Title: Quantum causal models

Date: Jul 30, 2018 10:00 AM

URL: http://pirsa.org/18070048

Abstract: From a brief discussion of how to generalise Reichenbachâ€™s Principle of the Common Cause to the case of quantum systems, I will develop a formalism to describe any set of quantum systems that have specified causal relationships between them. This formalism is the nearest quantum analogue to the classical causal models of Judea Pearl and others. At the heart of the classical formalism lies the idea that facts about causal structure enforce constraints on probability distributions in the form of conditional independences. I will describe a quantum analogue of this idea, which leads to a quantum version of the three rules of Pearlâ€™s do-calculus. If time, I will end with some more speculative remarks concerning the significance of the work for the foundations of quantum theory.

PERIMETER PI INSTITUTE FOR THEORETICAL PHYSICS

UNIVERSITY OF OXFORD

Quantum Causal Models

Jonathan Barrett

Foundations of Quantum Mechanics

July 2018

Introduction

The framework of classical causal models describes classical random variables with specified causal relationships between them.

The causal relationships induce constraints on probability distributions.

The framework is useful in many contexts. It enables us to make inferences about causal structure in cases where we have observed data, but the causal structure is unknown.

This in turn enables us to make deductions about what will happen in alternative scenarios, e.g., if I intervene and fix a variable to have the value that I want, what happens to the other variables? It also enables a rigorous account of counterfactual statements.

The main aim of this talk is to describe a framework for quantum causal models.

Summary

1) The classical notion of common cause.
2) The quantum notion of common cause.
3) Causal models.
4) Independence and causal structure.
5) Conclusions.

Assume finitely-valued random variables and finite dimensional Hilbert spaces throughout!

Reichenbach's principle

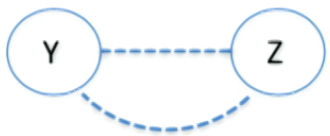

Y, Z are classical random variables.

Suppose they are *correlated*, i.e., $P(YZ) \neq P(Y)\,P(Z)$.

Then: one variable is a cause of the other, or there is a common cause, or both:

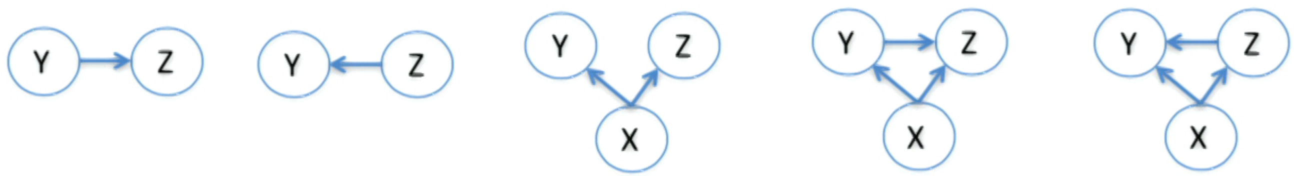

Reichenbach's principle

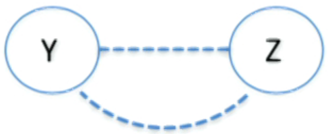

This one's special, because in this case, the principle also implies a constraint on the probabilities

Y, Z are classical random variables.

Suppose they are *correlated*, i.e., P(YZ) ≠ P(Y) P(Z)

Then: one variable is a cause of the other, or there is a common cause, or both:

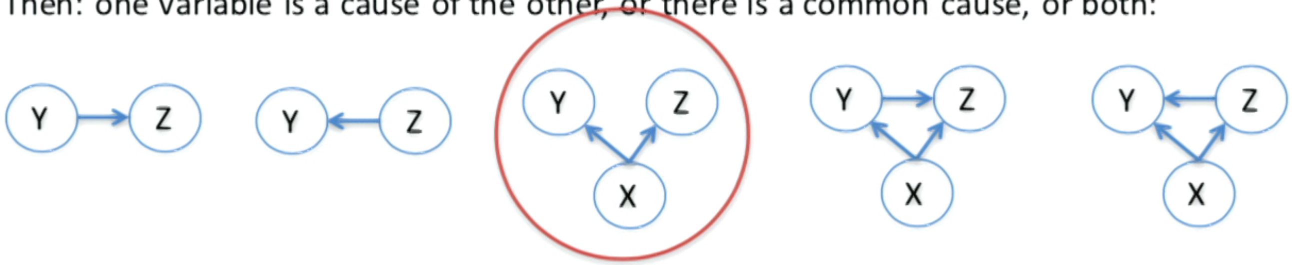

Reichenbach's principle

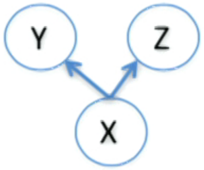

If X is a complete common cause of Y and Z, and Y is not a cause of Z and Z is not a cause of Y, then:

