# Identification of causal influences in quantum processes

Isaac Friend and Aleks Kissinger

Department of Computer Science, University of Oxford
{isaac.friend,aleks.kissinger}@cs.ox.ac.uk

June 21, 2022

**Abstract**

Causal identification is a type of causal inference problem concerned with recovering from observational data and qualitative assumptions the causal relationships generating the data, and hence the effects of hypothetical interventions. Though the topic is typically considered in the context of classical statistical models, recent years have seen great interest in extending causal inference techniques to quantum and generalised theories. A major obstacle to a theory of causal identification in the quantum setting is the question of what should play the role of "observational data," as any means of extracting data at a certain locus will almost certainly disturb the system. Hence, one might think *a priori* that quantum measurements are already too much like interventions, so that the problem of causal identification trivialises. This is not the case. Fixing a limited class of quantum instruments (namely the class of all projective measurements) to play the role of "observations," we note that as in the classical setting, there exist scenarios for which causal identification is not possible. We then present sufficient conditions for quantum causal identification, starting with an example of a quantum analogue of the well-known "front-door criterion" and finishing with a broader class of scenarios for which the effect of a single intervention is identifiable. These results arise from generalising the process-theoretic account of classical causal inference given by Jacobs, Kissinger, and Zanasi beyond the setting of Markov categories, and thereby treating the classical and quantum problems uniformly.

## 1 Introduction

The problem of causal inference is to deduce from statistical correlations among variables something about the causal mechanisms responsible for those correlations, where a causal mechanism is a process that answers *interventional* queries. Although the majority of the work in the field of causal inference has focused on classical, statistical models, it is interesting to consider causal inference problems in the quantum setting as well, where quantum systems play the role of classical random variables. One can ask, for example, whether it is possible to deduce using only certain limited operations whether agents are in a common-cause type setting (e.g., accessing two parts of a quantum entangled state) or a cause-effect type setting (e.g., accessing the same system at two points in time). Ried et al. presented a solution to this inference problem for specific scenarios involving two quantum systems, and asked how their scheme might generalize to scenarios with more systems [Rie+15]. This article begins to answer that question, using the logical conception of causality presented in [JKZ19] to reveal the common process-theoretic underpinnings of causal inference in both ordinary stochastic and quantum settings.

A theory of quantum causal inference requires first a mathematical model of quantum causal scenarios. Here, we will take a minimal notion of a quantum causal model consisting of a "circuit with holes," i.e., a directed acyclic string diagram where some wires have gaps in them allowing agents to apply local processes. This can be seen as a second-order process, or *comb* [CDP08], which maps local non-deterministic processes to probabilities.

This notion of causal model is relatively weak in that unlike the one studied in [All+17] and [BLO20], it doesn't seem to admit a relation of "complete common cause" whereby a single, tensor-inseparable quantum system can act as the sole source of correlations between two systems in its future. On the other hand, the quantum interventional models studied here do correspond to the quantum causal models of [CS16]. Though our definition of interventional causal model includes non-Markovian models, all the quantum combs falling under our definition can be given larger Markovian explanations in the sense of [CS16], and the quantum models in our Propositions 5 and 3 explicitly depict all the latent laboratories needed for such Markovian explanations. Our identification criteria can therefore be restated in terms of the directed acyclic graphs used in more traditional presentations of both classical [Pea09] and quantum [CS16] causal modeling.

As complete common causes can be pictured in the classical setting as "copying" a random variable and using it as input to two or more subsequent stochastic maps, it is difficult and somewhat subtle to make sense of a "complete quantum common cause" in the absence of a physically meaningful process of cloning or broadcasting quantum systems. Hence, it is interesting to see how much traction we can get on causal inference for a class of models which don't admit the explicit general representation of complete common causes. We will show here that, in the case of the particular problem of *quantum causal identifiability*, we can get relatively far without such a representation. We also recover, from an abstract perspective, theorems in classical statistical causal inference.

Causal identification, in the classical case, refers to the problem of identifying the effects of (often hypothetical) interventions on the basis of purely observational data [Pea09]. In contrast to related problems such as causal discovery, here the hypothesised causal structure of events (e.g., the directed acyclic graph)–reflecting assumptions that certain variables cannot possibly exert direct influence on certain others–is known in advance, but not the exact conditional probability distributions (or functional dependencies) governing the influence of individual variables on each other. Even with the causal structure given in advance, however, this problem can be highly non-trivial in the presence of confounding variables [Pea09] or selection bias [CTB19].

In generalising to quantum causal identification, one needs to fix a notion that stands in the place of "observation," as it is impossible to extract any data from a quantum system without causing a disturbance, which in some sense is already a type of intervention. Here, we fix the class of processes playing the role of "observations" as local projective measurements, whereas "interventions" can be arbitrary quantum instruments. The latter includes, for example, the process of discarding the incoming state of a system and preparing a fixed new state chosen by the agent, while the former does not.

