

## A graph-separation theorem for quantum causal models

This content has been downloaded from IOPscience. Please scroll down to see the full text.

2015 New J. Phys. 17 073020

(<http://iopscience.iop.org/1367-2630/17/7/073020>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 131.130.87.134

This content was downloaded on 09/08/2016 at 07:42

Please note that [terms and conditions apply](#).



## PAPER

## A graph-separation theorem for quantum causal models

Jacques Pienaar<sup>1,2</sup> and Časlav Brukner<sup>1,2</sup><sup>1</sup> Faculty of Physics, University of Vienna, Boltzmanngasse 5, A-1090 Vienna, Austria<sup>2</sup> Institute of Quantum Optics and Quantum Information, Austrian Academy of Sciences, Boltzmanngasse 3, A-1090 Vienna, AustriaE-mail: [jacques.pienaar@univie.ac.at](mailto:jacques.pienaar@univie.ac.at)

Keywords: causality, bayesian networks, quantum foundations

## RECEIVED

27 November 2014

## REVISED

27 April 2015

## ACCEPTED FOR PUBLICATION

26 May 2015

## PUBLISHED

14 July 2015

Content from this work  
may be used under the  
terms of the [Creative  
Commons Attribution 3.0  
licence](#).

Any further distribution of  
this work must maintain  
attribution to the  
author(s) and the title of  
the work, journal citation  
and DOI.



## Abstract

A causal model is an abstract representation of a physical system as a directed acyclic graph (DAG), where the statistical dependencies are encoded using a graphical criterion called ' $d$ -separation'. Recent work by Wood and Spekkens shows that causal models cannot, in general, provide a faithful representation of quantum systems. Since  $d$ -separation encodes a form of Reichenbach's common cause principle (RCCP), whose validity is questionable in quantum mechanics, we propose a generalized graph separation rule that does not assume the RCCP. We prove that the new rule faithfully captures the statistical dependencies between observables in a quantum network, encoded as a DAG, and reduces to  $d$ -separation in a classical limit.

## 1. Introduction

An essential problem faced by any scientist trying to make sense of the world is this: how do we infer causal relationships between the observed quantities, based only on information about their statistical dependencies? This problem is well known to statisticians and researchers working on artificial intelligence (AI), who have developed *causal models* as a tool for making causal inferences from a set of observed correlations. In most practical situations, the task is made easier by the availability of additional information and physical intuition. For example, in considering possible explanations for the observed correlation between smoking and cancer, we might consider it plausible that the two are independently caused by a common genetic factor, but few people would advocate the idea that having cancer causes people to smoke—not least because smoking tends to precede the onset of cancer, and we know that an effect cannot precede its cause. If we are simply told that two abstract variables  $X$  and  $Y$  have correlated values, the task is much more difficult. Such situations arise in theoretical work where one aims to relax the existing framework and construct more general models, or in practical applications like programming an AI to make causal inferences about data that it acquires.

In a causal model, defined in section 2, the random variables of interest are represented by nodes and causal influences between them are represented by lines with arrows, called directed edges. The laws of physics require that no effect can be its own cause, leading to the requirement that the graph be acyclic (i.e. free of directed loops). The resulting directed acyclic graph (DAG) provides a computationally useful tool for extracting information about the statistical relationships of variables. In particular, it allows us to determine whether one set of variables is independent of any other set, conditional on the values of a third set. This information can be obtained directly from the graph using a simple algorithm, based on a concept called  $d$ -separation. Two sets of variables will be independent conditional on a third set of variables if and only if they are  $d$ -separated by the third set in the graph.

The proof that  $d$ -separation allows one to extract all (and only) correct information about a list of conditional independence (CI) statements makes causal models particularly powerful tools for representing physical systems. Indeed, we are tempted to interpret the causal structure represented by the graph as 'out there in the world' in the same sense as we can take the classical space-time manifold (which encodes causal relations between events) to be an independent and objectively defined structure. However, the program faces significant conceptual difficulties when one attempts to apply it to quantum physics. In principle, any observed probability distribution can be explained by some causal model, if we allow the possibility of hidden variables. However, as

first clearly articulated by Bell [1], hidden-variable accounts of quantum mechanics can be challenged because they imply highly nonlocal behaviour of the model. This feature manifests itself in causal models in the form of *fine-tuning*, where one is forced to posit the existence of causal effects between variables whose statistics are independent. The fact that causal models of quantum systems require fine-tuning was recently shown by Wood and Spekkens [2].

These considerations revive an old question: what does causality really mean in the context of quantum mechanics? Do we accept that there exist nonlocal hidden variables whose direct influence is in principle unobservable at the statistical level? Or could it be that the classical concept of causality does not extend to quantum systems, and that we need a completely new way of determining whether two quantum events are causally related? Following the latter point of view, we define a causal model based on quantum networks and use it to derive a graph separation rule analogous to  $d$ -separation, for obtaining the CI relations between variables. Our approach differs from previous work that assigns quantum amplitudes to the nodes in the DAG [3], or that aims to replace the conditional probabilities at the nodes with some appropriate quantum analog [4]. Instead, we retain classical probability theory, but seek a physically motivated graphical representation of the causal structure that gives rise to the probability distributions predicted by quantum mechanics. Our approach is more closely aligned with previous work in which quantum network diagrams are used to obtain joint probabilities obeying standard probability theory [5–11]. Particularly relevant is the recent work by Fritz [9], in which a DAG representation of quantum correlations is proposed that encompasses our concept of a quantum causal model (QCM) as will be discussed in section 3. Our work takes the additional step of defining a specific representation and a graph separation rule within this framework.

Recently, another DAG representation for general networks was proposed by Henson *et al* [12] in which  $d$ -separation continues to hold between the observed variables representing classical data. This is achieved by adding extra nodes to the graph representing ‘unobserved’ variables, which ensure that the restriction of the CI relations to just the observed nodes produces the conditional independencies expected of a quantum network (or generalized probabilistic theory). Our approach differs from these authors, in that we consider all nodes to be in principle observable; this leads us instead to modify the criterion for obtaining the CI relations from the graph (see section 3.3). The comparison to [12] will be discussed further in section 4.

The paper is organized as follows: in section 2 we give a review of the relevant concepts concerning classical causal models (CCMs) and their graphical representation by DAGs. We include a discussion of the physical motivation for these models, and the meaning of the result in [2] that such models cannot faithfully represent quantum correlations. In section 3 we aim to find such a faithful representation by re-interpreting the DAG as a quantum network. We thereby derive a new graph separation rule that does not obey the version of ‘Reichenbach’s common cause principle (RCCP)’, which holds in the classical case, but instead obeys a weaker property we call the ‘quantum causality condition’. We show that the  $d$ -separation can be recovered in a suitably defined classical limit, and we note that super-quantum correlations exceeding Tsirelson’s bound still cannot be explained by our model without fine-tuning. We conclude in section 4 with a discussion about the physical interpretation of the result and possible directions for future work.

## 2. Review of CCMs

In this section, we review the basic definition of a causal model, here referred to as a CCM to emphasize that it is tied to physical assumptions motivated by classical systems. For more details on causal models and inference, see the book by Pearl and references found therein [13].

Before discussing the formal elements of these models, let us briefly recap their historical motivation. In science and statistics, one is often faced with the task of determining the causal relationships between random variables, given some sample data. We might observe that two variables are correlated, but this fact alone does not indicate the direction of the causal influence. If we are limited in our resources, we would like to know which set of follow-up experiments will most efficiently identify the direction of the causal influences, and which causal information can already be deduced from the existing data. Correlations between variables can be represented graphically, for example, we can require that  $X$  be independent of  $Y$  conditional on a set  $Z$  whenever the removal of the nodes  $Z$  and their connections from the graph renders the sets of nodes  $X$  and  $Y$  disconnected in the resulting graph. Such a rule for obtaining independence relations from a graph is referred to as a ‘graph separation rule’. The lists of CI relations are called *semi-graphoids* and they satisfy certain axioms, described in section 2.2.

Correlations can be regarded as restrictions on the possible causal relationships between the variables. These causal relations determine how the observed statistics change after an intervention on a system. When an external agent intervenes to change the probability distribution of some of the variables at will, the distributions of the remaining variables will be updated depending on the direction of the causal influences between them and

the manipulated variables; flicking a switch can cause a light to turn off, but extinguishing the light by other means will not affect the position of the switch. Causal information tells us more about the statistical relationships between the variables than can be obtained from correlations alone. It is therefore useful to design a graphical representation and a graph separation rule that captures causal information, not just correlations.

The directions of causal influences are represented by adding arrows to the edges in the graph. This supplements the information about correlations with further constraints on the conditional independencies. Every causal graph, up to an absolute ordering of the variables, is in one-to-one correspondence with a list of CI relations called a *causal input list*. The list can be thought of as guidelines for generating a probability distribution: one begins by generating the values of the independent variables, then computes the values of any variables that depend directly on them, then the variables that depend on those, and so forth. Hence every causal graph can be taken to represent a family of stochastic physical processes proceeding over many time steps. In practice, working with the causal input list can be cumbersome, so it is more efficient to obtain the conditional independencies directly from the graph using a graph separation rule called *d-separation*. In this work we will propose to upgrade the definitions of causal input list and *d-separation* to quantum systems.

## 2.1. Notation

Random variables, or sets of random variables, are denoted by capital roman letters, e.g.  $X$ ,  $Y$ ,  $Z$ , which take values from a set of possible outcomes. If  $\mathcal{E}_X$  is the space of all possible outcomes of  $X$ , let  $P(X)$  represent a probability distribution on  $\mathcal{E}_X$  and  $P(X = x)$  the probability that the variable  $X$  takes the value  $x \in \mathcal{E}_X$ . In many cases, we will use the term  $P(X)$  also to represent  $P(X = x)$ , except where it might cause confusion. The statement  $X = P(X)$  means that the random variable  $X$  is distributed over its outcomes according to the distribution  $P(X)$ . The joint probability  $P(X, Y, Z)$  represents a probability distribution over all the possible values of the random variables  $\{X, Y, Z\}$ . The *conditional* probability  $P(X|Y)$  is a set of probability distributions defined on  $\mathcal{E}_X$ , for the possible values of  $Y$ . Given a joint distribution  $P(X, Y)$ , a *marginal* probability  $P(Y)$  is defined by summing over all possible values of the other variables, i.e.

$$\begin{aligned} P(Y) &= \sum_X P(X, Y) \\ &:= \sum_{x \in \mathcal{E}_X} P(X = x, Y). \end{aligned} \quad (1)$$

These concepts are united by the law of total probability, which states that  $P(X, Y) = P(X|Y)P(Y)$ . Unless otherwise specified, we consider only variables with discrete outcome spaces.

## 2.2. Formal definitions for causal models

Let us consider a set of random variables whose values are governed by some joint probability function and which in general may be correlated. Formally, the statistical dependencies between variables are given by their CI relations:

**Definition 1. CI relations.** Let  $X$ ,  $Y$ ,  $Z$  be three disjoint sets of variables. The sets  $X$  and  $Y$  are said to be *conditionally independent* given  $Z$  if knowing  $Y$  provides no new information about  $X$  given that we already know  $Z$  (i.e.  $Z$  ‘screens-off’  $Y$  and  $X$  from each other). We write this as  $(X \perp Y|Z)$ , which stands for the assertion that  $P(X|Y, Z) = P(X|Z)$ . We will often use the shorthand  $(X \cup W \perp Z|Y) := (XW \perp Z|Y)$  when dealing with set unions in CI relations.

Any joint probability distribution  $P$  can be conveniently characterized by the complete set of CI relations that it implies for the variables. In fact, one only needs to specify a subset of CI relations, from which the rest can be obtained using the *semi-graphoid axioms*:

*Semi-graphoid axioms:*

- 1.a. Symmetry:  $(X \perp Y|Z) \Leftrightarrow (Y \perp X|Z)$
- 1.b. Decomposition:  $(X \perp YW|Z) \Rightarrow (X \perp Y|Z)$
- 1.c. Weak union:  $(X \perp YW|Z) \Rightarrow (X \perp Y|ZW)$
- 1.d. Contraction:  $(X \perp Y|ZW) \text{ and } (X \perp W|Z) \Rightarrow (X \perp YW|Z)$ .

Note that if  $(X \perp Y|Z)$  and  $(W \perp Y|Z)$  both hold for disjoint sets  $X, Y, Z, W$ , then  $(XW \perp Y|Z)$  does not necessarily hold. This might seem counter-intuitive, but examples where it fails are easy to construct<sup>3</sup>.

The semi-graphoid axioms can be derived directly from the axioms of probability theory. The interpretation of the axioms is given by the following excerpt from Pearl [13], (chapter 1.1):

‘The *symmetry* axiom states that, in any state of knowledge  $Z$ , if  $Y$  tells us nothing new about  $X$  then  $X$  tells us nothing new about  $Y$ . The *decomposition* axiom asserts that if two combined items of information are judged irrelevant to  $X$ , then each separate item is irrelevant as well. The *weak union* axiom states that learning irrelevant information  $W$  cannot help the irrelevant information  $Y$  become relevant to  $X$ . The *contraction* axiom states that if we judge  $Y$  irrelevant to  $X$  after learning some irrelevant information  $W$ , then  $Y$  must have been irrelevant before we learned  $W$ .’

**Definition 2. Semi-graphoid closure.** Given any set  $S$  of CI relations, the *closure* of  $S$  is the set  $\bar{S}$  that includes all CI relations derivable from  $S$  using the axioms 1.a–d.

Given a joint probability distribution  $P$ , let  $\bar{C}(P)$  denote the complete closed set of CI relations obtainable from  $P$ . In general, the CI relations do not uniquely fix the probability distribution; there may exist two distinct joint distributions  $P$  and  $P'$  for which  $\bar{C}(P) = \bar{C}(P')$ . Hence, the CI relations alone do not capture the full information about the statistics. In the following, we will supplement the CI relations with a causal structure and functional relations in the form of a CCM.

A CCM provides us with an algorithm to generate the precise statistics of the observables. It can therefore be regarded as an abstract description of a physical system: if the predictions match the actual observations, then the CCM provides a possible explanation of the data. Formally, a CCM consists of two ingredients: an ordered set of CI relations  $L_\mathcal{O}$ , and a set of functions  $F$  called the *model parameters*.

