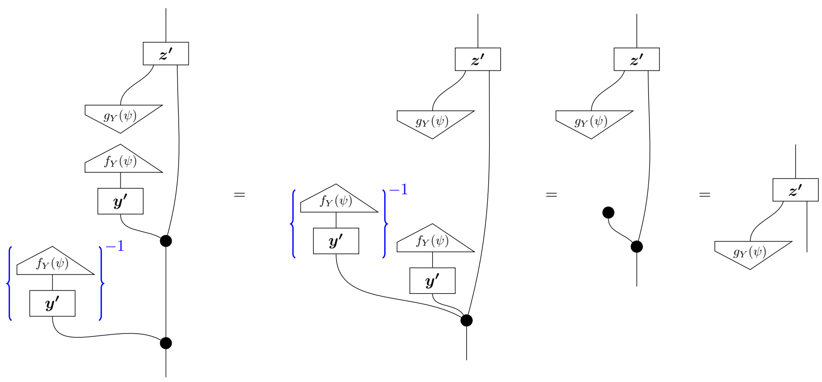

One now knows the latter process for an informationally complete set of states $g_X(\psi)$, and hence obtains the value of the process $\boldsymbol{z'}$. The known processes $\boldsymbol{x}$, $\boldsymbol{y'}$, and $\boldsymbol{z'}$ can now be composed to form the entire data-generating process in expression (11), completing the causal inference. ◀

B Choi-Jamiołkowski isomorphism and channel convolution

The Choi-Jamiołkowski isomorphism gives a bijective correspondence between linear super-operators $\mathcal{E} : L(\mathcal{H}_A) \to L(\mathcal{H}_B)$ and linear maps $\rho^{\mathcal{E}} : \mathcal{H}_B \otimes \mathcal{H}_A^* \to \mathcal{H}_B \otimes \mathcal{H}_A^*$. The linear map $\rho^{\mathcal{E}}$, called the *Choi matrix* of $\mathcal{E}$, can be defined explicitly in terms of a basis $\{|i\rangle_A\} \subset \mathcal{H}_A$ and its dual basis $\{|i\rangle_{A^*}\} \subset \mathcal{H}_A^*$ as follows:

$$\rho^{\mathcal{E}} := \sum_{ij} \mathcal{E}(|i\rangle_A\langle j|) \otimes |i\rangle_{A^*}\langle j| \tag{13}$$

Here we have used Dirac's "bra-ket" notation to write operators/matrices as products of basis vectors ("kets" $|i\rangle_A$) and their associated dual vectors ("bras" ${}_A\langle i|$).

The Choi-Jamiołkowski isomorphism states that $\mathcal{E}$ is completely positive if and only if $\rho^{\mathcal{E}}$ is positive. Now, for a quantum channel $\Phi : L(\mathcal{H}_A) \to L(\mathcal{H}_B) \otimes L(\mathcal{H}_C)$ define the following three positive operators:

$$\rho_{BC|A} := \rho^{\Phi} \qquad \rho_{B|A} := I_{\mathcal{H}_C} \otimes \mathrm{tr}_{\mathcal{H}_C}(\rho^{\Phi}) \qquad \rho_{C|A} := I_{\mathcal{H}_B} \otimes \mathrm{tr}_{\mathcal{H}_B}(\rho^{\Phi}) \tag{14}$$

We can regard each of these as an operator on $\mathcal{H}_B \otimes \mathcal{H}_C \otimes \mathcal{H}_A^*$ (note that we have suppressed "swap" maps above).

Theorem 2 in [1] implies the following, which will be used to establish Eq. 8 for autonomous quantum channels[6].

▶ **Proposition 31.** *If Φ is an autonomous quantum channel, then it satisfies*

$$\rho_{BC|A} = \rho_{B|A}\rho_{C|A} \tag{15}$$

We therefore proceed to study channels satisfying Eq. 15.

B.1 Channel convolution

Satisfying Eq. 15 will be shown equivalent to decomposing in a certain way with respect to the following convolution operation for superoperators. For (not necessarily completely positive) linear maps $\Phi_1, \Phi_2 : L(\mathcal{H}_A) \to L(\mathcal{H}_B)$, let $\Phi_1 * \Phi_2 : L(\mathcal{H}_A) \to L(\mathcal{H}_B)$ be a new linear map defined on basis elements $|i\rangle\langle j| \in L(\mathcal{H}_A)$ as follows:

$$\Phi_1 * \Phi_2(|i\rangle\langle j|) := \sum_k \Phi_1(|i\rangle\langle k|)\Phi_2(|k\rangle\langle j|)$$

First, we show that the Choi-Jamiołkowski isomorphism carries this operation to matrix multiplication.

▶ **Lemma 32.** *Let $\rho^{\Phi_1}, \rho^{\Phi_2}$, and $\rho^{\Phi_1 * \Phi_2}$ be the Choi matrices of the super-operators Φ_1, Φ_2, and $\Phi_1 * \Phi_2$, respectively. Then $\rho^{\Phi_1}\rho^{\Phi_2} = \rho^{\Phi_1 * \Phi_2}$.*

Proof. Unroll (13) and simplify. ◀

Although the convolution of an arbitrary pair of completely positive maps need not be completely positive, the convolution of a pair of completely positive maps that commute under convolution is completely positive:

---

[6] Theorem 2 in [1] is used in that article to motivate a definition of quantum Markovianity based on the idea that directed edges in a graph should indicate signaling relations between input and output systems of a unitary quantum channel.