While we do not intend to argue here that these notions of "observation" and "intervention" are fully conceptually justified, we will give strong evidence instead that this kind of quantum causal identification problem is *interesting*: we note that the problem can be hard in general, then give criteria under which it becomes easy.

In the quantum just as in the classical case, there exist scenarios in which causal identification is impossible, i.e., in which there is a pair of models which behave identically with respect to projective measurements, but whose behaviour differs under arbitrary interventions. Simple such pairs of models were mentioned in [Rie+15].

Our first result is an illustration of a quantum version of the front-door criterion for causal identifiability [Pea09]. This result is then generalised to a sufficient condition for identification that implies the quantum analogues of multiple sufficient conditions in the statistical causal modeling literature, including the front-door criterion and some cases covered by Galles and Pearl in [GP95] and by Tian and Pearl in [TP02]. The statements and proofs here invoke diagrammatic technology presented in [CK17] and previously applied to causal inference by Jacobs, Kissinger, and Zanasi [JKZ19], who indicated the possibility, realized in the present article, of "import[ing] results from classical causal reasoning to the quantum case" by changing the concrete process theory in which abstract causal diagrams are modeled.

Consequences of the work are both practical and conceptual. Practically, the results here guarantee that certain causal influences are identifiable in quantum networks of certain shapes, and describe how to identify them. This paper thus initiates a quantum parallel of the systematic general study of causal identification now codified in textbooks and routinely applied to the analysis of real data. As Ried et al. explained, their inference schemes "promise extensive applications in experiments exhibiting quantum effects." [Rie+15] The schemes presented here might be similarly applicable, particularly to the problem of detecting non-Markovianity in quantum information processing [AKP06]. On the conceptual side, the process-theoretic presentation here, yielding identification protocols formally identical to known classical ones, illuminates the structures and procedures–comb factorization, bases, and process tomography–that underpin causal inference regardless of whether the probability theory governing the variables is classical.

The isolation of these rudiments should not only help guide the further development of theories of causal inference for quantum and other special kinds of processes, but also motivate continued research in ordinary statistical causal modeling using the logical and compositional techniques of theoretical computer science.

## 2 Preliminaries

To treat classical and quantum theory on the same footing, we will use the language of *process theories* throughout. Process theories have been defined in slightly varying ways in the literature. Here we will define a process theory as follows.

**Definition 1.** A *process theory* is a symmetric monoidal category $(\mathcal{C}, \otimes, I)$ equipped with a distinguished family of *discarding* morphisms $d_A : A \to I$ for each object $A$ satisfying $d_{A\otimes B} = d_A \otimes d_B$ and $d_I = 1_I$.

To give a physical or computational interpretation to process theories, it is typical to refer to generic morphisms $f : A \to B$ as *processes*, morphisms of the form $\rho : I \to A$ as *states*, and morphisms of the form $\pi : A \to I$ as *effects*. Objects are also called *system-types*. Throughout the paper, we will adopt *string diagram* notation, where processes are depicted as boxes and objects as wires. We depict discarding using a black dot.

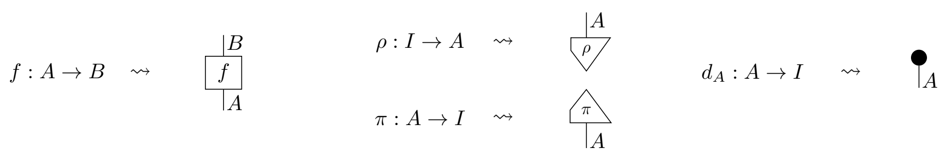


Note that we don't require *a priori* that the discarding maps satisfy any equations aside from the basic compatibility with $\otimes$. They play an important role, however, in identifying certain families of well-behaved maps within a process theory. The most important such condition is the following.

**Definition 2.** A map $f : A \to B$ is called *causal* if $d_B \circ f = d_A$, or diagrammatically:

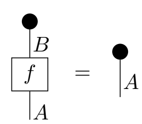


(1)

Intuitively, causality captures the fact that the only influence a map can have is on its "future," i.e., its output. If the output is discarded, then the actual causal process that took place is irrelevant.

Our main examples of process theories are $\mathbf{Mat}[\mathbb{R}_+]$ and $\mathbf{CPM}$, which contain (finite-dimensional) classical probability theory and quantum theory, respectively.

**Example 1.** The process theory $\mathbf{Mat}[\mathbb{R}_+]$ has as objects natural numbers and as morphisms $M : m \to n$ the $n \times m$ matrices whose entries are non-negative real numbers. The monoidal product is given by tensor product of matrices (a.k.a. Kronecker product), whose unit is the $1 \times 1$ 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 1's. 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 $P(i|j) := M_{ij}$).