Y and Z are *conditionally independent given X*: $P(YZ|X) = P(Y|X)P(Z|X)$

A question

Following J. Pearl, *Causality*, let us take:

-- causal relationships between variables to be facts about the world (*ontic*)
-- probabilities to be the degrees of belief of a rational agent (*epistemic)*

Then a question arises. Suppose a rational agent takes the causal relations to be like this:

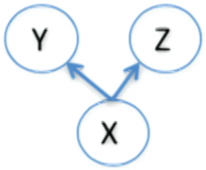

Why should the rational agent arrange their beliefs such that Y and Z are conditionally independent given X?

If the agent does not do this, are they *irrational*? Will they lose bets?

Dilation to functions

A (partial) answer:

Suppose that the agent *cannot observe all variables* and takes it that *causal relationships are fundamentally deterministic, i.e., functional.*

Suppose that the agent takes the functional relationships to be as follows:

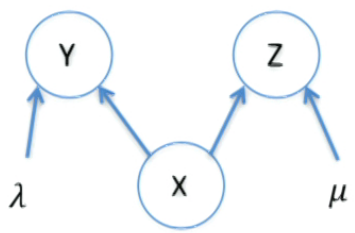

$$Y = f_Y(\lambda, X) \quad Z = f_Z(\mu, X)$$

i.e.,
Y does not depend on μ and Z does not depend on λ.

In this situation we will say that X is the complete common cause of Y and Z.

Suppose further that the agent assigns $P(\lambda X \mu) = P(\lambda)P(X)P(\mu)$.
Then it follows that the agent will assign $P(YZ|X) = P(Y|X)P(Z|X)$.

Dilation to functions

This goes in the other direction too. Hence:

Classical Reichenbach Theorem

Given a conditional distribution, *P(YZ|X)*, the following are equivalent:

(i) It is possible to define random variables λ, μ, and functions f_Y, f_Z, such that $Y = f_Y(\lambda, X), Z = f_Z(\mu, X)$, and $P(\lambda X \mu) = P(\lambda)P(X)P(\mu)$.

(ii) $P(YZ|X) = P(Y|X)P(Z|X)$

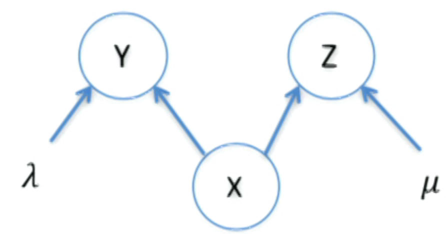

Same thing in circuit language

Think of *P(YZ|X)* as a channel.

P(YZ|X) = P(Y|X)P(Z|X) if and only if the channel can be dilated to a function g such that:

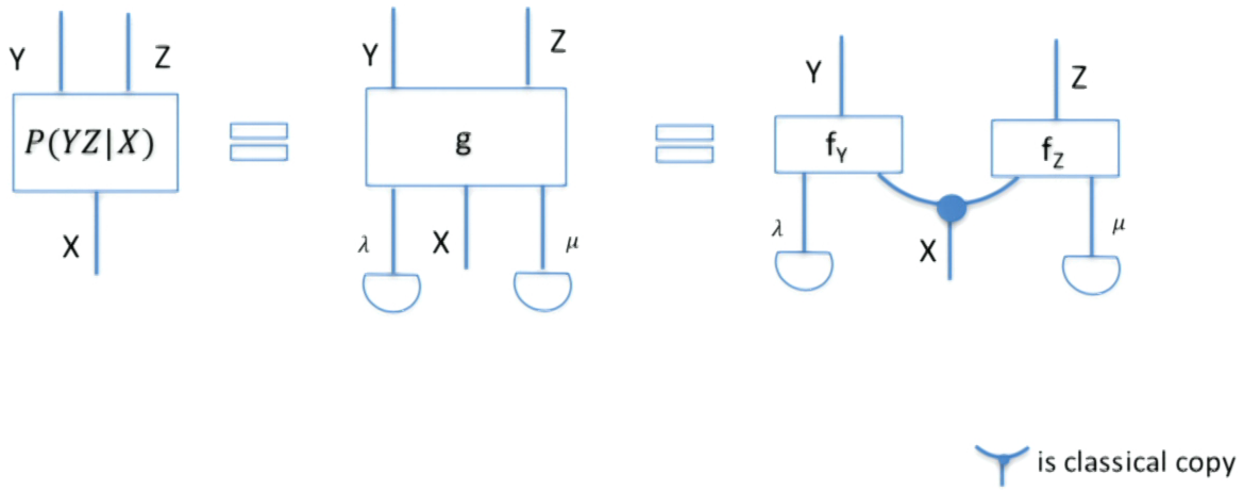

2
The quantum notion of common cause

Suppose that we take the causal situation to be as shown.

What do the arrows mean?

Following the classical discussion, we expect the arrows to be telling us, that A is (in some sense) a complete common cause of B and C.

Suppose that there is some quantum channel from A to BC.

We should then expect that "A is the complete common cause of B and C" places a constraint on this channel, analogous to the classical factorisation of P(YZ|X).

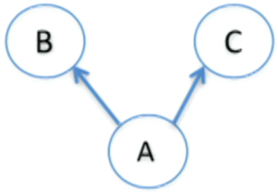

Notation:

Consider a quantum channel, with input *A* and output *B*, corresponding to a CP map *E*:

$$\rho_B = E(\rho_A).$$

Let the Choi-Jamiołkowski-isomorphic operator be given by:

$$\rho_{B|A} = \sum_{ij} E(|i\rangle_A\langle j|) \otimes |i\rangle_{A^*}\langle j|$$

$\rho_{B|A}$ is a positive operator, with $Tr_B(\rho_{B|A}) = I_{A^*}$.

Definition:

For a generic bipartite unitary *U*:

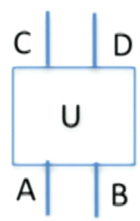

say that *B does not influence C* if:

- for all inputs ρ_A, the marginal ρ_C is independent of ρ_B
- equivalently, $Tr_D\ \rho_{CD|AB} = \rho_{C|A} \otimes I_B$

Definition:

Given a unitary U:

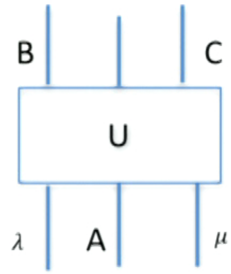

say that A is the *complete common cause* of B and C if:

- μ does not influence B
- λ does not influence C

Quantum Reichenbach Theorem (J.-M. Allen, et al. arXiv:1609.09487) :

Given $\rho_{BC|A}$, the following are equivalent:

(i) there exists a unitary dilation of $\rho_{BC|A}$, with latent systems λ and μ, such that A is the complete common cause of B and C, and $\rho_{\lambda A\mu} = \rho_\lambda \otimes \rho_A \otimes \rho_\mu$

(ii) $\rho_{BC|A} = \rho_{B|A}\rho_{C|A}$

NB Each of the following is also equivalent to the two conditions above:

(iii) I (B : C | A) = 0 (evaluated on the positive operator $(1\backslash \mathrm{d}_A)\rho_{BC|A}$)

(iv) $H_A = \bigoplus_i \left(H^i_{A_L} \otimes H^i_{A_R}\right)$, $\rho_{BC|A} = \sum_i \rho_{B|A^i_L} \otimes \rho_{C|A^i_R}$

Cf P. Hayden et al., Comm. Math. Phys. **246**, 359 (2004).

Same thing in circuit language

$\rho_{BC|A} = \rho_{B|A}\rho_{C|A}$ if and only if $\rho_{BC|A}$ can be dilated to a unitary U such that:

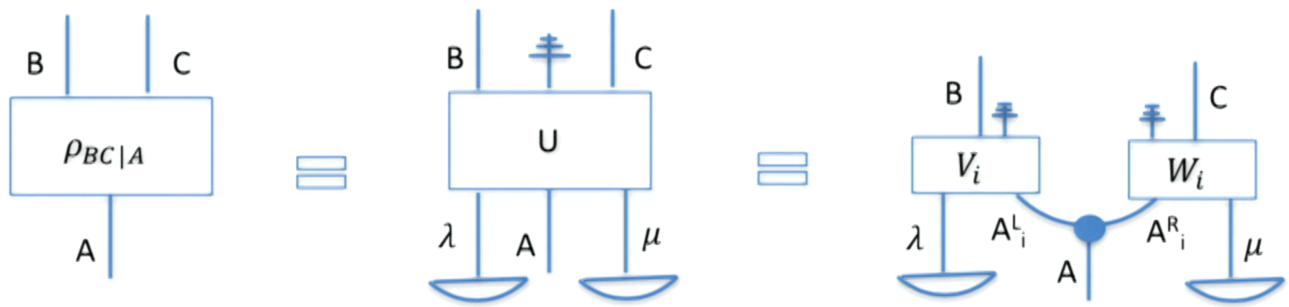

In quantum theory, there is no copy. So what is this?