**Definition 3. Causal input list.** Let  $\mathcal{O}$  be an ordering that assigns a unique integer in  $\{1, 2, \dots, N\}$  to each member of a set of  $N$  variables. Consider an ordered set of variables  $\{X_i; i = 1, 2, \dots, N\}$ , where the subset of variables  $X_j$  with  $j < i$  are called the *predecessors* of  $X_i$ . A *causal input list* is the ordered set of CI relations of the form:  $L_\mathcal{O} := \{(X_i \perp R(X_i) | \mathbf{pa}(X_i)) : i = 1, 2, \dots, N\}$ , where each set  $\mathbf{pa}(X_i)$  is a subset of the predecessors of  $X_i$  called the *parents* of  $X_i$ , and  $R(X_i)$  are the remaining predecessors of  $X_i$  excluding the parents.

**Definition 4. Ancestors and descendants.** Given a causal input list  $L_\mathcal{O}$ , consider the set of parents of  $X_i$ , their parents’ parents, and so on. These are called the *ancestors* of  $X_i$ . Similarly, the *descendants* of  $X_i$  are all variables for which  $X_i$  is an ancestor. We will use  $\mathbf{an}(X)$  to denote the union of the ancestors of a set  $X$ .

**Definition 5. Model parameters.** Given a causal input list  $L_\mathcal{O}$ , the *model parameters* are a set  $F := \{F_i; i = 1, 2, \dots, N\}$  consisting of  $N$  probabilistic functions  $F_i$ . Each  $F_i(X)$  is equivalent to applying a deterministic function  $f_i(X, U_i)$  with probability  $P(U_i)$  for some auxiliary variable  $U_i$ . The  $U_i$  are sometimes called *error variables* and by definition have no parents. Each function  $F_i$  determines the probability of  $X_i$  conditional on the values of its parents:


$$P(X_i | \mathbf{pa}(X_i)) = F_i(\mathbf{pa}(X_i)) = f_i(\mathbf{pa}(X_i), U_i) \quad \text{with probability } P(U_i). \quad (2)$$

For variables without any parents, called *exogenous variables*, the function  $F_i$  just specifies a probability distribution over the possible values of  $X_i$ , i.e.  $F_i(\emptyset) := P(X_i)$ . We assume that all exogenous variables, including any error variables  $U_i$ , are independently distributed (the Markovian assumption).

**Definition 6. Classical causal model.** A *classical causal model* on  $N$  variables is a pair  $\{L_\mathcal{O}, F\}$  containing a causal input list  $L_\mathcal{O}$  and model parameters  $F$  defined on those variables. Alternatively, a CCM can be specified by the pair  $\{G_L, F\}$ , where  $G_L$  is the graph generated by  $L_\mathcal{O}$  (see section 2.4).

Given a CCM, we can construct a joint probability by generating random variables in the order specified by  $\mathcal{O}$  and using the functions  $F$  to define the probabilities of each variable given its parents. These can then be used to construct a joint distribution from the CCM according to the law of total probability:

<sup>3</sup> Suppose that all sets represent binary variables  $\in \{0, 1\}$ , and  $Y = X \oplus W$  where  $\oplus$  is addition modulo 2. Clearly, knowledge of  $Y$  does not tell us anything about  $X$  or  $W$  individually. But knowing  $Y$  does reduce the set of possibilities for the joint set  $X \cup W$ , for example, if  $Y = 1$  then  $X \cup W$  cannot have  $X$  and  $W$  the same.



**Figure 1.** A DAG representing a simple classical causal model.  $Z$  is a direct cause of both  $Y$  and  $X$ ;  $Y$  is a direct cause of  $X$  only.

$$P(X_1, X_2, \dots, X_N) = \prod_{j=1}^N P(X_j | X_{j+1}, \dots, X_N). \quad (3)$$

The joint probability obtained in this way from a CCM  $\mathcal{M}$  is said to be *generated* by  $\mathcal{M}$  and is denoted  $P^{\mathcal{M}}$ . It satisfies the following property:

**Causal Markov condition:** given that  $P^{\mathcal{M}}$  is generated by a CCM  $\mathcal{M}$ , each variable  $X_i$  in  $P^{\mathcal{M}}$  is independent of its non-descendants, conditional on its parents.

*Note:* our definition of the causal Markov condition follows Pearl ([13]), in which it is proven to hold for any Markovian causal model (i.e. a model that is acyclic and whose exogenous variables are all independent). In the present work, a CCM is Markovian by construction, so the causal Markov condition holds. In the next section, we will use the causal Markov condition to motivate interpreting the parents of a variable as its direct causes.

**Example 1.** Consider three variables  $X, Y, Z$ . Suppose we have the ordering  $\mathcal{O}: \{X, Y, Z\} \rightarrow \{3, 2, 1\}$ , and the causal input list indicates that  $\text{pa}(X) = \{Y, Z\}$ ;  $\text{pa}(Y) = \{Z\}$ ; and  $\text{pa}(Z) = \emptyset$ . It will be shown in section 2.4 that this generates the graph shown in figure 1. Suppose the model parameters are  $F = \{f_x, f_y, f_z\}$ , where  $f_x, f_z$  are deterministic and  $F_y(A) := f_y(A, U_Y)$  with probability  $P(U_Y)$ . Then we obtain the joint probability  $P(X, Y, Z)$  as follows: first, generate the lowest variable in the ordering,  $Z$ , using the random function  $P(Z)$ . Next, generate  $U_Y$  using  $P(U_Y)$  and then apply  $f_y(Z, U_Y)$  to obtain the value of  $Y$ . Finally, use  $f_x(Y, Z)$  to obtain the value of  $X$ , the last variable in the ordering. The statistics generated by this procedure are given by:

$$P(X, Y, Z) = \sum_{U_Y} P(X|Y, Z) P(Y|Z, U_Y) P(Z) P(U_Y),$$

where

$$\begin{aligned} P(X|Y, Z) &= P(X = x | Y = y, Z = z) = \delta(x, f_x(y, z)), \\ P(Y|Z, U_Y) &= P(Y = y | Z = z, U_Y = u_y) = \delta(y, f_y(z, u_y)). \end{aligned} \quad (4)$$

### 2.3. Physical interpretation

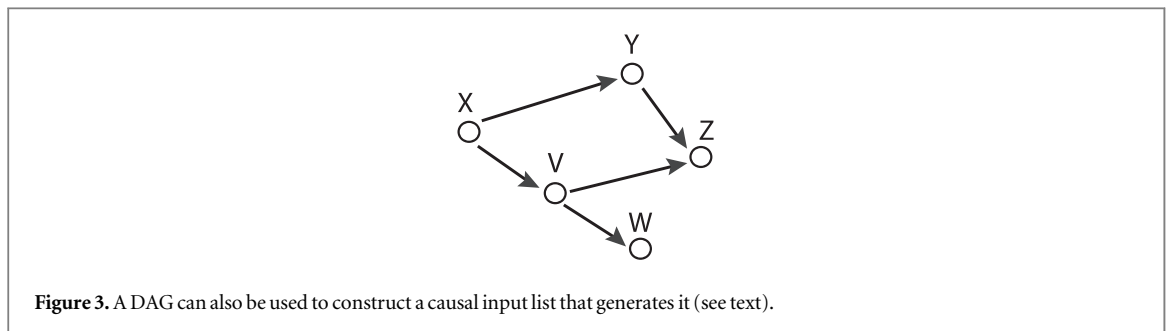
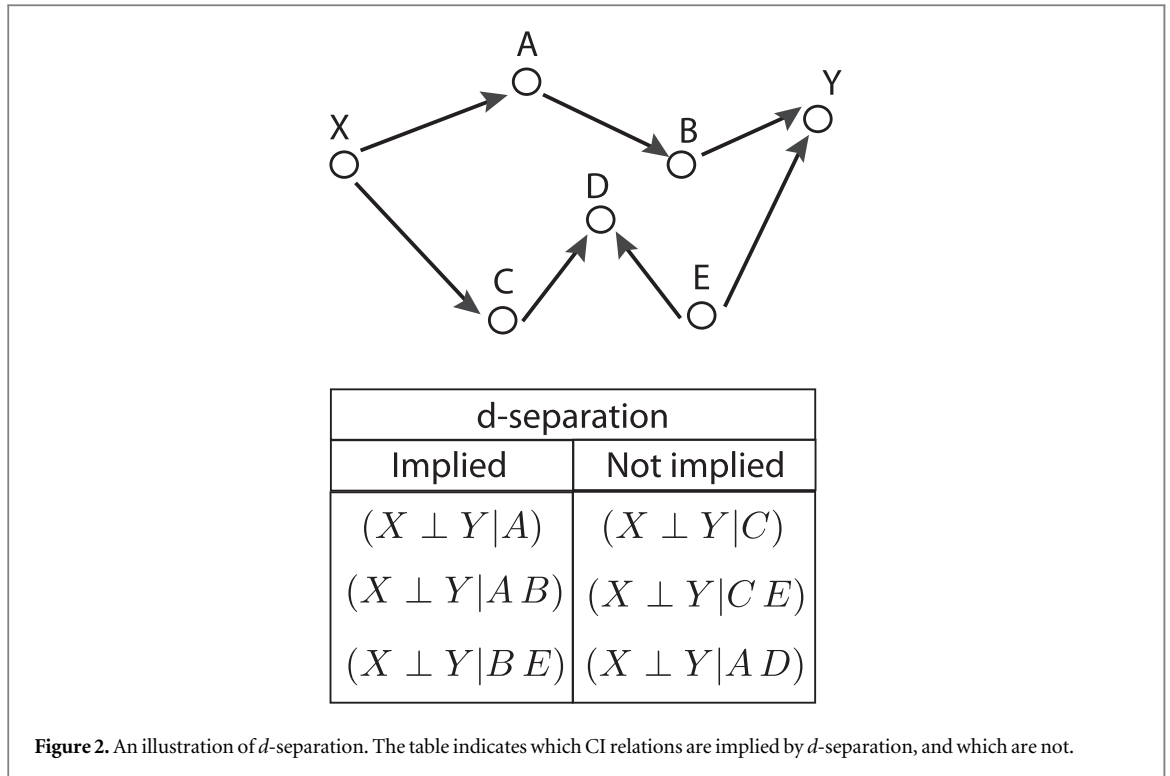
In the previous section, we gave a formal definition of a CCM and described how it generates a probability distribution over the outcomes of its random variables. Since these variables represent physical quantities, we would like to supplement this mathematical structure with a physical interpretation of a CCM, as describing the causal relationships between these physical quantities. To do so, we make the following assumption that connects the intuitive concept of a ‘direct cause’ with its mathematical representation.

**Assumption 1. A variable’s parents represent its direct causes.** Physically, we expect that knowledge of the direct causes renders information about indirect causes redundant. Hence the direct causes should *screen off* the indirect causes in the sense of definition 1. We therefore define the *direct causes* of  $X_i$  as the parents of  $X_i$  and the *indirect causes* as the remaining (non-parental) ancestors of  $X_i$ ; the screening-off property then follows from the causal Markov condition.

The above assumption leads to the following physically intuitive properties of a CCM:

**Conditioning on common effects:** in a CCM, two variables that are initially independent (i.e. conditional on the empty set) may become dependant conditional on the value of a common descendant. This reflects our intuition that two independent quantities may nevertheless be correlated if one ‘post-selects’ on a future outcome that depends on these quantities. For example, conditional on the fact that two independent coin tosses happened to give the same result, knowing the outcome of one coin toss allows us to deduce the outcome of the other.

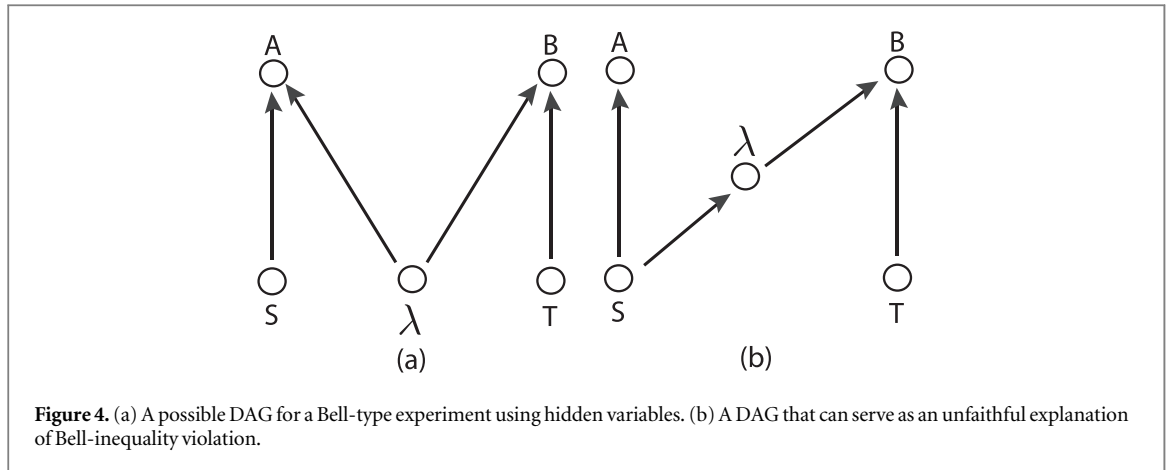




*Reichenbach's common cause principle (RCCP)*: if two variables are initially correlated (i.e. conditional on the empty set) and *causally separated* (neither variable is an ancestor of the other), then they are independent conditional on the set of their *common causes* (parents shared by both variables).

It is not immediately obvious that the RCCP follows from the causal Markov condition. For a proof using the DAG representation (discussed in the next section) see [14]. We note that there exist in the literature numerous definitions of the RCCP, so our chosen definition deserves clarification. It was pointed out in [15] that a general formulation of the principle encompasses two main assumptions. The first states that causally separated correlated variables *must* share a common cause (called the 'principle of common cause', or PCC), and the second states that the variables must be screened-off from each other by their common causes (the 'factorization principle' or FP). Our definition of the RCCP refers *only* to the factorization property FP. By contrast, [9] takes the RCCP as being equivalent to just the PCC, while the FP is regarded as a separate additional property that happens to hold for classical correlations. Note that the FP and the PCC are both consequences of the definition of a CCM—they follow directly from what we have called the assumption of Markovianity and the fact that the variables are only functionally dependant on their parents.