**Example 2.** The process theory $\mathbf{CPM}$ has as objects finite-dimensional Hilbert spaces $\mathcal{H}, \mathcal{K}, ...$ 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 $L(\mathbb{C}) \cong \mathbb{C}$. States $\rho : \mathbb{C} \to L(\mathcal{H})$ are 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.

We will furthermore find it convenient to assume that our process theory has a (self-dual) compact structure, meaning that 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:

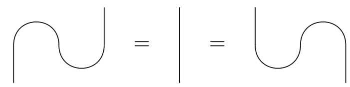


This structure enables us easily to represent higher-order maps as first order ones. For example, we can represent a process that takes processes of type $A \to A'$ and produces processes of type $B \to B'$ as a normal, first-order process $f : B \otimes A' \to A \otimes B'$. We then indicate its higher-order interpretation by drawing $f$ as a box with a “hole” in it, and use cups and caps to define “plugging” another box into that hole:

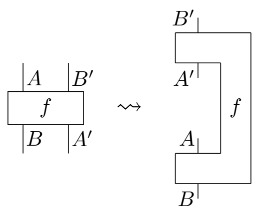


(2)

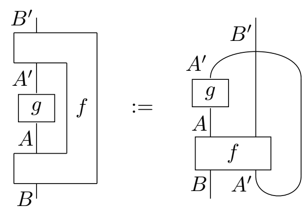


(3)

In [JKZ19], the authors furthermore assumed the structure of a CDU category–a minor variation on the notion of a Markov category [Fri20]–which captures an abstract notion of probabilistic maps by assuming every object carries a “copying” (a.k.a. “broadcasting”) map [CS12]. In particular, this allows one to capture causal models based on Bayesian networks as certain functors between CDU categories.

The famous no-cloning/no-broadcasting theorems of quantum theory, however, rule out a Markov-like structure in the category **CPM** of quantum maps. Hence, we adopt a weaker notion of causal model, consisting of a formal string diagram (i.e., a morphism in the free category over a signature) and an interpretation of that diagram into a concrete process theory (of, e.g., probabilistic or quantum maps).

# 3 Interventional causal models

A causal model consists of two parts: (i) a formal string diagram capturing our causal hypotheses, and (ii) an associated interpretation in a concrete process theory (i.e., $\mathbf{Mat}[\mathbb{R}_+]$ or **CPM**).

We define a formal string diagram as a morphism of a particular form in the free symmetric monoidal category $\mathbf{Free}(\Sigma)$ over some signature $\Sigma$. For a fixed set of objects $\{X_1, \ldots, X_n\}$ in $\Sigma$, we call a diagram $D : X_1 \otimes \ldots \otimes X_n \to X_1 \otimes \ldots \otimes X_n$ a *circuit with holes* if it is a morphism in the free symmetric monoidal category and furthermore has the property that joining each input $X_i$ to its corresponding output $X_i$ yields another morphism in the free SMC (i.e., it doesn’t introduce a directed cycle).

The intuition is that each of the input/output pairs is a “hole” in the diagram, which we call an *intervention locus*, or simply *locus* (plural loci), where a local process can be plugged in. For example:

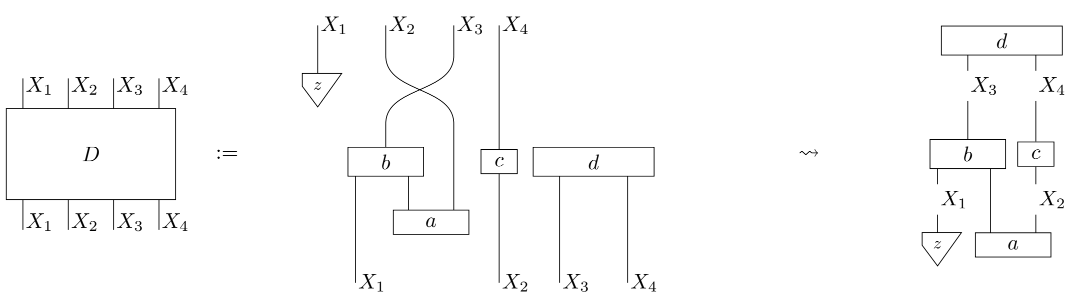


As do Barrett, Lorenz, and Oreshkov [BLO20], we require a locus’s input and output system-types to be identical in order to accommodate the special “trivial intervention,” which joins a locus’s input and output with an identity wire.

We can now introduce a notion of causal model that is similar in spirit to that of [JKZ19], but no longer relies on the CDU structure needed to capture Bayesian networks.

**Definition 3.** For any process theory $\mathcal{C}$, an interventional causal model consists of a pair $(D, \Phi)$ where $D$ is a circuit-with-holes in $\mathbf{Free}(\Sigma)$, $\Phi$ is a causal process in $\mathcal{C}$, and there exists a symmetric monoidal functor $F : \mathbf{Free}(\Sigma) \to \mathcal{C}$ such that $F(D) = \Phi$.