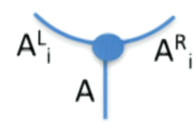

3
Causal models

Classical causal models

A *directed acyclic graph* with random variables on nodes:

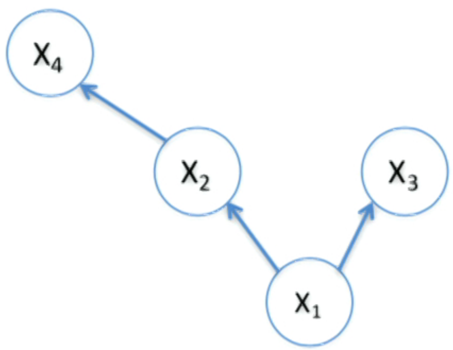

A set of conditional probabilities:
For each *i*, $P(X_i|Pa_i)$

Joint distribution:
$P(X_1 \dots X_k) = \Pi_i \, P(X_i|Pa_i)$

Pa_i denotes the parents of X_i, that is the set of nodes X_k such that there is an arrow from X_k to X_i.

Quantum causal models

A *directed acyclic graph.*
Each node is associated with a Hilbert space $H_i \otimes H_i^*$

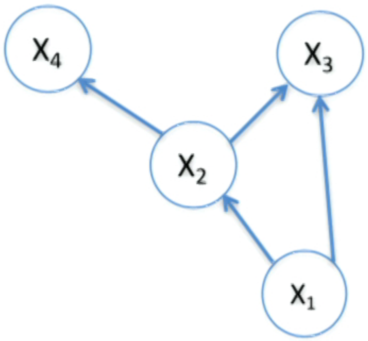

- A set of channels: $\rho_{X_i|Pa_i} \in B(H_i \otimes_{k \in Pa_i} H_k^*)$

such that for all *i,j* $[\rho_{X_i|Pa_i}, \rho_{X_j|Pa_j}] = 0$.

- Form a process matrix by taking the product of these channel operators.

- E.g., for the graph on the left:

$$\sigma = \rho_{X_4|X_2}\, \rho_{X_3|X_2 X_1}\, \rho_{X_2|X_1}\, \rho_{X_1}$$

- In general, $\sigma \in B\,(\otimes_i\, (H_i \otimes H_i^*))$.

Quantum Reichenbach Theorem (J.-M. Allen, et al. arXiv:1609.09487) :

Given $\rho_{BC|A}$, the following are equivalent:

(i) there exists a unitary dilation of $\rho_{BC|A}$, with latent systems λ and μ, such that A is the complete common cause of B and C, and $\rho_{\lambda A \mu} = \rho_\lambda \otimes \rho_A \otimes \rho_\mu$

(ii) $\rho_{BC|A} = \rho_{B|A}\rho_{C|A}$

NB Each of the following is also equivalent to the two conditions above:

(iii) I (B : C | A) = 0 (evaluated on the positive operator $(1\backslash \mathrm{d}_A)\rho_{BC|A}$)

(iv) $H_A = \bigoplus_i \left(H^i_{A_L} \otimes H^i_{A_R}\right)$, $\rho_{BC|A} = \sum_i \rho_{B|A^i_L} \otimes \rho_{C|A^i_R}$

Cf P. Hayden et al., Comm. Math. Phys. **246**, 359 (2004).

Remarks

- The approach of two Hilbert spaces per node, and the resulting σ operator, has much in common with the multi-time formalism, process matrices, and quantum combs. The novel aspect is the constraints that we get on σ from a particular causal structure.
- Allows for possibility of interventions at nodes. An intervention corresponds to a quantum instrument, mediating between the input Hilbert space (H_i) and the output Hilbert space (H_i^*). Joint probabilities for the outcomes of these interventions are obtained from σ via a trace rule.
- No intervention at a node corresponds to identity channel between input and output Hilbert spaces. Marginalization in this case corresponds to tracing out both Hilbert spaces, with an operator corresponding to the identity channel multiplying σ.