It is a topic of debate in the literature whether (and to what extent) Reichenbach's principle applies to quantum mechanics, particularly in light of Bell's theorem, where the factorization property often appears as an explicit assumption governing the statistics of a local hidden-variable model [2, 16]. As discussed in section 2.6, we will take the position that the violation of Bell's inequalities in experiments shows that the correlations produced by quantum entanglement cannot be given a causal explanation using the 'local' causal graph of figure 4 (a), if one also assumes that (i) there is no *fine-tuning* of the data (section 2.5); (ii) standard probability theory holds; (iii) the FP holds for quantum systems. In this work we will pursue the possibility—though it is by no means the only possibility—that it is precisely (iii) that is violated by quantum mechanics. Hence our



framework does call for a rejection of the RCCP<sup>4</sup>. Note that the PCC need not be violated, since (without conditioning on effects) two independent quantum systems can only become correlated through direct interactions, hence they share a common source.

A natural question to ask is whether some other physics principle might serve to distinguish quantum correlations from generalized probabilistic theories. Interesting work along these lines can be found in [17–20]. Here we are concerned with finding a graphical separation rule for causal models that could accommodate such relaxations in which the FP no longer applies. The reasons why a new graph-separation rule is useful in this context will be discussed further in section 2.6.

Finally, we note that one can interpret the ordering of variables  $\mathcal{O}$  as representing the time-ordering of the variables, such that each variable represents a physical quantity localized to an event in space-time. However, this interpretation is not strictly necessary for what follows. Indeed, it may be interesting to consider alternative interpretations in which some causal influences run counter to the direction of physical time, such as in the retro-causal interpretation of quantum mechanics [21].

## 2.4. The DAG representation of a CCM

It is useful to represent  $L_{\mathcal{O}}$  using a DAG, which can be thought of as the causal ‘skeleton’ of the model. In the DAG representation of  $L_{\mathcal{O}}$ , there is a node representing each variable and a directed arrow pointing to the node from each of its parents. The DAG  $G_L$  constructed in this way is said to be *generated* by the causal input list  $L_{\mathcal{O}}$ . The parents of a node in a DAG are precisely those nodes that are directly connected to it by arrows pointing towards it. It is straightforward to see that the ancestors of  $X$  are represented by nodes in the DAG that each have a *directed path* (i.e. a path along which all arrows are connected head-to-tail) leading to  $X$ , and the descendants of  $X$  are those nodes that can be reached by a directed path from  $X$ . In example 1, the system is represented by the DAG shown in figure 1.

To establish a correspondence between a DAG and the causal input list  $L_{\mathcal{O}}$  that generates it, we need an algorithm for reconstructing a list of CI relations  $L_{\mathcal{O}}$  from a DAG, such that the list generates the DAG. For this purpose, one uses *d-separation* (see figure 2).

**Definition 7. d-separation.** Given a set of variables connected in a DAG, two disjoint sets of variables  $X$  and  $Y$  are said to be *d-separated* by a third disjoint set  $Z$ , denoted  $(X \perp Y | Z)_d$ , if and only if every undirected path (i.e. a path connecting two nodes through the DAG, ignoring the direction of arrows) connecting a member of  $X$  to a member of  $Y$  is *rendered inactive* by a member of  $Z$ . A path connecting two nodes is rendered inactive by a member of  $Z$  if and only if:

- (i) the path contains a chain  $i \rightarrow m \rightarrow j$  or a fork  $i \leftarrow m \rightarrow j$  such that the middle node  $m$  is in  $Z$ , or
- (ii) the path contains an inverted fork (*head-to-head*)  $i \rightarrow m \leftarrow j$  such that the node  $m$  is not in  $Z$ , and there is no directed path from  $m$  to any member of  $Z$ .

A path that is not rendered inactive by  $Z$  is said to be *active*.

<sup>4</sup> Although the work of [9] claims not to reject Reichenbach’s principle, the author defines this principle to be essentially what we have called the PCC. If one accounts for the difference in definitions, the present work is entirely consistent with [9].



By assuming that all  $d$ -separated nodes are independent conditional on the separating set (see below), we can then obtain CI relations from the DAG. In principle, the rules for  $d$ -separation can be derived from the requirement that it produces the CI relations contained in the list  $L_{\mathcal{O}}$  that generates the DAG. However,  $d$ -separation also provides an intuitive graphical representation of the physical principles discussed in the previous section. In particular, a path between two nodes in the graph is rendered inactive by a set  $Z$  in precisely those situations where we would physically expect the two variables to be independent conditional on  $Z$ : when we are not conditioning on any common effects (head-to-head nodes); when we are conditioning on a common cause (as in the RCCP); or when we are conditioning on a node that is a link in a causal chain (screening off indirect causes). With the physical interpretation in mind, we are motivated to use  $d$ -separation to obtain CI relations using the correspondence:

$$(X \perp Y | Z)_d \Rightarrow (X \perp Y | Z), \quad (5)$$

i.e. we assume that if  $X$  and  $Y$  are  $d$ -separated by  $Z$  in a DAG, then they are conditionally independent given  $Z$  in the semi-graphoid closure of any list  $L_{\mathcal{O}}$  that generates the DAG.

Formally, let  $G$  be a DAG, and let  $C(G)$  be the set of CI relations obtainable from  $G$  using  $d$ -separation, and  $\bar{C}(G)$  the closure of this set. We then have the following theorem:

**Theorem 1 (Verma and Pearl [22]):** *let  $G_L$  be the DAG generated by  $L_{\mathcal{O}}$ . Then  $C(G_L) = \bar{L}_{\mathcal{O}}$ . That is, a CI relation is implied by the DAG  $G_L$  if and only if it is deducible from  $L_{\mathcal{O}}$  and the semi-graphoid axioms.*

Theorem 1 implies that  $d$ -separation is *sound*, since every CI relation obtainable from the DAG is in the closure of the causal input list ( $C(G_L) \subseteq \bar{L}_{\mathcal{O}}$ ), and *complete*, since there are no CI relations implied by the causal input list that are not obtainable from the DAG ( $\bar{L}_{\mathcal{O}} \subseteq C(G_L)$ ). It also follows that the set of CI relations obtainable from a DAG  $G$  is closed under the semi-graphoid axioms,  $C(G) = \bar{C}(G)$ .

A DAG  $G$  defines a *partial ordering* of the variables. A partial ordering is a weaker notion than a total ordering: given any pair of variables  $X, Y$ , it either assigns a strict ordering,  $X < Y$  or  $Y < X$ , or else leaves them unordered,  $X \sim Y$ . In the case of a DAG, if two variables connected by a directed path then they are strictly ordered according to the direction of the arrows, otherwise they are left unordered.

One can always find a causal input list that generates  $G$  as follows: choose a total ordering  $\mathcal{O}$  that is consistent with the partial ordering imposed by  $G$ , and then write down the ordered list of CI relations of the form  $(X_i \perp R(X_i) | \text{pa}(X_i))$ , where the parents of each variable are the same as the parents of its representative node in  $G$ . Moreover, the list obtained in this way is unique, modulo some freedom in the ordering of causally separated events (e.g. if the variables represent events in relativistic spacetime, this freedom corresponds to a choice of reference frame). It is this feature that allows us to replace the causal input list with its corresponding DAG in the definition of a CCM (definition 6).

**Example 2.** Consider the DAG shown in figure 3. The graph assigns parents as follows:

$$\begin{aligned} \text{pa}(X) &= \emptyset, \\ \text{pa}(V) &= X, \\ \text{pa}(Y) &= X, \\ \text{pa}(W) &= V, \\ \text{pa}(Z) &= V \cup Y. \end{aligned} \quad (6)$$

A total ordering consistent with this graph is  $\mathcal{O}: \{X, V, W, Y, Z\} \rightarrow \{1, 2, 3, 4, 5\}$ . Therefore, we obtain the causal input list:

$$L_{\mathcal{O}} = \{(X \perp \emptyset | \emptyset), (V \perp \emptyset | X), (W \perp X | V), (Y \perp \emptyset | X), (Z \perp XW | YV)\}. \quad (7)$$

Finally, we have the following useful definitions:

**Definition 8. Independence maps and perfect maps.** Given  $P(X)$  and a DAG  $G$  on the same variables, we will call  $G$  an *independence map* (I-map) of  $P$  iff  $\bar{C}(G) \subseteq \bar{C}(P)$ . If equality holds,  $\bar{C}(G) = \bar{C}(P)$ , then  $G$  is called a *perfect map* of  $P$ .

The intuition behind this definition can be understood as follows. A DAG is an independence map of a probability distribution iff every CI relation implied by the DAG is also satisfied by the distribution. That means that if two variables are not causally linked in the DAG, they must be conditionally independent in the distribution. However, the converse need not hold: the arrows in a DAG represent only the *possibility* of a causal influence. In general, depending on the choice of model parameters, it is possible for two variables to be

connected by an arrow and yet still be conditionally independent in the probability distribution. Equivalently, one can find a probability distribution that satisfies *more* conditional independencies than those implied by its DAG. A DAG is a perfect map iff it captures *all* of the CI relations in the given distribution, i.e. every causal dependence implied by the arrows in the DAG is manifest as an observed dependence in the statistics. Interestingly, there exist distributions for which no DAG is a perfect map. This fact forms the basis for the criterion of *faithfulness* of a CCM, discussed in the following section.

## 2.5. Faithful explanations and fine-tuning

Suppose we obtain the values of some physical observables over many runs of an experiment and the statistics can be modelled (to within experimental errors) by a CCM. Then we can say that the CCM provides a causal explanation of the data, in the sense that it tells us how physical signals and information might be propagating between the physical observables. In particular, it allows us to answer counterfactual questions (what would have happened if this observable had taken a different value?) and predict how the system will respond to interventions (how will other quantities be affected if a given variable is forcibly altered?). These notions can be given a rigorous meaning using causal models, and they constitute a formal framework for making causal inferences from observed data. In the present work, we will be primarily concerned with defining a QCM that can be given a useful graphical representation in DAGs, so we will not discuss interventions and inference in causal models (the interested reader is referred to [13] for inference in the classical case, and [23, 24] for attempts at quantum generalizations).

Before we consider quantum systems, it will be useful to review some caveats to the question of whether a CCM provides an adequate description of some given data, and in particular, whether a given CCM gives a *faithful* account of the observed statistics. The first caveat has to do with the possibility of hidden, or *latent* variables. Suppose that we have a probability distribution  $P(X)$  for which there is no CCM that generates it (this can occur, for example, if some exogenous variables in the model are found to be correlated with each other, thereby violating the basic property of Markovianity required for a CCM). Rather than giving up the possibility of a causal explanation, we might consider that there exist additional variables that have not been observed, but whose inclusion in the data would render the statistics explainable by some CCM. Formally, suppose there is an extension of  $P(X)$  to some larger distribution  $P'(X, \lambda)$  that includes latent variables  $\lambda$ , such that the observed statistics are the marginal probabilities obtained by summing over the unobserved variables:

$$P(X) = \sum_{\lambda} P'(X, \lambda). \quad (8)$$

If there exists a CCM  $\mathcal{M}(X, \lambda)$  such that  $P^{\mathcal{M}}(X, \lambda) = P'(X, \lambda)$ , then we can say that this CCM explains the probability distribution  $P(X)$  with the aid of the latent variables  $\lambda$ . The admittance of hidden variables in causal models seems to lead to a problem: it turns out that *every* probability distribution  $P(X)$  can be explained by a CCM, with the aid of a sufficient number of hidden variables! For this reason, we further constrain the possible explanations of the observed data by requiring that the models be *faithful* to the data:

**Definition 9. Faithfulness.** Consider a distribution  $P(X)$  and a CCM  $\mathcal{M}(X) = \{G, F\}$  that generates  $P(X)$ . The explanation offered by  $\mathcal{M}(X)$  is called *faithful* to  $P(X)$  iff the DAG derived from  $\mathcal{M}(X)$  is a perfect map of  $P(X)$ , i.e.  $\bar{C}(G) = \bar{C}(P)$ .

*Latent variables:* Suppose there is no CCM  $\mathcal{M}(X)$  that is faithful to  $P(X)$ . Consider instead a CCM  $\mathcal{M}(X, \lambda) = \{G', F'\}$ , which obtains  $P(X)$  by summing the generated distribution  $P'(X, \lambda)$  over the hidden variables  $\lambda$ . This extended CCM is considered faithful to  $P(X)$  iff every CI relation in  $P(X)$  is implied by the extended DAG  $G'$ , i.e.  $\bar{C}(P) \subseteq \bar{C}(G')$ .

The motivation for this definition is that a faithful explanation of the observed statistics is a better candidate for describing the ‘real causal structure’ of the system than an unfaithful explanation, because it accurately captures all dependencies in the observed statistics. If there exists no faithful explanation of the observed statistics, but one can obtain a faithful explanation using hidden variables, then we can interpret the statistics of the observed variables as the marginal statistics arising from ignoring the unobserved variables. Note that not all probability distributions can be *faithfully* reproduced by some CCM, even with the aid of hidden variables; this will be relevant when we consider quantum mechanics in section 2.6.

Geiger and Pearl [25] showed that, for every DAG  $G$ , one can explicitly construct a distribution  $P$  such that  $G$  is faithful for  $P$ , i.e. such that  $\bar{C}(G) = \bar{C}(P)$  holds. Furthermore, if there exists a DAG  $G$  that is faithful for a given  $P$ , and if we assume there are no latent variables, then it can be shown that there exists a set of model parameters  $F$  such that the CCM  $\{G, F\}$  generates  $P$  [26].

Finally, faithfulness can be equivalently defined as the rejection of *fine-tuning* of the model parameters, as noted in [2, 13]. In particular, if a DAG is unfaithful, this implies that there exists at least one CI in the statistics

that is not implied by the DAG. It can be proven that within the set of probability distributions compatible with the DAG, those that satisfy such additional CI relations are a set of measure zero [13]. Thus, the only way that such additional CI relations could arise is by a kind of ‘conspiracy’ of the model parameters  $F$  to ensure that the extra CI holds, even though it is not indicated by the causal structure. In this sense, fine-tuning represents causal influences that exist at the ontological level but cannot be used for signalling at the level of observed statistics, due to the careful selection of the functional parameters; any small perturbation to these parameters would result in a signal.