When $\mathcal{C} = \mathbf{Mat}[\mathbb{R}^+]$ we call $(D, \Phi)$ a *classical interventional causal model*, whereas when $\mathcal{C} = \mathbf{CPM}$, we call it a *quantum interventional causal model.*

**Definition 4.** In a classical or quantum interventional causal model with loci $X_1, \ldots, X_n$, the *interventional channel from $X_i$ to $X_j$* is the process obtained by filling in all loci other than $X_i$ and $X_j$ with identity interventions, and inputting a normalized, i.e., causal, state to the wire leaving locus $X_j$.

The interventional channel is a process of the form

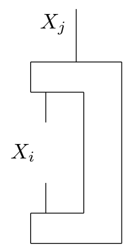


which maps intervention–or non-deterministic intervention outcome–$f : X_i \to X_i$ at locus $X_i$ to the state on system $X_j$ resulting from the combination of intervention $f$ at $X_i$ and trivial (identity) interventions at all loci other than $X_i$ and $X_j$. In particular, the interventional channel gives the consequence for $X_j$ of forcibly setting the state leaving $X_i$ to $\psi$:

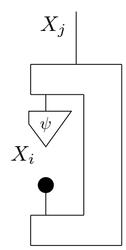


Thus the interventional channel yields what in ordinary causal modeling is called an "interventional distribution," or "the causal effect of $X_i$ on $X_j$." The shift in focus from distributions to channels is in line with a trend toward channel-based accounts of probabilistic reasoning [JZ18; CJ19], and facilitates the study of conditional actions and stochastic policies.

A key commonality between the classical and quantum processes studied in this work is that they can be completely specified by the numbers that result when they are locally composed with states and effects.

**Proposition 1.** *The theories $\mathbf{Mat}[\mathbb{R}_+]$ and $\mathbf{CPM}$ have local process tomography: any process*

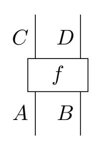


*is determined by numbers*

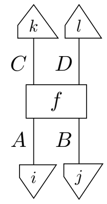


(4)

*where $i, j, k,$ and $l$ index finite, informationally complete sets of states or effects on the appropriate system-types.*

We call the set of numbers in equation (4) the *generalised matrix elements* associated with a process $f$. A local process tomography protocol for causal–i.e., probability-preserving–maps in $\mathbf{Mat}[\mathbb{R}_+]$ and **CPM** uses observed probabilities of combinations of measurement outcomes conditioned on combinations of causal state preparations. In the quantum case, though one cannot obtain all of the generalised matrix elements using a single choice of measurement basis, it is always possible to obtain them from the measurement statistics of multiple projective measurements. Local process tomography for comb-shaped quantum processes corresponds to Ried et al.'s [Rie+15] "causal tomography."

Local process tomography for a classical or quantum interventional model typically relies on probabilities given by filling intervention loci with maps of the form

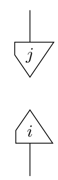


where $i$ and $j$ index the same set of states (and their adjoints), but may be distinct for a single outcome. An outcome with distinct $i$ and $j$ arise in an experiment when an observation is recorded and then a new state is prepared according to the will of the experimenter.

In contrast, what we call observational, or non-experimental, data arise when the only outcomes of this form that can be implemented are those satisfying $i = j$: there is no possibility of recording a locus's incoming state but then feeding forward a different state.

The general problem of causal identification is to use qualitative assumptions about the causal scenario to compute quantitative causal influences given statistics from only a highly restricted set of interventions. Usually the allowed interventions are "passive observations," which non-deterministically implement the aforementioned "observational" outcomes, and thereby teach the observer the probabilities of those outcomes. There is no quantum process appropriately called passive observation; for the purposes of this paper, the quantum interventions allowed as "observations" are exactly the projective measurements, which include identity processes (totally uninformative measurements). Thus the observational outcomes whose joint probabilities are available for inference include identity maps and maps composed of an effect followed by its adjoint.

The class of actual deterministic processes–i.e., the measurement processes–that non-deterministically result in what we call observational outcomes is closely related to a criterion, called informational symmetry, whereby Ried et al. characterized certain interventions in both classical and quantum causal scenarios as mere observations.[Rie+15] (Informational symmetry depends on both the intervening process and the prior state. Here we desire a criterion applying only to the intervening process itself.) Note that (aside from the trivial measurements) the measurement processes themselves, which are implemented with certainty, are not part of our process theories. Instead we depict the individual, generally non-deterministic outcomes.