Justifying the definition

- We have already argued that in the simple case of 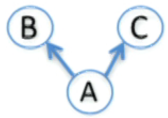

, the DAG is to be interpreted in terms of an underlying unitary evolution in which A is the complete common cause of B and C.

- It followed, by one direction of the quantum Reichenbach theorem, that the channel $\rho_{BC|A}$ should satisfy $\rho_{BC|A} = \rho_{B|A}\, \rho_{C|A}$.

- This leads to a process matrix of the form $\sigma = \rho_{B|A}\, \rho_{C|A}\, \rho_A$, which is consistent with the general definition on the previous slide.

- A generalization of the quantum Reichenbach theorem plays a similar role in the case of an arbitrary DAG, which is interpreted in terms of an underlying unitary circuit.

The quantum Reichenbach theorem for arbitrary causal structures

Consider a DAG G, with nodes $X_1, \ldots, X_n$, and a Hilbert space $H_i \otimes H_i^*$ for each X_i.
Consider also a positive operator $\sigma \in B(\bigotimes_i H_i \otimes H_i^*)$.

Then (JB, R. Lorenz, O. Oreshkov, forthcoming):

σ is of the form $\sigma = \Pi_i \rho_{X_i|Pa_i}$ for a pairwise commuting set of channel operators $\rho_{X_i|Pa_i}$.

⟺

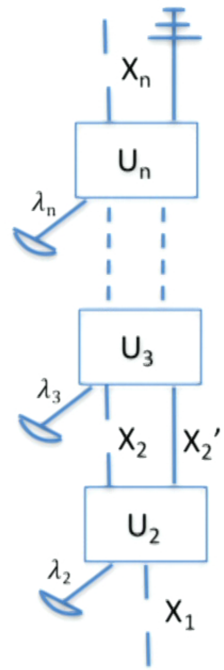

There exists a unitary comb, with causal structure corresponding to the DAG G, such that σ is returned when we marginalize over latent variables.

4
Independence and causal structure

Conditional independence in classical causal models

- In classical causal models, the structure of the DAG places constraints on the joint probability distribution in the form of conditional independences.

- In fact (turning things around) these conditional independences allow us to make inferences about what the underlying causal structure might be in cases where we don't know it, but do have some observational data. This in turn allows us to answer important questions like: what if, next time round, I don't just observe these variables, but actively intervene, and fix one so that it has the value I want it to have? What happens to the other variables?

- Pearl's *do-calculus* formalises some of this.

Definition:

Consider a DAG G. Let S, T and U be disjoint subsets of nodes of G. A path from S to T is an undirected path in the DAG, which starts on an S node and ends on a T node. A path from S to T is blocked by U if any of the following hold:

(i) The path contains a fork at a node in U.
(ii) The path contains a traversal at a node in U.
(iii) The path contains a collider at a node which is not in U, and which does not have any descendants in U.

Say that *S and T are d-separated by U* if all paths from S to T are blocked by U.

Theorem (see, e.g., Pearl, *Causality*):

d-separation is sound and complete for P(ST|U) = P(S|U) P(T|U).

Classical split node models

A *directed acyclic graph.*

Each node is associated with two random variables, X_i^{in} and X_i^{out}.

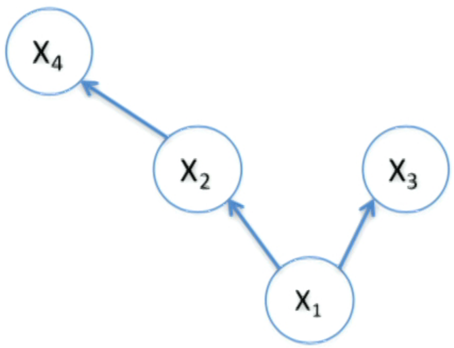

A set of conditional probabilities:
For each *i*, $P(X_i^{in}|Pa_i^{out})$

Classical "process matrix"
$K = \Pi_i \, P(X_i^{in}|Pa_i^{out})$

Formally, K is a conditional probability distribution over input variables, given output variables.

Allows for joint probabilities of outcomes to be calculated, when agents make arbitrary interventions at nodes.