**Example 3.** Recall the system of example 1 and suppose that we observe  $(X \perp Y|Z)$  in  $P$ . This CI relation is not found in  $\bar{L}_\emptyset$ ; in fact, there is a directed edge from  $Y$  to  $X$  in the DAG of  $L_\emptyset$  (figure 1). The only way to account for this discrepancy is if the model parameters  $F$  are chosen such that the predicted signal from  $Y$  to  $X$  is obscured. This could happen if  $X, Y, Z$  are positive integers and we choose model parameters  $f_y[z, u_y] = u_y + z$  and  $f_x[y, z] = y + z - k$  for some integer constant  $k$ , and  $F_y$  is restricted such that  $P(u_y = k) = 1$ . Then we obtain the joint probability:

$$\begin{aligned} P(X, Y, Z) &= \sum_{U_Y} \delta(x, f_x[y, z]) \delta(y, f_y[z, u_y]) P(U_Y) P(Z), \\ &= \delta(x, y + z - k) \delta(y, k + z) P(Z), \\ &= \delta(x, 2z) \delta(y, k + z) P(Z), \end{aligned} \quad (9)$$

which factorizes into  $P(X|Z)P(Y|Z)P(Z)$ . We see that the value of  $X$  tells us nothing about the value of  $Y$ , because the  $k$ 's conveniently cancel out in  $P(X|YZ)$ , leading to the observed independence  $(X \perp Y|Z)$ . This can be understood as fine-tuning of  $F_y$  such that its error variable is constant. Indeed, if  $P(U_Y)$  was distributed over several values, then  $X$  would still carry information about  $Y$  in accordance with the causal structure. Hence, absence of fine-tuning in  $P$  with respect to a causal input list  $L_\emptyset$  can also be defined as the requirement that the CI relations observed in  $P$  should be robust under changes in the model parameters consistent with  $L_\emptyset$  [13].

## 2.6. Does quantum mechanics require fine-tuning?

Consider a probability distribution  $P(A, B, S, T)$ , satisfying the generating set of CI relations  $K := \{(S \perp T|\emptyset), (A \perp T|S), (B \perp S|T)\}$ . Suppose that the closure of this set contains all the CI relations satisfied by  $P$ , i.e.  $\bar{K} = \bar{C}(P)$ . This represents a generic Bell-type experiment: the setting variables  $S, T$  are independent of each other, and there is no signalling from  $S$  to  $B$  or from  $T$  to  $A$ , but the outputs  $A, B$  are correlated. (Here, the absence of signalling is a constraint on the allowed probability distributions. It refers to the fact that the marginal distribution of the outcomes on each side must be conditionally independent of the values of the setting variables on the opposite side, a requirement often called ‘signal locality’ in the literature. It does not forbid the possibility of signalling at the ontological level, as in nonlocal hidden variables, etc.)

Note that we cannot explain these correlations by a CCM without latent variables, because  $A$  and  $B$  are correlated without a common cause. Hence let us consider the extended distribution  $P'(A, B, S, T, \lambda)$  satisfying  $\sum_\lambda P'(A, B, S, T, \lambda) = P(A, B, S, T)$ . Of course, we require that  $\bar{K} \subseteq \bar{C}(P')$ , but we can impose additional physical constraints on the hidden variable  $\lambda$ . In particular, we expect  $\lambda$  to be independent of the settings  $S, T$  and, in keeping with the no-signalling constraint, to represent a common cause subject to Reichenbach’s principle. This leads to the extended set of constraints  $K' := \{(S \perp T\lambda|\emptyset), (T \perp S\lambda|\emptyset), (A \perp T|S\lambda), (B \perp S|T\lambda)\}$  and we assume that  $\bar{K}' \subseteq \bar{C}(P')$ , i.e. that the extended distribution satisfies at least these constraints. We then ask whether there exists a CCM that can *faithfully* explain the observed correlations.

If  $A$  and  $B$  are independent conditional on the hidden variable  $\lambda$  in the distribution  $P'$ , i.e. if  $(A \perp B|\lambda)$  holds in  $\bar{C}(P')$ , then it is easy to see that  $\lambda$  qualifies as a common cause of  $A$  and  $B$  and the correlations can be explained by a CCM with the DAG  $G'$  shown in figure 4 (a). It can be shown that this occurs whenever  $P$  satisfies Bell’s inequality [2]. Conversely, Wood and Spekkens showed that if  $P$  violates Bell’s inequality, there is no CCM that can faithfully explain  $P$ , even allowing hidden variables. Of course, one can find numerous un-faithful explanations, such as the DAG in figure 4 (b) and fine-tuning of the model parameters to conceal the causal influence of  $S$  on  $B$ . This result implies that, in general, CCMs cannot faithfully explain the correlations seen in entangled quantum systems.

**Example 4.** Consider the deBroglie–Bohm interpretation of quantum mechanics. This interpretation gives a causal account of Bell inequality violation using super-luminal influences, one possible variant of which is depicted in figure 4(b). Here,  $\lambda$  is a hidden variable that carries information about the setting  $S$  faster than light to the outcome  $B$ . The model posits a CCM that generates a distribution  $P'(A, B, S, T, \lambda)$  and the observed statistics  $P(A, B, S, T)$  are interpreted as the marginal obtained from this distribution by summing over  $\lambda$ . The

no-signaling CI relations  $(A \perp T|S)$ ,  $(B \perp S|T)$  hold in the observed statistics, however they do not follow from the DAG figure 4(b) which includes the hidden variables, hence the CCM that generates  $P'$  using this graph is not a faithful explanation for  $P$ . In general, the deBroglie–Bohm interpretation and its variants appear to require fine-tuning [2].

How should we interpret this result? On one hand, we might take it as an indication that faithfulness is too strong a constraint on the laws of physics, and that nature allows hidden variables whose causal influences are concealed at the statistical level by fine-tuning. Alternatively, we could take it to indicate that the class of physical models describable by CCMs is not universal, and that a new type of causal model is needed to give a faithful account of quantum systems. Along these lines, we could choose to interpret figure 4(a) as a quantum circuit, where  $\lambda$  now stands for the preparation of an entangled pair of quantum systems and the arrows stand for their distribution across space. In doing so, we implicitly shift our perception of quantum mechanics from something that needs to be explained, to something that forms part of the explanatory structure. We no longer seek to explain quantum correlations by an underlying causal mechanism, but instead we incorporate them as a fundamental new addition to our causal structure, which can then be used to model general physical systems. This approach entails that we no longer require  $(A \perp B|\lambda)$  to hold in the extended set of constraints  $K'$  for the ‘common cause’  $\lambda$ , and hence that we abandon RCCP (specifically, the factorization property). Since we aim to use a DAG like figure 4(a) as the explanation, this implies that  $d$ -separation is no longer the correct criterion for reading CI relations from the graph, now that the DAG is interpreted as a quantum circuit. In what follows, we will propose a new criterion that serves this purpose, leading to the concept of a QCM.

### 3. Quantum causal models

#### 3.1. Preliminaries

We begin by considering quantum networks modelled as a DAGs, in which the nodes represent state preparations, transformations and measurements. Based on this interpretation, we obtain a corresponding notion of a quantum input list and a graph separation criterion that connects the DAG to the list that generates it. We mention that there exist other approaches to quantum computation in which it would be interesting to explore causal relations, such as measurement-based quantum computation. For efforts along these lines, see e.g. [27, 28].

The general theory of quantum networks as given in [7] provides a DAG representation in which nodes represent completely general quantum operations. Below, we define a canonical form of a general quantum network in order to cleanly separate the classical apparatus settings from the measurement outcomes, to facilitate the definition of a graph separation criterion. Given a DAG, we divide the nodes into three classes: as before, those with only outgoing edges are called *exogenous*; those with only ingoing edges are called *drains*; those with ingoing and outgoing edges are called *intermediates*. We will typically restrict our attention to connected graphs, i.e. graphs in which there are no *isolated nodes* having no parents or children. Such variables can usually be ignored since they are irrelevant to all other variables of interest. However, in cases where they do arise (e.g. when removing nodes in the proof of theorem 3 later on in section 3.3), they can be interpreted as indicating an operation performed on a system not connected to anything else—like the value of a dial on a broken telephone.

We assign the following interpretations to the elements of the DAG:


*Edges*: every edge in a DAG is associated with a Hilbert space of finite dimension  $d > 1$ . Thus, we can associate a physical system with finite degrees of freedom to every edge. The Hilbert space dimension is allowed to be different for different edges in the DAG.

*Exogenous nodes*: every exogenous node is associated with a random variable. Each possible value of the variable corresponds to the preparation of a normalized quantum state, represented by a density matrix. The set of states need not be orthogonal—in fact they may even be degenerate, with more than one value of the variable corresponding to preparation of the same state. The only requirement is that the states exist in a Hilbert space with dimension equal to  $\mathcal{H}^{(\text{out})}$ , which is the tensor product of the Hilbert spaces of all the outgoing edges.

*Intermediate nodes*: every node with both incoming and outgoing edges is associated with a random variable. Each value of the variable represents a general transformation (a CPT map)  $C: \mathcal{H}^{(\text{in})} \rightarrow \mathcal{H}^{(\text{out})}$  from the input to output Hilbert spaces.

*Drains*: every drain is associated with a random variable, which may be classed as either a *setting* or an *outcome* (see section 3.2). For outcomes, each value of the variable corresponds to one of the possible outcomes of a general measurement (POVM) on the incoming systems  $\mathcal{H}^{(\text{in})}$ .

In the alternative case where the drain node is a setting, it is treated just like an intermediate node (each value corresponds to a CPT map on  $\mathcal{H}^{(\text{in})}$ ) except that the outgoing systems are discarded. Since these variables will




**Figure 5.** One possible way to convert a quantum circuit (a) into a DAG, by replacing feed-forwards with unitary interactions, followed by measurements in a fixed basis and replacing general measurements with unitary coupling to ancilla states, followed by measurement in a fixed basis, as shown in (b). The resulting DAG is shown in (c).


turn out to have no dependence on any other variables in the system, they could be ignored just like isolated nodes, although we will find it useful to include them in the proof of theorem 3 in section 3.3.

The above definitions allow us to associate a quantum network to any DAG. Conversely, every quantum network has a representation as a DAG of this form.

**Example 5.** Consider the circuit in figure 5(a). This describes the preparation of two qubits as mixtures  $\rho_1 = \gamma_1 |\psi_0\rangle\langle\psi_0| + (1 - \gamma_1)|\psi_1\rangle\langle\psi_1|$  and  $\rho_2 = \gamma_2 |\psi_0\rangle\langle\psi_0| + (1 - \gamma_2)|\psi_1\rangle\langle\psi_1|$  in an arbitrary orthogonal basis  $\{|\psi_0\rangle, |\psi_1\rangle\}$ . These are followed by a measurement of the first qubit in the computational  $\sigma_Z$  basis  $\{|0\rangle, |1\rangle\}$  and the subsequent application of a  $\sigma_X$  gate to the second qubit, conditional on the outcome of the first measurement. Finally a POVM is applied to the second qubit by coupling it via unitary interaction (either  $U_1$ ,  $U_2$  or  $U_3$ ) to a third ancilla qubit  $|\phi\rangle$ . The ancilla is traced out and the remaining qubit measured in the  $\sigma_Z$  basis. In figure 5(b) the feed-forward has been replaced with a unitary interaction (a CNOT) followed by tracing out the first qubit (all feed-forwards can be described in this way to ensure that the setting variables, representing the choice of input state and unitary, remain independent of each other). The tracing-out of the ancilla qubit is replaced with a measurement in the  $\sigma_Z$  basis, whose outcome can be ignored. In this form, the circuit can be cast directly into a DAG, as shown in figure 5(c). The variables  $X$  and  $Y$  take values corresponding to the basis states  $\{|\psi_0\rangle, |\psi_1\rangle\}$ , distributed with probabilities so as to produce the mixed states  $\rho_1, \rho_2$ .  $Z$  and  $S$  are single-valued,



**Figure 6.** ‘Chained’ variables: S and Y are each connected to T and Z by a W-chain. However X and R are W-detached from {Z, T}, because the collider Y is not in the set W.



q-separation	
Implied	Not implied
$(T \perp Z   VW)$	$(T \perp Y   XVU)$
$(X \perp Z   UW)$	$(T \perp Z   YVU)$
$(T \perp W   V)$	$(V \perp W   YZ)$
$(U \perp W   XYZ)$	$(U \perp W   V)$

**Figure 7.** An illustration of  $q$ -separation. The table contains CI relations that are either implied or not implied by the DAG using the rules of  $q$ -separation. Note the differences to  $d$ -separation (figure 2).

corresponding to the state  $|\phi\rangle$  and the unitary CNOT respectively.  $T$  has three values corresponding to the three possible unitaries, and is distributed according to the probability of each unitary being implemented. Finally,  $U$ ,  $V$ ,  $W$  are all binary-valued, corresponding to outcomes  $\in\{0, 1\}$  and whose probabilities are given by quantum mechanics (see section 3.2).

The example illustrates that it is also possible (although not necessary) to restrict all operations to pure states, unitaries and projective measurements and treat more general operations as epistemic mixtures of these. In the rest of the paper we allow the model parameters to be arbitrary, without confining ourselves to any particular convention.

Note that a single DAG can represent *any* member of the class of quantum networks with the same basic topology. Thus, in a quantum causal model, the preparations, transformations and measurements are taken as the model parameters and the DAG provides the causal structure, as explained in the next section.

### 3.2. Quantum input lists and model parameters

Recall that the classical causal input list  $L_{\mathcal{O}}$  represents a set of CI relations between variables in a CCM, from which a DAG can be easily constructed. The motivation for the causal input list comes from its physical interpretation, discussed in section 2.3, which embodies principles like the RCCP that we expect to hold for classical physical systems. Hence, to define the quantum analog of a causal input list, we should begin by asking: for variables in a quantum network, what physical principles constrain their statistical dependencies?



First, we note that the observables in a quantum network fall naturally into two distinct categories: *settings*  $S_i$  and *outcomes*  $O_i$  (we continue to use  $X$  to denote a generic variable or set of variables). The settings  $S_i$  determine the states produced at the sources and the transformations applied in the network, while the  $O_i$  represent the outcomes of general measurements. Since the settings  $S_i$  play the same role as the exogenous variables in a CCM, we assume that they are all distributed independently of each other; however, unlike in a CCM, this property now also applies to variables represented by intermediate nodes. This assumption is the analog of the Markovianity assumption for a CCM, and it embodies one aspect of the common cause principle that is retained in quantum mechanics, namely, that correlated variables (conditional on the empty set) must share a common source, or must have interacted previously. It is in this sense that the RCCP can be said to hold for quantum correlations in [9] (recall the discussion of section 2.3).