To apply our proofs of sufficient conditions for identifiability to the classical stochastic setting, we need not characterize the complete classical stochastic analogue of the quantum class of observation outcomes, but only posit that classical observation outcomes include identity matrices and matrices that have all zero entries except 1 in a single position on the main diagonal. (When classical probability theory is viewed as a sub-theory of quantum theory, what we call classical observation outcomes are in fact identified with projective measurement outcomes.) The latter kind of matrix represents an outcome of what is normally called "observing a random variable." By marginalization, identity interventions in the classical setting can be simulated from the probabilities of such projections onto pure causal states (point distributions). Thus our proofs of classical identifiability really appeal to no intervention procedures other than ordinary maximally informative classical observation.

A stochastic or quantum interventional channel, respectively, will be called identifiable from an abstract string diagram if for any *positive* stochastic or quantum model of the string diagram, the interventional channel can be computed from the probabilities of arbitrary combinations of observation outcomes at all intervention loci of the model. Positivity is defined as follows:

**Definition 5.** A positive stochastic or quantum interventional model is a model whose composition with any non-zero state and any non-zero effect gives a strictly positive number.

The states and effects composed with a model may in particular be products of those implemented at individual intervention loci. For a positive model, therefore, any combination of observational outcomes occurs with non-zero probability. The positivity condition in our process-theoretic account serves the same purpose as the common requirement in ordinary causal modeling that a probabilistic causal model induce a strictly positive joint distribution on all variables. Positivity ensures that all relevant conditional probabilities are defined, and that detecting an arbitrary state at a locus after intervening at another locus is at least possible– if it were not, asking for the corresponding interventional probability would make no sense. The definition of identifiability from an abstract string diagram captures the notion that the assumptions of no direct influence between loci disconnected in the abstract diagram–equivalent to assumptions of absence of certain arrows in a directed acyclic graph representing a classical [Pea09] or quantum [CS16] causal structure–suffice for inference: in any model satisfying at least the constraints implied by the string diagram, the quantity in question can be deduced from observational outcome statistics.

Circumscribing the class of allowed interventions raises the question of whether the restrictions are strong enough to rule out schemes like causal tomography that would always allow causal identification. The answer is affirmative, as Ried et al. [Rie+15] noted, and is evident from string diagrams like

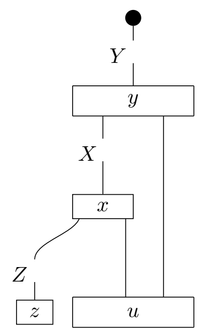


for which the interventional channel

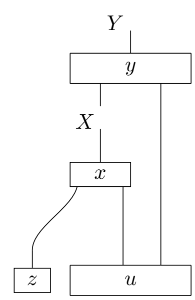


from $X$ to $Y$ is not identifiable. Two models yielding different interventional channels but identical observational outcome statistics are constructed via functorial interpretation according to Definition 3: in both models, $u$ is interpreted as the Bell state $|\Psi^+\rangle$ on two qubits, $z$ as a fixed quantum state with full support (i.e., a state whose composition with any non-zero effect is non-zero), and $x$ as the quantum map that discards its left-hand input and outputs its right-hand input unchanged. In the first model, $y$ is interpreted as the map that discards its right-hand input and applies to its left-hand input a projective measurement followed by a depolarizing channel with parameter $\lambda$. In the second model, $y$ is interpreted as the map that discards its left-hand input and applies to its right-hand input the same projective measurement followed by the same depolarizing channel as in the first model. The interpretations of $y$ in the two models are

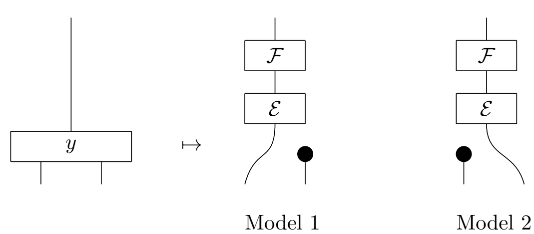

where

$$\mathcal{E}(\rho) = |0\rangle\langle 0|\rho|0\rangle\langle 0| + |1\rangle\langle 1|\rho|1\rangle\langle 1|$$
$$\mathcal{F}(\sigma) = (1-\lambda)\sigma + \lambda(\frac{1}{2}|0\rangle\langle 0| + \frac{1}{2}|1\rangle\langle 1|).$$

These models are positive because $z$ has full support, the reduced state arriving at $X$ is maximally mixed, and the state arriving at $Y$ includes a maximally mixed state with weight $\lambda$ no matter what outcome has occurred at $X$. The interventional channel is not identifiable from the abstract string diagram because it cannot be computed for *every* positive model. The two-variable quantum identification schemes of [Rie+15], however, might sometimes help in three-variable situations–perhaps some scenarios that involve coherence or entanglement and also include an instrumental variable, which would correspond here to the locus $Z$.