▶ **Corollary 33.** *For completely positive maps* $\Phi_1, \Phi_2 : L(\mathcal{H}_A) \to L(\mathcal{H}_B)$*, the super-operator* $\Phi_1 * \Phi_2$ *is completely positive if and only if* $\Phi_1 * \Phi_2 = \Phi_2 * \Phi_1$*.*

Proof. Suppose $\Phi_1 * \Phi_2 = \Phi_2 * \Phi_1$. Then, using Theorem 32, we have: $\rho^{\Phi_1}\rho^{\Phi_2} = \rho^{\Phi_1 * \Phi_2} = \rho^{\Phi_2 * \Phi_1} = \rho^{\Phi_2}\rho^{\Phi_1}$. So $\rho^{\Phi_1 * \Phi_2}$ is the product of commuting positive operators ρ^{Φ_1} and ρ^{Φ_2}, and hence positive itself. Therefore $\Phi_1 * \Phi_2$ is completely positive. Conversely, if $\Phi_1 * \Phi_2$ is completely positive then $\rho^{\Phi_1 * \Phi_2}$ is positive. Hence, by Theorem 32, $\rho^{\Phi_1}\rho^{\Phi_2}$ is also positive, which is only possible if the positive operators ρ^{Φ_1} and ρ^{Φ_2} commute. This in turn implies that $\rho^{\Phi_1 * \Phi_2} = \rho^{\Phi_2 * \Phi_1}$. By an inverse application of the Choi-Jamiołkowski isomorphism, we conclude that $\Phi_1 * \Phi_2 = \Phi_2 * \Phi_1$. ◀

Now, we can get a fully channel-based version of Eq. 15. For Hilbert spaces $\mathcal{H}_A, \mathcal{H}_B$, we define the (un-normalized) depolarizing channel as follows for all states $\rho \in L(\mathcal{H}_A)$:

$$\boldsymbol{d}_{A,B}(\rho) = \mathrm{tr}(\rho) I_{\mathcal{H}_B}$$

This channel is equivalent to the identity operator under the Choi-Jamiołkowski isomorphism, so it behaves as a unit for channel convolution.

For a channel $\Phi_{BC|A} := \Phi$, we define the reduced channels $\Phi_{B|A}$ and $\Phi_{C|A}$ simply by applying $\boldsymbol{d}$ to the appropriate output:

$$\Phi_{B|A} := (1_{\mathcal{H}_B} \otimes \boldsymbol{d}_{\mathcal{H}_C}) \circ \Phi \qquad \Phi_{C|A} = (\boldsymbol{d}_{\mathcal{H}_B} \otimes 1_{\mathcal{H}_C}) \circ \Phi$$

One can straightforwardly check that the Choi matrices of these channels are the reduced states in (14). From that fact and Lemma 32, we can immediately conclude:

▶ **Lemma 34.** *A channel* $\Phi : L(\mathcal{H}_A) \to L(\mathcal{H}_B) \otimes L(H_C)$ *has Choi matrices satisfying Eq. 15 if and only if* $\Phi_{BC|A} = \Phi_{B|A} * \Phi_{C|A}$*.*

▶ **Definition 35.** *The convolution inverse of a completely positive map* Φ *is a linear map* $\Phi^{(-1)}$ *satisfying* $\Phi * \Phi^{(-1)} = \Phi^{(-1)} * \Phi = \boldsymbol{d}_{A,B}$*.*

A completely positive map has a convolution inverse if and only if its associated Choi matrix ρ_Φ is invertible; in that case, the convolution inverse is the linear map defined by the usual matrix inverse ρ_Φ^{-1} under the Choi-Jamiołkowski isomorphism. For classical positive matrices, we can define the convolution inverse similarly, and it is given concretely by the positive matrix whose elements are $(M^{(-1)})_{i,j} := 1/M_{i,j}$.

From the graphical rules for convolution in Section 6, we can derive the condition for a quantum map to satisfy Eq. 8:

▶ **Proposition 36.** *The Choi matrix of channel* $\Phi : L(\mathcal{H}_A) \to L(\mathcal{H}_B) \otimes L(\mathcal{H}_C)$ *factors according to Eq. 15 if and only if*

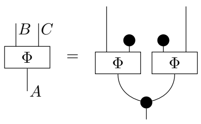

Proof. Writing Φ as $\Phi_{BC|A}$, we apply Lemma 34, then simplify:

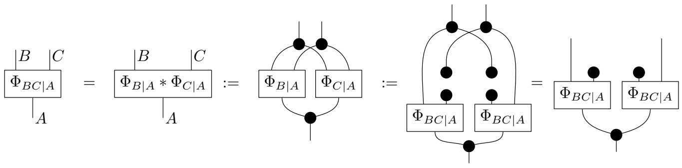

◀