Also as before, we assume an absolute ordering of the variables and enforce the physical assumption of causality (no causal loops) and we again assign a set of parents to each variable, representing the connections in the network and (implicitly) the possibility of a causal influence. However, unlike the case of a CCM, we are not able to interpret the parents of a variable as its direct causes. This is because the values of the settings by definition do not have any causes in the network (they are chosen by external factors, like experimental intervention). Furthermore, the parents of an outcome no longer screen it off from its other ancestors: the influence of an initial state preparation on the measurement outcome cannot in general be screened off by a choice of intermediate transformation. We leave it as an open question whether one can formulate a quantum network in a manner that respects this property of CCMs; we will find it more convenient simply to abandon it. Indeed, since the variables representing the preparation and transformation are assumed to be independent, they cannot carry any information about each other, nor can the variable representing the transformation reveal any information about the quantum system on which it acts. The assignment of parents to the variables therefore places much weaker constraints on the correlations than in the classical case. However, the following physical assumption is still justified in a quantum network:

**Assumption 2.** *The possible causes of an outcome are its ancestors.* (Compare to assumption 1.)

This assumption reflects our intuition that it is only the operations performed on a quantum system leading up to its measurement that can have a causal effect the measurement outcome. Indeed, it is also argued in [9] that there is no reason to maintain the distinction between direct and indirect causes in any generalized model that goes beyond classical correlations.

It is clear from our discussions in section 2 that the causal Markov condition is not expected to hold in a quantum network, since the RCCP no longer holds. Instead, we expect it to be replaced by a weaker property:

*Quantum causality condition:* an outcome is independent (conditional on the empty set) of all settings that are not its causes and all outcomes that do not share a common cause.

This property expresses the fact that outcomes should be independent of any settings from which they are causally disconnected and should be correlated only with other outcomes that share a common cause. This property holds also in the classical case, but unlike the classical case, we now do not require sets of outcomes to be independent of each other *conditional* on their common causes—instead we allow them to still be dependent, admitting violations of the factorization property of the RCCP (recall section 2.3). In addition to this weakened version of the RCCP, we still have the classical feature that independent variables can become dependant conditioning on common effects. Thus, for two variables to be independent, we will still have to avoid conditioning on certain colliders.

To make these ideas formal, let us consider a set of random variables partitioned into *outcomes*  $O$  and *settings*  $S$ . The following definitions will also be useful (see figure 6):

**Definition 10.  $O'$ -chain.** Given a set of outcomes  $O'$ , two other sets of outcomes  $O_1$  and  $O_2$  are said to be *connected by an  $O'$ -chain* iff  $O_1$  shares an ancestor with a member of  $O'$  that shares an ancestor with another member of  $O'$ , (etc), that shares an ancestor with  $O_2$ . A set of settings  $S_1$  is linked to  $O_2$  by an  $O'$ -chain iff  $S_1$  has a descendant in  $O_2$ , or a descendant  $O'$  that is connected by an  $O'$ -chain to  $O_2$ . Similarly,  $S_1$  and  $S_2$  are connected by an  $O'$ -chain iff they both have descendants in  $O'$  that are connected in this way.

**Definition 11.  $O'$ -detached.** Given a set of outcomes  $O'$  and some variables  $V$ , the set of all variables not connected to  $V$  by an  $O'$ -chain are said to be  *$O'$ -detached* from  $V$ , denoted  $\mathbf{dt}_{O'}(V)$ .

Note that when  $O'$  is the empty set, the detached variables  $\mathbf{dt}_{\emptyset}(V)$  are just those outcomes that do not share an ancestor with outcomes in  $V$ , and those settings that are not ancestors of outcomes in  $V$ —which are precisely the variables identified by the quantum causality condition when applied to the outcomes in  $V$ .

Intuitively, the  $O'$ -chain tells us when conditioning on  $O'$  will make two independent variables dependant. If two variables in a DAG are connected by an  $O'$ -chain, assuming there is no directed path from one to the other, it means either they are connected by a path on which every collider has a directed path to  $O'$ , or they are both settings whose descendants in  $O'$  share a common cause. Hence, the mutually detached variables are those nodes in the graph for which every path between them contains at least one collider that does *not* lead to  $O'$ , or else both nodes are settings and at least one of them has no descendants in  $O'$ . This will be useful later when we consider graph separation.

*Notation:* let  $\neg X$  denote the complement of a set  $X$ , let  $X/S$  denote only the outcomes in  $X$ , and let  $\neg_S X$  be all settings not in  $X$ , i.e.  $\neg_S X := S \cap \neg X$ . Under a choice of ordering  $\mathcal{O}$ , let  $S(< X_i)$  denote the set of predecessors of  $X_i$  in  $S$ . Using these definitions, we propose the following characterization of the CI relations in a quantum network:

**Definition 12. Quantum input list.** A quantum input list  $Q_{\mathcal{O}}$  is a pair  $\{\text{PA}_{\mathcal{O}}, Q\}$ , containing:

- (i) An ordered list of parents,  $\text{PA}_{\mathcal{O}} := \{\text{pa}(X_i) : i = 1, 2, \dots, N\}$ , where each set of parents  $\text{pa}(X_i)$  is a subset of  $S(< X_i)$ . Members of  $O$  cannot have children. Ancestors, descendants, etc, are defined from the list of parents in the usual way.
- (ii) A set of CI relations denoted  $Q$ , constructed as follows. For every subset of settings  $S'$  and outcomes  $O'$ , there is a CI relation of the form  $(S' \perp \text{dt}_{O'}(S') | O')$  and a CI relation of the form  $(O' \text{an}(O') \perp \text{dt}_{\emptyset}(O') | \emptyset)$  in  $Q$ .

The first CI relation in the above definition expresses the physical requirement of setting independence, modulo the possibility of correlating the settings by conditioning on their effects. In particular, it says that settings are guaranteed independent except when conditioned on an  $O'$ -chain that connects them. The second CI relation simply expresses the quantum causality condition, since  $\text{dt}_{\emptyset}(O')$  contains both the ‘non-causes’ of  $O'$ ,  $\neg_S \text{an}(O')$ , and the outcomes that do not share a common cause with  $O'$ .

The quantum input list  $Q_{\mathcal{O}}$  is said to be *compatible* with a given probability distribution  $P$  iff  $\bar{Q} \subseteq \bar{C}(P)$ . Given a quantum input list, we can construct a DAG in the usual way, by drawing a directed edge to each variable from each of its parent nodes. The DAG  $G_Q$  constructed in this way is said to be *generated* by the list  $Q_{\mathcal{O}}$ . As usual, the ancestors of  $X$  are those nodes in the graph that have a directed path leading to  $X$ . The quantum input list defines the causal constraints on the variables, based on their interpretation as settings and outcomes in a quantum network. All conditional independencies in this list are expected to hold in a quantum circuit, when the circuit is expressed as a DAG as outlined in section 3.1. In addition, we conjecture that the list captures *all* such relations:

*Conjecture:* if a CI relation holds in every quantum network represented by a DAG  $G$ , then it is implied by  $Q$  in any quantum input list  $Q_{\mathcal{O}}$  that generates  $G$ .

One way to prove this conjecture would be to show that it is always possible to find parameters (i.e. states, operations and measurements) on a quantum network that violate any relation not implied by its quantum input list. We content ourselves with the following plausibility argument: suppose that there exists some CI relation  $R$  that holds for a quantum circuit with DAG  $G$ , but which does not appear in a quantum input list that generates  $G$ . This suggests that  $R$  is not implied by the quantum causality condition, hence that it expresses the independence of two outcomes that share a common ancestor, or the independence of an outcome and a setting that is its ancestor. There seems nothing preventing us from making a new circuit with the same DAG in which all outcomes with a common ancestor are correlated with each other, and in which each outcome is dependent on all of its ancestors (assuming no restrictions on the size of the Hilbert spaces involved). Since this would invalidate  $R$ , we conclude that there is no relation that holds in all circuits with the same DAG, but does not appear in the quantum input list, vindicating the conjecture. However, a rigorous proof is still needed.

So far, we have only specified the causal structure and independence relations. To obtain a full joint probability distribution from a quantum input list, we need to supplement it with model parameters specifying the pure state preparations, transformations, and measurements that correspond to the variables. These parameters define the space of possible quantum circuits that are described by a given DAG:

**Definition 13. Quantum model parameters.** Consider a set of variables  $X_i$  with outcome spaces  $\mathcal{E}_{X_i}$ , with the drain nodes partitioned into settings or outcomes, connected in a DAG  $G_Q$ . Then the *quantum model parameters*  $F_q$  consist of:

- (i) For each edge, a Hilbert space of finite dimension  $n > 1$ ;
- (ii) For every outcome drain node  $X_i$ , a POVM with an outcome for every value in  $\mathcal{E}_{X_i}$ ;

- (iii) For every exogenous node  $X_i$ , a normalized quantum state for every value in  $\mathcal{E}_{X_i}$ ;
- (iv) For every intermediate node (and every setting drain node)  $X_i$ , a CPT map  $\mathcal{C}: \mathcal{H}^{(\text{in})} \rightarrow \mathcal{H}^{(\text{out})}$  for every value in  $\mathcal{E}_{X_i}$ ;
- (vi) A marginal probability distribution on the outcome space  $\mathcal{E}_{X_i}$  of every variable that is not an outcome drain node. These marginal distributions are all mutually independent.

The states and operators mentioned above apply to the Hilbert spaces of their respective nodes, as determined using (i) and the number of outgoing and ingoing edges. The distributions given in (vi) represent the experimenters' setting choices and/or environmental conditions. These are used to determine the resulting probability distributions of the outcome variables according to the usual laws of quantum mechanics. This is made precise using the following definition:

**Definition 14. Quantum causal model.** A QCM on a set of variables  $X$  is a pair  $\{Q_\mathcal{O}, F_q\}$  consisting of a quantum input list  $Q_\mathcal{O}$  for the set  $X$ , and a set of quantum model parameters  $F_q$  for the DAG  $G_Q$  generated by the input list.

Every QCM defines a joint probability distribution over its variables according to the following procedure. Consider a QCM on  $N$  ordered variables  $\{X_i: i = 1, 2, \dots, N\}$ , and partition of the set  $\{1, 2, \dots, N\} := S \cup O$  such that  $i \in S$  labels the exogenous variables, intermediate variables, and setting drain variables, and  $i \in O$  labels the variables corresponding to outcome drains in the DAG  $G_Q$ . From  $F_q$  we obtain the mutually independent marginal distributions  $P(X_i) \forall i \in S$ , which includes all setting variables. The joint probability of the outcomes conditional on the settings,  $P(\bigcup_{i \in O} X_i | \bigcup_{i \in S} X_i)$ , is computed in the usual way from the quantum circuit obtained from the DAG  $G_Q$  using the states and transformations associated with the settings  $\{\bigcup_{i \in S} X_i\}$ . One thus obtains the total joint probability:

$$P(X) = \prod_{i \in S} P(X_i) P\left(\bigcup_{j \in O} X_j \middle| \bigcup_{j \in S} X_j\right). \quad (10)$$

*Note:* Our definition of a QCM on a DAG  $G$  can be regarded as a concrete example of the more general notion of a *quantum correlation* on the graph  $G$ , as defined by Fritz [9]. In particular, working in the category  $\mathcal{C}$  of completely positive maps, where Hilbert spaces are the objects and CP maps are the morphisms, we assign Hilbert spaces to the edges of the graph and CP maps to the nodes. The model parameters are just the set of functions from outcomes to morphisms that define the *C-instruments* in the language of [9], allowing us to compute probabilities. Our model distinguishes these functions according to the placement of their nodes in the graph (exogenous, intermediate or drain). This convention does not represent a limitation of our model, but are used for convenience. By contrast, our restriction to the case of variables with finite outcome spaces is a limitation of our model, but we expect the generalization to continuous variables following [9] to be straightforward.

Now that we have defined a QCM, we would like to have a graph separation rule analogous to  $d$ -separation that would allow us to recover all the CI relations implied by  $Q_\mathcal{O}$  from the DAG  $G_Q$ . This is proposed in the next section.

### 3.3. Graph separation in quantum networks

In general, because of the failure of the RCCP, we can never guarantee that two outcomes will be independent conditional on their common causes. However, there are still situations in which variables are expected to be conditionally independent of each other; we examine the possibilities below.

Two settings are already assumed to be chosen independently, so they can only become dependent on each other by conditioning on a common effect (which is an outcome), or conditioning on a connected chain of such effects. This applies also to conditioning on common effects in a CCM (recall section 2.3). In the case of a setting and an outcome, these might be dependent on each other if the setting is already a possible cause of the outcome, since the causal influence cannot in general be screened-off by other variables. On the other hand, if there is no directed path from the setting  $S_1$  to the outcome  $O_2$  and no chain of conditioned effects, one would expect the two to be independent. We must be careful, however: it is also possible for  $O_2$  to be correlated with another outcome that is descended from  $S_1$ , such that conditioning on the latter outcome correlates  $S_1$  with the causally separated outcome  $O_2$ . To ensure their independence, therefore, one should not also not condition on any outcomes that are descended from the setting. Finally, two outcomes should be independent unless they share a common cause, or are connected by a chain of conditioned effects.

These considerations lead us to the following graph separation criterion:

**Definition 15.  $q$ -separation.** Given a DAG representing a quantum network, two disjoint sets of variables  $X$  and  $Y$  are said to be  $q$ -separated by a third disjoint set  $Z$ , denoted  $(X \perp Y | Z)_q$ , iff every undirected path between  $X$  and  $Y$  is rendered inactive by a member of  $Z$ . A path connecting two variables is rendered inactive by  $Z$  iff at least one of the following conditions is met (see figure 7):

- (i) both variables are settings, and at least one of the settings has no directed path to any outcome in  $Z$ ;
- (ii) one variable is a setting and the other is an outcome, and there is no directed path from the setting to the outcome, or to any outcome in  $Z$ ;
- (iii) the path contains a collider  $i \rightarrow m \leftarrow j$  where  $m$  is not an outcome in  $Z$ , and there is no directed path from  $m$  to any outcome in  $Z$ .

Of course, the heuristic motivation given above does not necessarily guarantee that  $q$ -separation captures all of the CI relations that are implied by a quantum input list, nor is it obvious that the input list contains all CI relations implied by  $q$ -separation. A proof that  $q$ -separation is sound and complete for quantum input lists is given in the next section.