Because $Z$ does not influence $X$ or $Y$ in these models, this example is essentially equivalent to those two-variable scenarios in [Rie+15] for which the desired interventional channel was noted to be unidentifiable. We show three variables to detach our example from the conceptually unique two-variable case–for which classical stochastic causal effects are in a sense never identifiable. Moreover, we explicitly note that identification is impossible even for positive models of the string diagram.

# 4 Front-door scenarios

It is generally impossible to tell from observational data whether two correlated random variables, one of which is known *not* to be a descendant of the other–i.e., one of which comes "after" the other–stand in a cause-effect relation or are instead descendants of an unobserved common cause. If, however, there is a third observed variable along the possible path of causal influence between the first two, such inference may be possible. This is the content of the "front-door criterion" for causal identifiability. The criterion has a quantum analogue, the simplest instances of which, along with their classical counterparts, are captured by the following result, derived using common features of the quantum and classical process theories.

In this and the following section, each system represented by an uppercase letter may be a composite of multiple smaller systems, and similarly each box could be a composite of smaller boxes. Thus, a single locus in one of our diagrams might correspond to a list of several classical variables or quantum laboratories [CS16] occupying several nodes of a more traditional causal diagram. What we call an intervention at a locus would then correspond to (possibly choreographed) interventions at all those nodes.

**Proposition 2.** *For quantum or stochastic models of a string diagram*

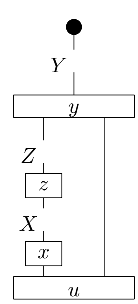


(5)

*the interventional channel*

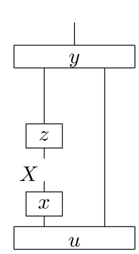


*from $X$ to $Y$ is identifiable.*

*Proof.* First, we compute the process $z$, determined by its generalised matrix elements, which we obtain by introducing a non-zero scalar factor and its inverse (where the inverse is indicated by a diagram inside $\{-\}^{-1}$), then using the causality equation (1) to transform into the following quantity:

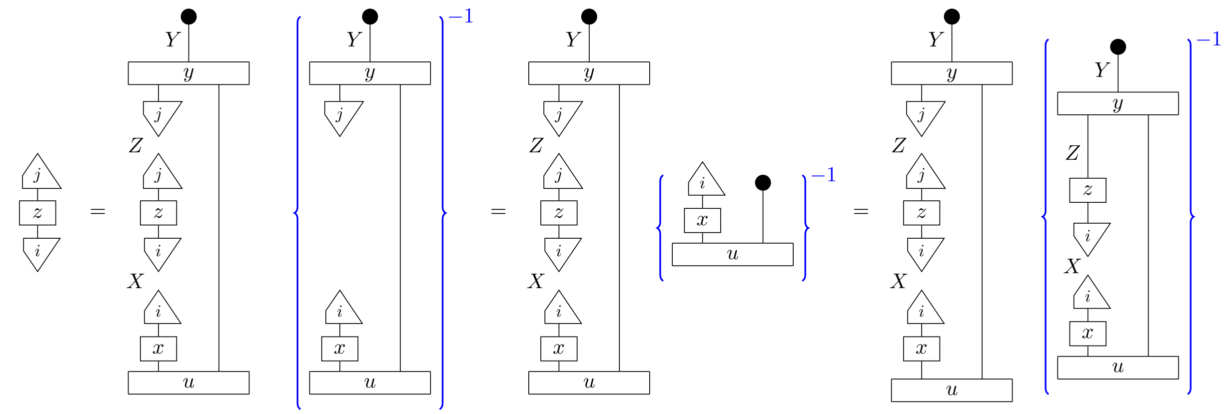


Note that the scalars being inverted are indeed non-zero, by positivity of the whole interventional model. Furthermore, the rightmost diagram above consists of quantities that can be computed purely from projective measurements at all of the loci (including the identity/trivial measurement at $Z$).

Once we have computed the generalised matrix elements of $z$, we can use them to compute those of the outer comb-shaped process by adjusting for $z$:

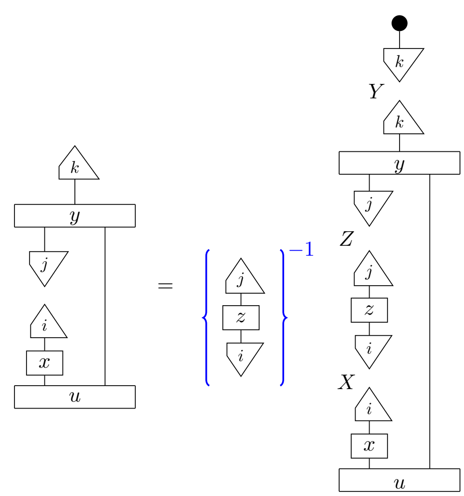


Finally, we compose the two processes at $Z$, leaving the $X$ input and output, to obtain the interventional channel. □

Thus, in the quantum just as in the classical case, observation at a locus $Z$ lying on the path between $X$ and $Y$ "blocks" that path and allows control of the confounding influence of $u$.