### 3.4. The $q$ -separation theorem

In this section we prove the soundness and completeness of  $q$ -separation. The proof approximately follows that of Pearl and Verma [22] for  $d$ -separation. Consider the set of CI relations obtainable from a DAG  $G$  using the  $q$ -separation criterion (definition 15) and assuming  $(X \perp Y | Z)_q \Rightarrow (X \perp Y | Z)$ . Let this set be denoted  $C_q(G)$ , with  $\bar{C}_q(G)$  its closure. In fact, we prove in appendix that  $q$ -separation satisfies the semi-graphoid axioms, and therefore this set is equal to its closure:  $C_q(G) = \bar{C}_q(G)$ . Hence,  $C_q(G)$  and  $\bar{C}_q(G)$  are used interchangeably in what follows. If we replace  $d$ -separation with  $q$ -separation in definition 8, we obtain analogous criteria for  $G$  to be an I-map or a perfect map of a given distribution  $P$ . We can now prove the following useful theorem:

**Theorem 2.** : Let the DAG  $G$  be a perfect map of a distribution  $P(X)$  under  $q$ -separation, i.e.  $\bar{C}_q(G) = \bar{C}(P)$ . Then there is a quantum input list  $Q_\Theta$  compatible with  $P$  that generates the DAG  $G$ .

**Proof.** The DAG  $G$  imposes a partial order on the variables  $X$ . Choose any total order  $\Theta$  that is consistent with this. Label the nodes in  $G$  as outcomes  $O_i \in O$  if they are drains, and settings  $S_i \in S$  otherwise. Define the parents  $\text{pa}(X_i)$  in  $\text{PA}_\Theta$  to be the nodes with directed edges pointing to  $X_i$  in the graph. A path between a setting and any other variable is rendered inactive by outcomes  $O'$  if there is at least one collider on the path not in  $O'$  and with no directed path to  $O'$ , or if the other variable is also a setting and  $O'$  does not contain descendants of both settings (condition (i) of  $q$ -separation). This is true for all variables that are  $O'$ -detached from the setting, hence the CI relation  $(S' \perp \text{dt}_{O'}(S') | O')$  is implied by  $G$ . For each set of outcome nodes and their ancestors,  $O' \text{an}(O')$ , a path from this set to the non-ancestors  $\neg_S \text{an}(O')$  can only be activated by conditioning on an outcome. Furthermore, a path from  $O' \text{an}(O')$  to  $\text{dt}_{O'}(O')/S$  must contain a collider, or else  $O'$  and  $\text{dt}_{O'}(O')/S$  would share an ancestor (a contradiction), so it too can only be activated by conditioning on an outcome. Hence these paths are rendered inactive by the empty set, and since  $\neg_S \text{an}(O') \cup \text{dt}_{O'}(O')/S = \text{dt}_{O'}(O')$ , the CI relation  $(O' \text{an}(O') \perp \text{dt}_{O'}(O') | \emptyset)$  is implied by  $G$ . According to definition 12, these ingredients are sufficient to specify a quantum input list. By construction, this list also generates the DAG  $G$ .  $\square$

**Corollary 1.** Let  $Q_\Theta$  be any quantum input list that generates a DAG  $G$ . Then every relation in the closure  $\bar{Q}$  is implied by  $G$  via  $q$ -separation, i.e.  $\bar{Q} \subseteq \bar{C}_q(G)$ . To see why, let  $Q'_\Theta$  be the input list generated from  $G$  by the method described in theorem 2. According to that theorem,  $\bar{Q}'_\Theta \subseteq \bar{C}_q(G)$  for this input list. The corollary follows from the fact that every quantum input list that generates  $G$  must be equivalent to  $Q'_\Theta$ , up to a choice of total ordering consistent with  $G$ .

The next theorem provides the key result.

**Theorem 3.** Given a quantum input list  $Q_\Theta$  that is compatible with some distribution  $P(X)$ , the DAG  $G$  generated by  $Q_\Theta$  is an I-map of  $P$ , that is,  $\bar{C}_q(G) \subseteq \bar{C}(P)$ .

**Proof.** We prove the result by induction on the number of variables. First we show that the result holds for  $k$  variables, given that it holds for  $k - 1$  variables. Then we note that the result holds trivially for one variable; hence, by induction, it holds for any number of variables.

Let  $P$  be a distribution on  $k$  variables and  $Q_{\mathcal{O}}$  a compatible quantum input list, which generates the DAG  $G$ . Let  $n$  be the last variable in the ordering  $\mathcal{O}$ ; let  $\bar{C}(P - n)$  be the closed set of CI relations formed after removing from  $\bar{C}(P)$  all CI relations involving  $n$ ; let  $P - n$  be any probability distribution having exactly the closed set of CI relations  $\bar{C}(P - n)$  (such a distribution can always be constructed [13]); and let  $G - n$  be the DAG formed by removing the node  $n$  and all its connected edges from the graph  $G$ .

Consider the list obtained from  $Q_{\mathcal{O}}$  by removing every CI relation involving  $n$  from  $Q$  and removing  $\mathbf{pa}(n)$  from the list of parents, and let the reduced list of parents be denoted  $\mathbf{PA}_{\mathcal{O}} - n$  (Note that this procedure might result in one or more settings that are isolated or drain nodes in the graph). Let the resulting list be denoted  $Q - n$ . Then the pair  $Q_{\mathcal{O}} - n := \{\mathbf{PA}_{\mathcal{O}} - n, Q - n\}$  is a valid quantum input list on  $k - 1$  variables. Furthermore, by construction,  $Q_{\mathcal{O}} - n$  generates the DAG  $G - n$ .

Let us now assume that  $G - n$  is an I-map of  $P - n$ :  $\bar{C}_q(G - n) \subseteq \bar{C}(P - n)$ . We aim to prove that, under this assumption,  $G$  is also an I-map of  $P$ ,  $\bar{C}_q(G) \subseteq \bar{C}(P)$ . To do so, we will consider each CI relation of  $\bar{C}_q(G)$  and show that it exists also in  $\bar{C}(P)$ .

The CI relations of  $\bar{C}_q(G)$  can be divided into three cases:

- (1)  $n$  does not appear in the CI relation;
- (2)  $n$  appears in the first position in the CI relation, e.g.  $(nX \perp Y | Z)$ ;
- (3)  $n$  appears in the last position in the CI relation, e.g.  $(X \perp Y | nZ)$ .

Note that, if  $n$  appears in the second position in the CI relation, we can use symmetry (semi-graphoid axiom 1.a.) to move it into the first position and thereby convert it into case (2) above. We now prove the result for each case separately.  $\square$

**Lemma 1.** Let  $X, Y, Z$  be disjoint sets of variables and let  $R \in \bar{C}_q(G)$  be a relation of the form  $(X \perp Y | Z)$  that does not contain the variable  $n$ . Then  $R \in \bar{C}(P)$ .


**Proof.** Since  $R$  is in  $\bar{C}_q(G)$ , it must also be in  $\bar{C}_q(G - n)$ . If it were not, then there would be a path between  $X$  and  $Y$  that is active in  $\bar{C}_q(G - n)$  but rendered inactive by  $Z$  in  $\bar{C}_q(G)$ . But this is impossible, because an active path cannot be rendered inactive just by adding a node and its associated edges to the graph. Since  $G - n$  is an I-map of  $P - n$ , the relation  $R$  must also be contained in  $\bar{C}(P - n)$ , and since  $\bar{C}(P - n)$  is a subset of  $\bar{C}(P)$ ,  $R$  is also contained in  $\bar{C}(P)$ .  $\square$

For the remaining two cases, we first consider the special instance where  $n$  is a setting variable:

**Lemma 2.** Let  $R \in \bar{C}_q(G)$  be either of the form  $(nX \perp Y | Z)$  or  $(X \perp Y | nZ)$  and  $n$  is a setting variable. Then  $R \in \bar{C}(P)$ .

**Proof.** If either choice of  $R$  holds in  $G$ , then  $(X \perp Y | Z)$  must hold in  $G$ . If this were not true for  $R = (nX \perp Y | Z)$ , there would be an active path between  $X$  and  $Y$  conditional on  $Z$ , but this would imply an active path between  $nX$  and  $Y$  conditional on  $Z$ , contradicting  $R$ . It must also be true for  $R = (X \perp Y | nZ)$  because conditioning on a setting  $n$  cannot deactivate a previously active path. Since  $(X \perp Y | Z)$  holds in  $G$ , by lemma 1 it holds in  $\bar{C}(P)$ . Let  $Z_{\mathcal{O}}(Z_S)$  denote the outcomes (settings) in  $Z$ . Since  $n$  is a setting with no descendants, it is  $O'$ -detached from any subset of variables, for any set of outcomes  $O'$ ; in particular,  $(n \perp \mathbf{dt}_{Z_{\mathcal{O}}}(n) | Z_{\mathcal{O}})$  holds in the quantum input list  $Q_{\mathcal{O}}$ , hence in  $\bar{C}(P)$ . Noting that  $XYZ_S$  is a subset of  $\mathbf{dt}_{Z_{\mathcal{O}}}(n)$  and using the semi-graphoid axioms, we obtain:

$$\begin{aligned} (n \perp \mathbf{dt}_{Z_{\mathcal{O}}}(n) | Z_{\mathcal{O}}) &\in \bar{C}(P), \\ \Rightarrow (n \perp XYZ_S | Z_{\mathcal{O}}) &\in \bar{C}(P), \\ \Rightarrow (n \perp Y | XZ) &\in \bar{C}(P). \end{aligned}$$



Combining this with  $(X \perp Y|Z) \in \tilde{C}(P)$  by the contraction axiom we obtain  $(nX \perp Y|Z) \in \tilde{C}(P)$ , which also implies  $(X \perp Y|nZ) \in \tilde{C}(P)$ , so either way  $R$  is in  $\tilde{C}(P)$ .  $\square$

We now examine the last two cases under the assumption that  $n$  is an outcome.

**Lemma 3.** Let  $R \in \tilde{C}_q(G)$  be a relation of the form  $(nX \perp Y|Z)$ . Then  $R \in \tilde{C}(P)$ .

**Proof.** First, we partition the sets  $X, Y, Z$  into disjoint sets of outcomes and settings, e.g.  $Z = Z_O \cup Z_S$  where  $Z_O$  contains only outcomes and  $Z_S$  only settings.

Define the set  $\mathbf{zo}(O')$  as the members of  $Z_O$  that are connected to outcomes  $O'$  by a  $Z_O$ -chain. Let  $-\mathbf{zo}(O')$  denote its complement in  $Z_O$ . Next, consider  $O_x := nX_O \mathbf{zo}(nX_O)$ .

We can write the ancestors  $\text{an}(O_x)$  as the union of four disjoint sets,  $\text{an}(O_x) = A_X \cup A_Y \cup A_Z \cup A$ , where  $A_X := \text{an}(O_x) \cap X$ , and similarly for  $A_Y$  and  $A_Z$ . Any remaining members of  $\text{an}(O_x)$  not contained in any of  $X, Y, Z$  are contained in  $A$ . Likewise, let us decompose  $\neg_S \text{an}(O_x)$  into disjoint sets:  $\neg_S \text{an}(O_x) = B_X \cup B_Y \cup B_Z \cup B$  where e.g.  $B_X := \neg_S \text{an}(O_x) \cap X$  and similarly for  $B_Y, B_Z$  (see figure 8). Note that  $X_S = A_X \cup B_X$  and analogously for  $Y$  and  $Z$ . The CI relation  $(O_x \text{an}(O_x) \perp \mathbf{dt}_\emptyset(O_x) | \emptyset)$  must hold in  $Q$  and hence in  $\tilde{C}(P)$ . Using the above definitions:

$$\begin{aligned}
 (O_x \text{an}(O_x) \perp \mathbf{dt}_\emptyset(O_x) | \emptyset) &\in \tilde{C}(P) \\
 \Rightarrow (O_x \perp \neg_S \text{an}(O_x) [\mathbf{dt}_\emptyset(O_x)/S] | \text{an}(O_x)) &\in \tilde{C}(P) \\
 \Leftrightarrow (O_x \perp B_X B_Y B_Z B [\mathbf{dt}_\emptyset(O_x)/S] | A_X A_Y A_Z A) &\in \tilde{C}(P).
 \end{aligned} \tag{11}$$

The set  $-\mathbf{zo}(nX_O)$  must be a subset of  $\mathbf{dt}_\emptyset(O_x)/S$  (since the members of  $-\mathbf{zo}(nX_O)$  by definition cannot share an ancestor with  $O_x$ ). The same goes for  $Y_O$ , otherwise there would be a path connecting  $Y_O$  to  $nX_O$  on which every collider is in  $Z_O$  or has a directed path to  $Z_O$ , and they could not be  $q$ -separated in  $G$  as is required for



$R$  to be true. Hence, using the semi-graphoid axioms:

$$\begin{aligned} (O_x \perp B_X B_Y B_Z B \text{ dt}_{\emptyset}(O_x)/S | A_X A_Y A_Z A) &\in \bar{C}(P) \\ \Rightarrow (O_x \perp B_Y B_Z Y_O -\mathbf{zo}(nX_O) | X_S A_Y A_Z) &\in \bar{C}(P) \\ \Rightarrow (n \perp B_Y Y_O | XZ A_Y) &\in \bar{C}(P). \end{aligned} \quad (12)$$

No member of  $Y$  can be an ancestor of  $nX_O$  in  $Q$ , or else there would be a directed path from a setting in  $Y_S$  to an outcome in  $nX$  and they could not be  $q$ -separated by  $Z$  in  $G$ , contradicting our initial premise  $R$ . Therefore  $A_Y = \emptyset$  and  $B_Y = Y_S$ , and (12) implies  $(n \perp Y | XZ) \in \bar{C}(P)$ . The relation  $R$  implies  $(X \perp Y | Z) \in \bar{C}_q(G)$  and hence (by lemma 1):  $(X \perp Y | Z) \in \bar{C}(P)$ . Combining this with (12) and the semi-graphoid axioms, we obtain the desired result  $(nX \perp Y | Z) \in \bar{C}(P)$ .  $\square$

**Lemma 4.** Let  $R \in \bar{C}_q(G)$  be a relation of the form  $(X \perp Y | nZ)$ . Then  $R \in \bar{C}(P)$ .

**Proof.** Note that  $n$  cannot share an ancestor with both  $X_O \mathbf{zo}(X_O)$  and  $Y_O \mathbf{zo}(Y_O)$ , or else there would be a path connecting  $X_O$  to  $Y_O$  on which every collider has a descendant in  $nZ_O$ , preventing them from being  $q$ -separated in  $G$ . We therefore assume without loss of generality that  $n$  does not share an ancestor with  $Y_O \mathbf{zo}(Y_O)$ . By a similar argument, no member of  $X_O$  can share an ancestor with  $Y_O \mathbf{zo}(Y_O)$ ; hence  $nX_O -\mathbf{zo}(Y_O) \in \text{dt}_{\emptyset}(Y_O \mathbf{zo}(Y_O))$ . Now, either  $n$  shares an ancestor with  $X_O \mathbf{zo}(X_O)$ , or it does not. If it does, then  $Y$  cannot contain any ancestors of  $O_x$  (defined as in the previous Lemma). Furthermore  $Y_O$  cannot share any ancestors with  $O_x$ , since by assumption it shares no ancestors with  $nX_O$ , and if it shared an ancestor with  $\mathbf{zo}(nX_O)$  it could not be  $q$ -separated from  $X$  given  $nZ$ , contradicting the premise  $R$ . Hence we can use the same procedure as in lemma 3 to obtain the result  $(nX \perp Y | Z) \in \bar{C}(P)$ , and use semi-graphoid axiom 1.c. to obtain the desired result:  $(X \perp Y | nZ) \in \bar{C}(P)$ .

In the remaining case,  $n$  does not share an ancestor with  $X_O \mathbf{zo}(X_O)$  so we have the relation  $nY_O -\mathbf{zo}(X_O) \in \text{dt}_{\emptyset}(X_O \mathbf{zo}(X_O))$ . Let  $O_{xz} := X_O \mathbf{zo}(X_O)$  and consider the relation  $(O_{xz} \mathbf{an}(O_{xz}) \perp \text{dt}_{\emptyset}(O_{xz}) | \emptyset)$  that holds in  $Q$  and hence in  $\bar{C}(P)$ . Using the above properties, and the fact that  $Y$  cannot contain any ancestors of  $O_{xz}$  (for the usual reason that this would imply an active path between  $X$  and  $Y$  in  $G$ ), we obtain:

$$\begin{aligned} (O_{xz} \mathbf{an}(O_{xz}) \perp \neg_S \mathbf{an}(O_{xz}) \text{ dt}_{\emptyset}(O_{xz})/S | \emptyset) &\in \bar{C}(P) \\ \Rightarrow (X_O \mathbf{zo}(X_O) \perp Y_S nY_O -\mathbf{zo}(X_O) | X_S Z_S) &\in \bar{C}(P) \\ \Rightarrow (X_O \perp Y | nX_S Z) &\in \bar{C}(P). \end{aligned} \quad (13)$$

Let us partition  $Z_S = D_Z \cap E_Z$ , where  $D_Z$  contains the members of  $Z_S$  that are detached from  $X_S$  by  $nZ_O$ , and  $E_Z$  contains the rest. Consider the CI relation  $(X_S E_Z \perp \text{dt}_{nZ_O}(X_S E_Z) | nZ_O)$  that holds in  $Q$  and hence in  $\bar{C}(P)$ . The set  $\text{dt}_{nZ_O}(X_S E_Z)$  must contain  $\text{dt}_{nZ_O}(X_S)$ . If not, there would be a member of  $\text{dt}_{nZ_O}(X_S)$  that is not detached from  $X_S E_Z$ , hence it would be connected by an  $nZ_O$ -chain to  $E_Z$ . But since  $E_Z$  has a chain to  $X_S$ , it could not be a member of  $\text{dt}_{nZ_O}(X_S)$ , implying a contradiction. Hence  $(X_S E_Z \perp \text{dt}_{nZ_O}(X_S) | nZ_O) \in \bar{C}(P)$ . But note that  $Y$  must be detached from  $X_S$  by  $nZ_O$ , otherwise there would be a path connecting  $X$  to  $Y$  in  $G$  on which every collider has a directed path to  $nZ_O$ , contradicting  $R$ . Thus  $D_Z Y \in \text{dt}_{nZ_O}(X_S)$ , and we obtain:

$$\begin{aligned} (X_S E_Z \perp \text{dt}_{nZ_O}(X_S) | nZ_O) &\in \bar{C}(P) \\ \Rightarrow (X_S \perp Y | nZ) &\in \bar{C}(P). \end{aligned} \quad (14)$$

Combining this result with (13) and the contraction axiom 1.d, we obtain  $(X \perp Y | nZ) \in \bar{C}(P)$  as desired.  $\square$

Lemmas 1–4 together imply that  $G$  is an I-map of  $P$ , i.e.  $\bar{C}_q(G) \subseteq \bar{C}(P)$ , provided that  $G - n$  is an I-map of  $P - n$ . The latter condition can be guaranteed using the same logic:  $G - n$  is an I-map of  $P - n$  provided  $G - n - m$  is an I-map of  $P - n - m$ , where  $m$  is now the second-last variable in the chosen ordering. Continuing this process, every graph in the hierarchy is an I-map of its corresponding distribution provided that we can prove that  $G$  for a single variable is an I-map of a probability distribution  $P$  on a single variable. But  $\bar{C}_q(G) \subseteq \bar{C}(P)$  is trivially satisfied for a single variable, because both these sets are empty. This completes the proof of theorem 3.

The result of theorems 2 and 3 is that  $q$ -separation is sound and complete for quantum input lists. That is, if  $Q_{\emptyset}$  is a quantum input list, the DAG  $G$  generated by  $Q_{\emptyset}$  is a perfect map of the semi-graphoid closure  $\bar{Q}$ . Hence a CI relation follows from the DAG using  $q$ -separation if and only if it can be obtained from  $Q$  using the semi-graphoid axioms.

### 3.5. Correspondence to classical models

In a quantum circuit, one can obtain a ‘classical limit’ by restricting all states and operators to a subspace of Hilbert space. In order to define the classical limit of a QCM, we must ensure that, in this limit, we recover the assumptions listed in section 2. We expect that Reichenbach’s common cause principle will be recovered by restricting our operations to a classical subspace of Hilbert space, since this will rule out the possibility of entanglement. However, the causal Markov condition requires that direct causes ‘screen-off’ indirect causes, which will not in general be true after restricting the circuit to a classical subspace. To recover this principle, therefore, we need to transform the DAG of the QCM into a form that respects this property. This can be done by assigning the ancestors of every outcome (the ‘possible causes’ by assumption 2) to either direct or indirect causes. There may be many ways to do this, so we will adopt the simplest solution and make them all direct causes. This procedure is described in the following definition:

**Definition 16. Classical limit for DAGs.** Given a DAG  $G$  interpreted as a quantum network, the *classical limit* of  $G$  is a new DAG  $G_C$  obtained by the following procedure:

- (1) Draw a directed edge from every setting  $S_i$  (non-drain node) to every outcome (drain node)  $O_i$  that is descended from  $S_i$  in  $G_C$ , unless such an edge already exists.
- (2) Eliminate all edges that connect pairs of settings to each other.

The first step makes every setting a direct cause of every outcome that is descended from it, while the second step uses the fact that the settings are mutually independent to eliminate redundant edges. The resulting DAG is consistent with the causal structure of the original DAG in the following sense: there is a causal chain from one variable to another in  $G_C$  *only if* there existed such a chain in the DAG of the quantum network  $G$ . The screening-off property is enforced since there are no intermediate nodes left in the DAG. This allows us to recover  $d$ -separation in the classical limit:

**Theorem 4.** *Given any DAG  $G$  with classical limit  $G_C$ , every CI relation obtainable from  $G_C$  by  $q$ -separation is also implied by  $d$ -separation, i.e.  $\bar{C}_q(G_C) \subseteq \bar{C}_d(G_C)$ .*

**Proof.** Suppose  $(S_A \perp S_B | V)_q$  holds under  $q$ -separation for disjoint sets  $S_A, S_B, V$  where  $S_A, S_B$  are settings. Since all paths connecting two settings in  $G_C$  must contain at least one collider that is an outcome,  $q$ -separation implies that at least one of these colliders is not in  $V$ , but this also implies that the settings are  $d$ -separated. Similarly, if  $(S_A \perp O_B | V)_q$  holds under  $q$ -separation for a set of outcomes  $O_B$ , this implies that  $O_B$  is not descended from  $S_A$ . It follows that every path between them must contain at least one collider, and at least one of these colliders is not in  $V$ , which also implies  $d$ -separation. Finally, if  $(O_A \perp O_B | V)_q$  holds for two sets of outcomes, then every path connecting them must contain a collider that is not in  $V$ , again implying  $d$ -separation. □

Given this result, it is straightforward to convert a QCM into a CCM: we restrict the quantum model parameters to a classical subspace and obtain a classical circuit. This circuit defines a set of functions  $F$  that determine the values of the outcomes given the values of their ancestors (these become their parents in  $G_C$  obtained from  $G_Q$ ). The pair  $\{G_C, F\}$  then satisfies the requirements of a CCM.

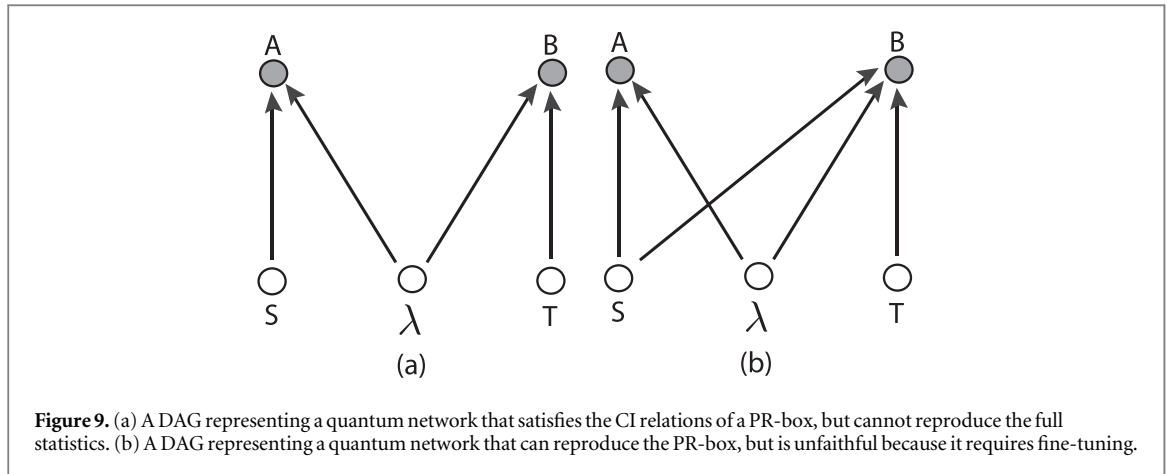
### 3.6. More general correlations

We note that the ‘Bell-type’ experiment described in section 2.6 also applies more generally to any joint probability distribution with settings and outcomes that obey the no-signalling criterion. In particular, one can find a joint probability distribution  $P$  on the variables  $A, B, S, T$  that satisfies the CI relations  $K$  (implied by setting independence and no-signalling) but which cannot be generated by any QCM defined on the same variables. For example, let all variables be binary variables taking values in  $\{0, 1\}$ , let  $\oplus$  represent addition modulo 2, and consider the joint distribution:

$$P(A, B, S, T) := P(S)P(T)P(A, B|S, T), \quad \text{where}$$

$$P(A, B|S, T) = \frac{1}{2}(1 \oplus A \oplus B)(ST \oplus 1) + \frac{1}{2}(A \oplus B)ST. \quad (15)$$

One can check that this distribution satisfies the CI relations  $K$  and that the probabilities sum to one. This distribution characterizes a *Popescu–Rorlich box* (PR-box), also called a nonlocal or non-signalling box [29].



A PR-box defines correlations that are stronger than quantum correlations, in the sense that they violate Bell's inequality to its algebraic maximum. How do super-quantum correlations fit into the present framework of QCMs?

There is no QCM on just the variables  $A$ ,  $B$ ,  $S$ ,  $T$  that provides a faithful explanation of the probability distribution  $P$ . This is trivially true because the constraints in  $K$  require that there is no directed edge from  $S$  to  $B$  or from  $T$  to  $A$  in the corresponding DAG (or else they would not be  $q$ -separated by any subset of variables) and no outgoing edges from  $A$  and  $B$ , which are outcomes—but this splits the DAG into two disconnected parts, which are necessarily independent, implying no correlations between  $A$  and  $B$ . As before, we can try to explain the correlations by introducing a hidden variable  $\lambda$  and extending the constraints to the set  $K'$  (see section 2.6). For a QCM, this is equivalent to supposing that the outcomes  $A$  and  $B$  can depend on a shared entangled resource, in a pure state specified by  $\lambda$ ; the corresponding DAG is shown in figure 9(a). It should therefore come as no surprise that a QCM based on this DAG fails to reproduce the distribution  $P$ . Whichever pure states one chooses for the values of  $\lambda$ , the resulting statistics obtained from the QCM must obey quantum mechanics and hence the joint probability generated by the QCM must be able to violate Bell's inequality only up to Tsirelson's bound.