## 5 A more general case of a single intervention

The identification criterion of Proposition 2 can be generalised, using the same proof technique, to a quantum version of Jacobs, Kissinger, and Zanasi's Theorem 8.1 [JKZ19].

**Proposition 3.** *For quantum or stochastic models of a string diagram*

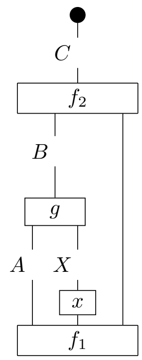


*the interventional channel*

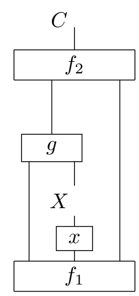


*from X to C is identifiable.*

*Proof.* First, we compute the generalised matrix elements of $g$, similarly to before:

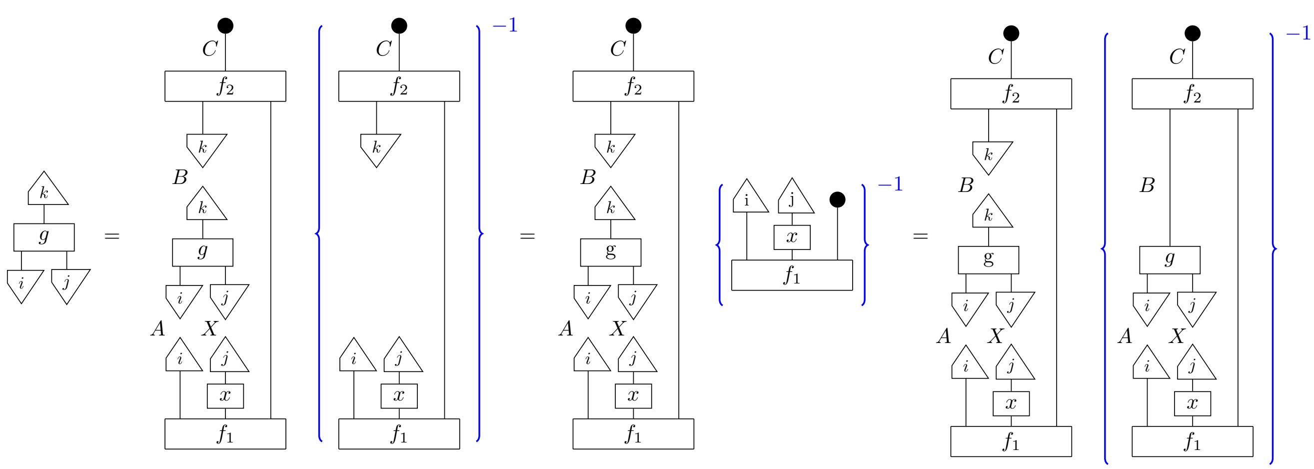

Once we have the generalised matrix elements for $g$, we can again generate those of the outer comb by adjusting for $g$:

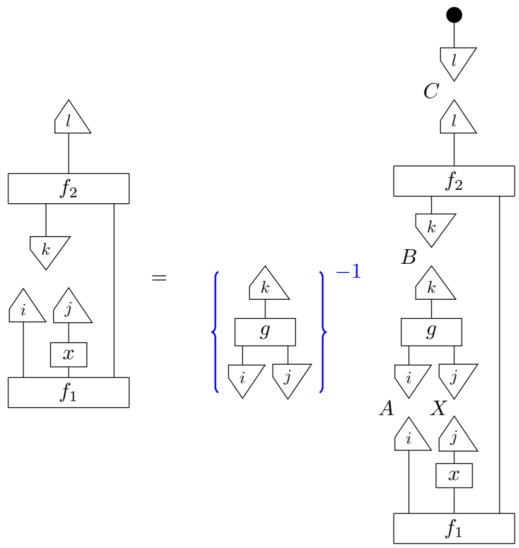


Finally, we compose the two processes at $A$ and $B$, leaving the $X$ input and output, to obtain the interventional channel. □

# 6 Conclusion

Any abstract, categorical account of established theory raises the question of what use may be made of the translation to the abstract framework. The present article answers that question with regard to the categorical study of causal inference initiated in [JKZ19]. The framework generalises the problems and facts of causal inference in such a way that the generalisations are immediately applicable to the quantum realm, in which the subject has been studied little. In addition to constituting new domain knowledge, the quantum results obtained from the process-theoretic treatment promote confidence in the usefulness of the logical foundations of causal inference presented in [JKZ19], foundations which may even be relied upon in the course of further work in ordinary probabilistic reasoning.