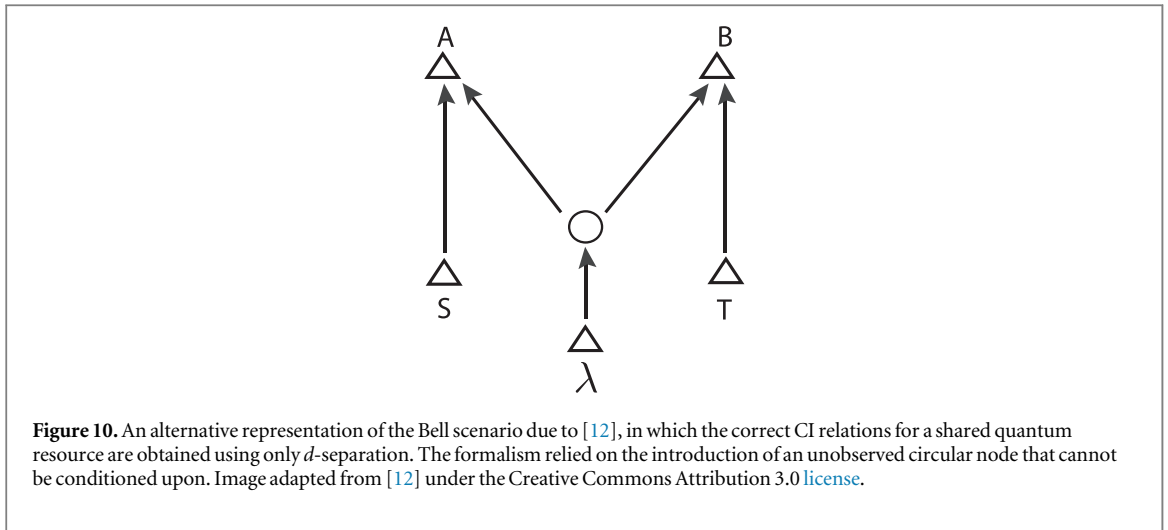
Of course, there is no reason to restrict ourselves to this DAG—just as in the classical case, we can consider DAGs such as the one shown in figure 9(b), in which there is a signal from  $S$  to  $B$ , representing a 'hidden' link in the underlying quantum network (*Note: if we interpret this DAG as a CCM, it also serves as an alternative to figure 9(b) for explaining quantum or PR-Box correlations using fine-tuning*). This additional link can be used to send a single bit of information from  $S$  to  $B$ , which can be exploited to simulate the desired probability distribution. But, just as in the classical case, this would imply that  $S$  and  $B$  should not be  $q$ -separated, so the CI relation  $(S \perp B | T\lambda) \in K'$  can only be due to fine-tuning of the model parameters and the explanation is not faithful. Just as a CCM cannot faithfully explain quantum correlations, a QCM also cannot faithfully explain super-quantum correlations.

Note that this observation follows from the definition of post-quantum theories which violate Bell's inequalities up to the algebraic maximum while obeying the same causal constraints. It is a topic for future work to find out whether information theoretic or entropic constraints characteristic of quantum correlations, such as those discussed in [17–20], can be incorporated into the graphical representation in a useful way.

## 4. Discussion and conclusions

Given a probability distribution, a CCM describes the causal connections of events compatible with the observed correlations under the assumption that these correlations are generated by classical physics. We have shown that the same can be done under the assumption that the underlying physics is quantum. We gave a suitable definition of a QCM on a DAG that represents a quantum network and is consistent with related work, namely [9]. We showed that it is possible to deduce the CI relations implied by the quantum network using a criterion called  $q$ -separation, which we proved to be sound and complete. In principle, an algorithm based on  $q$ -separation could be used to program an AI to make inferences about the connections between the components of a quantum network, given only the observed correlations between the variables in the network. It is left to future work to investigate whether such an algorithm presents any practical advantage over other approaches.

It is interesting to compare the approach to graph separation taken here to that of [12], in which a DAG representation of generalized probabilistic theories was proposed that retains the  $d$ -separation rule. If we



consider the Bell scenario in this latter framework, we obtain a graph like that of figure 10, where now  $\lambda$  represents the in-principle observable preparation of an entangled resource, as in our formalism, but there is an additional ‘unobserved’ node, depicted as a circle, representing the quantum nature of the resource. The remaining nodes are depicted as triangles to indicate that they are observed. The presence of the unobserved node ensures that the two outcomes are no longer  $d$ -separated by any subset of the *observed* variables  $S$ ,  $T$ ,  $\lambda$ .

Of course, if it were possible to condition on the unobserved node, we would not have progressed from a CCM. Rather, the purpose of these unobserved nodes is to restrict us to a special subset of the CI relations obtainable by  $d$ -separation from the graph, which is then proven to be the correct set of CI relations for a quantum network (or a generalized probabilistic theory). A possible advantage of retaining  $d$ -separation in this way is that the existing algorithms for extracting conditional independencies from a graph still apply and can be used in a practical setting by a computer program. However, it could be argued that this approach misses something of the structure that underlies the CI relations in a quantum setting, which is made explicit in the present work through the definition of  $q$ -separation. Thus, the present work is complementary in that it elucidates those constraints on the causal structure of quantum networks that remain after the factorization property of the RCCP is dropped.

We have argued that there still exist non-trivial physical constraints on the conditional independencies between variables in a general quantum network, even without the RCCP. Our observation that super-quantum correlations are differentiated from quantum correlations only by their model parameters (and not by their CI relations) in the Bell scenario indicates that the same constraints may apply generally, and that they stem from the Markovianity condition; this is supported by the work of [9]. Unfortunately, this also indicates that relaxations of the RCCP alone may not distinguish quantum theory from more general probabilistic theories. However, we have not proven that  $q$ -separation is sound and complete for a suitably chosen DAG formulation of any generalized probabilistic theory, so it is left as an open question, as is the question of whether one can generalize the RCCP to a principle that can distinguish quantum from general probabilistic theories, not just from CCMs.

Finally, we speculate that the type of causal model discussed here might present a starting point for investigations into the quantum nature of space-time. After all, if the space-time manifold of classical general relativity is to give way to a more fundamental structure at the Planck scale, then it seems plausible that this structure should consist of something like a causal network, which supplies the essential geometric information about some discrete set of fundamental events. It would be interesting to see whether such a construction could make connections to existing work on quantum gravity, such as spin foams and causal sets, and whether it can be generalized to include more exotic effects, such as closed time-like loops or the recently proposed phenomenon of ‘indefinite causality’ [6, 11].

## Acknowledgments

We thank M Pusey for drawing our attention to a flaw in an earlier draft of the paper (now fixed), and an anonymous referee for insightful comments. This work has been supported by the European Commission Project RAQUEL, the John Templeton Foundation, FQXi, and the Austrian Science Fund (FWF) through CoQuS, SFB FoQuS, and the Individual Project 2462.

## Appendix Proof that $q$ -separation satisfies the semi-graphoid axioms

### 1.a. Symmetry: $(X \perp Y|Z)_q \Leftrightarrow (Y \perp X|Z)_q$

Two sets  $X$  and  $Y$  are  $q$ -separated by  $Z$  iff every path connecting them is rendered inactive by  $Z$  according to definition 15. Since an undirected path between  $X$  and  $Y$  in the graph is manifestly symmetric (it is the same as saying ‘a path between  $Y$  and  $X$ ’), the symmetry property follows automatically.

### 1.b. Decomposition: $(X \perp YW|Z)_q \Rightarrow (X \perp Y|Z)_q$

If every path between the sets  $X$  and  $YW$  is rendered inactive by  $Z$ , then so is every path between  $X$  and any subset of  $YW$ . The decomposition property follows automatically.

### 1.c. Weak union: $(X \perp YW|Z)_q \Rightarrow (X \perp Y|WZ)_q$

The lhs implies all paths between  $X$  and  $Y$  are inactive conditional on  $Z$ . To obtain the rhs, it is sufficient to show that these paths remain inactive when we condition on  $W$  as well. This is trivially true for all settings in  $W$  (since by  $q$ -separation it is impossible to activate an inactive path by conditioning on a setting) so we need only prove the result for outcomes in  $W$ .

*Case I.* Suppose a path between specific nodes  $X' \in X$  and  $Y' \in Y$  is inactive due to the presence of a collider such that condition (iii) of definition 15 applies (in particular, this means there is at least one collider on the path). Divide the possibilities into the following special sub-cases:

- (a)  $X'$  and  $Y'$  are both settings,
- (b)  $X'$  is a setting and  $Y'$  an outcome,
- (c)  $X'$  is an outcome and  $Y'$  a setting,
- (d)  $X'$  and  $Y'$  are both outcomes.

Now suppose that conditioning on the set  $W$  renders this path between  $X'$  and  $Y'$  active. This can only happen if all colliders on the path are either in  $W$  or in  $Z$  or have a directed path to either of these. In cases (a) and (b) the assumed violation of conditions (i) and (ii) of  $q$ -separation imply that  $X'$  has a directed path to either  $W$  or  $Z$  or both; but then  $X'$  could not be  $q$ -separated from  $W$  given  $Z$ , contradicting the lhs. For case (c), the assumed violation of (ii) implies a directed path from  $Y'$  to  $W$ . But this implies a path connecting  $X'$  to  $W$ , and the assumed violation of (iii) implies that there are no colliders on this path that do not terminate in  $W$  or  $Z$ , so  $X'$  could not be  $q$ -separated from  $W$  given  $Z$  and we have a contradiction with the lhs. For case (d), the assumed violation of (iii) after conditioning on  $W$  again implies an active path between outcomes  $X$  and  $W$ , given  $Z$ . We conclude that conditioning on  $W$  cannot activate any path that is inactive due to condition (iii).

*Case II.* Suppose that a path between  $X'$  and  $Y'$  is inactive because of condition (ii) in definition 15. The only possible cases are (b) and (c). Now suppose that conditioning on  $W$  renders this path active. In case (b) this can only occur if there is a directed path from  $X'$  to  $W$ , but then they could not be  $q$ -separated given  $Z$ , contradicting the lhs. In case (c), this can only occur if there is a directed path from  $Y'$  to  $W$  and if condition (iii) of  $q$ -separation fails to hold between  $X'$  and  $Y'$  given  $W$ . However, these conditions together ensure an active path between  $X'$  and  $W$ , contradicting the lhs. We conclude that conditioning on  $W$  cannot activate any path that is inactive due to condition (ii).

*Case III.* Suppose that a path between  $X'$  and  $Y'$  is inactive because of condition (i) in definition 15, which implies that both are settings (case (a)). If conditioning on  $W$  activates this path, it implies that  $X'$  has a directed path to either  $W$  or  $Z$  or both, but then  $X'$  could not be  $q$ -separated from  $W$  given  $Z$ , contradicting the lhs. (In the case where  $X'$  has a directed path to  $Z$ , we can construct an active path from  $X'$  to  $W$  as follows. First, if the path between  $X'$  and  $Y'$  is inactive due to a collider not in  $Z$ , then Case I applies, so we can assume that any colliders on the path are in  $Z$ . Second, for the path to be inactive given  $Z$ ,  $Y'$  cannot have a directed path to any outcome in  $Z$ , but since the path is activated by  $ZW$ ,  $Y'$  must have a directed path to an outcome in  $W$ . Hence there is a path between  $X'$  and an outcome in  $W$  via  $Y'$ , on which any colliders are in  $Z$ , so  $X'$  cannot be  $q$ -separated from  $W$  given  $Z$ ). Hence conditioning on  $W$  cannot activate any path that is inactive due to condition (i).

From the above, we conclude that conditioning on  $W$  cannot activate any path that is previously inactive between  $X$  and  $Y$  given the premise  $(X \perp YW|Z)_q$ , hence axiom 1.c. holds.

### 1.d. Contraction: $(X \perp Y|ZW)_q$ and $(X \perp W|Z)_q \Rightarrow (X \perp YW|Z)_q$

A path that is active conditional on a set  $Z$  cannot be rendered inactive by any set that contains  $Z$ . This is because, if all of the conditions (i)–(iii) of  $q$ -separation are false conditional on  $Z$ , they remain false conditional on any larger set  $ZW$ . Therefore  $(X \perp Y|ZW)_q \Rightarrow (X \perp Y|Z)_q$ . Combining this with  $(X \perp W|Z)_q$  we conclude that every path between  $X$  and  $Y$  and every path between  $X$  and  $W$  is rendered inactive by  $Z$ . Thus every

path between  $X$  and  $YW$  is inactive given  $Z$ , so  $(X \perp YW | Z)_q$  follows (we are indebted to an anonymous referee for suggesting this proof).  $\square$

## References

- [1] Bell J 1976 *Epistemological Lett.* **9** 11–24
- [2] Wood C and Spekkens R 2015 *New J. Phys.* **17** 033002
- [3] Tucci R 2007 arXiv:[quant-ph/0701201](#)
- [4] Leifer M S and Spekkens R W 2013 *Phys. Rev. A* **88** 052130
- [5] Coecke B and Paquette E O 2009 *New Structures for Physics* (Berlin: Springer) pp 173–286
- [6] Oreshkov O, Costa F and Brukner C 2012 *Nat. Commun.* **3** 1092
- [7] Chiribella G, D'Ariano G M and Perinotti P 2009 *Phys. Rev. A* **80** 022339
- [8] Hardy L 2007 *J. Phys. A: Math. Theor.* **40** 3081
- [9] Fritz T 2014 arXiv:[1404.4812](#)
- [10] Laskey K 2007 *AAAI Spring Symposium: Quantum Interaction* **2007** 142–9
- [11] Oreshkov O and Cerf N 2014 arXiv:[1406.3829](#)
- [12] Henson J, Lal R and Pusey M 2014 *New J. Phys.* **16** 113043
- [13] Pearl J 2000 *Causality: Models Reasoning and Inference* (Cambridge: Cambridge University Press)
- [14] Arntzenius F 2010 *The Stanford Encyclopedia of Philosophy* (Stanford, CA: The Metaphysics Research Lab, Centre for the Study of Language and Information)
- [15] Cavalcanti E and Lal R 2014 *J. Phys. A: Math. Theor.* **47** 424018
- [16] Wiseman H M 2014 *J. Phys. A: Math. Theor.* **47** 424001
- [17] Dahlsten O C O, Lercher D and Renner R 2012 *New J. Phys.* **14** 063024
- [18] Al-Safi S W and Short A J 2011 *Phys. Rev. A* **84** 042323
- [19] Chaves C, Majenz C and Gross D 2015 *Nat. Commun.* **6** 5766
- [20] Pawłowski M *et al* 2009 *Nature* **461** 1101
- [21] Price H 2008 *Stud. Hist. Phil. Sci. B* **39** 752
- [22] Verma T and Pearl J 1988 *Proc. 4th Conf. Uncertainty in Artificial Intelligence (UAI1988)*
- [23] Ried K *et al* 2015 *Nat. Phys.* **11** 414–20
- [24] Chaves R *et al* 2014 *Proc. 30th Conf. Uncertainty in Artificial Intelligence (UAI2014)* pp 112–21
- [25] Geiger D and Pearl J 1988 *Proc. 4th Conf. Uncertainty in Artificial Intelligence (UAI1988)*
- [26] Druzdzel M and Simon H 1993 *Proc. 9th Conf. Uncertainty in Artificial Intelligence (UAI)* 93
- [27] Morimae T 2014 *Phys. Rev. A* **90** 010101
- [28] Miyazaki M H J and Murao M 2015 *Phys. Rev. A* **91** 052302
- [29] Popescu S and Rohrlich D 1994 *Found. Phys.* **24** 379