Future work will first explicate the relationship between the interventional models presented here and the graph-based quantum causal models of [CS16] and [BLO20], with a view to stating full quantum analogues of sufficient graphical conditions for identifiability from [GP95] and [TP02]. Next, the results here can be generalised to cases of multiple interventions at loci that may be causally related, as suggested in [JKZ19]. Moreover, there is the problem of establishing necessary conditions for identifiability, and comparing such conditions to necessary conditions from the statistical causal inference literature. The overall research program is to discover exactly how much of causal inference is about shapes of diagrams in a syntactic category.

**Acknowledgements.** The authors would like to thank Rob Spekkens, Jon Barrett, and Frederick Eberhardt for useful discussions about causal structures, observation, common causes, control of confounding, and categories in computer science. AK acknowledges support of Grant No. 61466 from the John Templeton Foundation as part of the QISS project. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation.

# References

[GP95] David Galles and Judea Pearl. "Testing Identifiability of Causal Effects". In: *UAI*. 1995.

[TP02] Jin Tian and Judea Pearl. "A general identification condition for causal effects". In: *AAAI/IAAI*. 2002.

[AKP06] Dorit Aharonov, Alexei Kitaev, and John Preskill. "Fault-Tolerant Quantum Computation with Long-Range Correlated Noise". In: *Phys. Rev. Lett.* 96 (5 Feb. 2006), p. 050504. DOI: 10.1103/PhysRevLett.96.050504. URL: https://link.aps.org/doi/10.1103/PhysRevLett.96.050504.

[CDP08] G. Chiribella, G. M. D'Ariano, and P. Perinotti. "Quantum Circuit Architecture". In: *Phys. Rev. Lett.* 101 (6 Aug. 2008), p. 060401. DOI: 10.1103/PhysRevLett.101.060401. URL: https://link.aps.org/doi/10.1103/PhysRevLett.101.060401.

[Pea09] Judea Pearl. *Causality*. 2nd ed. Cambridge University Press, 2009. DOI: 10.1017/CBO9780511803161.

[CS12] Bob Coecke and Robert W Spekkens. "Picturing classical and quantum Bayesian inference". In: *Synthese* 186.3 (2012), pp. 651–696.

[Rie+15] Katja Ried et al. "A quantum advantage for inferring causal structure". en. In: *Nature Physics* 11.5 (May 2015). Bandiera_abtest: a Cg_type: Nature Research Journals Number: 5 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Quantum information;Quantum optics;Single photons and quantum effects Subject_term_id: quantum-information;quantum-optics;single-photons-and-quantum-effects, pp. 414–420. ISSN: 1745-2481. DOI: 10.1038/nphys3266. URL: https://www.nature.com/articles/nphys3266 (visited on 09/20/2021).

[CS16] Fabio Costa and Sally Shrapnel. "Quantum causal modelling". In: *New Journal of Physics* 18.6 (June 2016), p. 063032. DOI: 10.1088/1367-2630/18/6/063032. URL: https://doi.org/10.1088/1367-2630/18/6/063032.

[All+17] John-Mark A. Allen et al. "Quantum Common Causes and Quantum Causal Models". In: *Phys. Rev. X* 7.3 (July 2017). Publisher: American Physical Society, p. 031021. DOI: 10.1103/PhysRevX.7.031021. URL: https://link.aps.org/doi/10.1103/PhysRevX.7.031021.

[CK17] 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.

[JZ18] Bart Jacobs and Fabio Zanasi. "The Logical Essentials of Bayesian Reasoning". In: *CoRR* abs/1804.01193 (2018). arXiv: 1804.01193. URL: http://arxiv.org/abs/1804.01193.

[CJ19] Kenta Cho and Bart Jacobs. "Disintegration and Bayesian inversion via string diagrams". In: *Mathematical Structures in Computer Science* 29.7 (2019), pp. 938–971. DOI: 10.1017/S0960129518000488.

[CTB19] Juan D Correa, Jin Tian, and Elias Bareinboim. "Identification of causal effects in the presence of selection bias". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 33. 01. 2019, pp. 2744–2751.

[JKZ19] Bart Jacobs, Aleks Kissinger, and Fabio Zanasi. "Causal Inference by String Diagram Surgery". In: *Foundations of Software Science and Computation Structures*. Ed. by Mikołaj Bojańczyk and Alex Simpson. Cham: Springer International Publishing, 2019, pp. 313–329. ISBN: 978-3-030-17127-8.

[BLO20] Jonathan Barrett, Robin Lorenz, and Ognyan Oreshkov. "Quantum Causal Models". In: *arXiv:1906.10726 [quant-ph]* (Nov. 2020). arXiv: 1906.10726. URL: http://arxiv.org/abs/1906.10726 (visited on 10/15/2021).

[Fri20] Tobias Fritz. "A synthetic approach to Markov kernels, conditional independence and theorems on sufficient statistics". In: *Advances in Mathematics* 370 (2020), p. 107239.