If there is no intervention at a node, put a delta so that X^{out} takes the same value as X^{in}

Independence in classical split node models

Reminder: Given three random variables, X,Y,Z, the following are equivalent ways of defining "Y and Z are conditionally independent given X":

P(YZ|X) = P(Y|X) P(Z|X)
I(Y:Z|X) = 0
P(XYZ) P(X) = P(YX) P(ZX)
P(XYZ) = $\alpha(YX)\beta(ZX)$, for real valued functions α and β.

Definition: Given a split node classical causal model, consider disjoint subsets of nodes S, T, U.

Say that *S and T are strongly independent relative to U* if $K_{STU} = \alpha_{SU}\ \beta_{TU}$.

Equivalent operational statement: for all possible interventions at S, T and U nodes, the resulting joint distribution over input and output variables satisfies:

$$P\left(S^{in}S^{out}T^{in}T^{out}\,|U^{in}U^{out}\right) = P\left(S^{in}S^{out}\,|U^{in}U^{out}\right)P(T^{in}T^{out}\,|U^{in}U^{out})$$

Example of a situation where *weak independence* holds but not *strong independence*

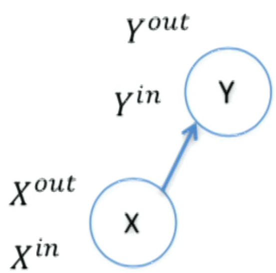

$Y^{in} = f(X^{out})$, where f(0) = 0, f(1) = 0, f(2) = 1.

$$P\left(X^{in} = 0\right) = P\left(X^{in} = 1\right) = 1/2.$$

With no interventions, Y^{in} = 0 with certainty, hence P($X^{in}Y^{in}$) = P(X^{in})P(Y^{in})
Weak independence holds.

But it is false that the classical process matrix can be written $K_{ST} = \alpha_S \, \beta_T$.
(Consider: an agent at X can signal to Y if they want to, by sometimes fixing X^{out} = 2.)
Strong independence fails.

Independence in classical split node models

Theorem (JB, R. Lorenz, O. Oreshkov, forthcoming)

d-separation is sound and complete for strong relative independence.

Theorem (JB, R. Lorenz, O. Oreshkov, forthcoming)

d-separation is sound and complete for weak relative independence.

Independence in quantum causal models

Definition: Given a quantum causal model, consider disjoint subsets of nodes S, T, U.

Say that *S and T are strongly independent relative to U* if $\sigma_{STU} = \alpha_{SU}\,\beta_{TU}$,
for Hermitian operators α and β.

Theorem (JB, R. Lorenz, O. Oreshkov, forthcoming): d-separation is sound and complete for quantum strong relative independence.

Remarks

- In the case of σ being diagonal with respect to a product basis (i.e., everything is classical), a quantum causal model reduces to a split node classical model, and quantum strong relative independence reduces to classical strong relative independence.

- In the case of a tripartite state, ρ_{ABC}, the definition of quantum strong relative independence reduces to the usual definition of quantum conditional independence:

  $\rho_{ABC} = \alpha_{BA}\beta_{CA}$ iff $I(B{:}C|A)=0$

Conclusions

- Following a close look at the notions of common cause and independence, we have given a definition for quantum causal models. NB The most similar approach in the literature is probably that of F. Costa and S. Shrapnel, New J. Phys. **18**, 063032 (2016). But there are some significant differences.
- A look at classical split node models reveals various senses in which two sets of nodes might be regarded as independent relative to a third.
- At least one of these (strong relative independence) has a quantum analogue, with a d-separation theorem.
- In fact, this theorem is a special case of the first rule of the *quantum do-calculus*, which supplies quantum analogues for all three rules of Pearl's classical do-calculus. (But I haven't discussed this!)

More foundational and/or speculative remarks

- The classical formalism has a natural interpretation, wherein causal structure is explained in terms of underlying functional relationships between variables. The functional relationships are taken to be facts about the world (*ontic*). Probabilities arise when an agent does not know the values of all variables, hence expresses degrees of belief with a probability distribution (*epistemic*).

- The existence of a compelling quantum analogue to the classical formalism, with unitaries replacing functions, lends support to the view that the unitaries (and the causal structure they define) are ontic, and that the positive operators are epistemic.

- But the positive operators express … information about what?

- Outcomes of interventions or the ontic states of some underlying theory?

More foundational and/or speculative remarks

- Is there a quantum analogue of weak relative independence?
- In the classical split node case, weak relative independence essentially expresses the ordinary conditional independence of the variables (in the case that no one intervenes).
- In the quantum case, in the absence of "some underlying theory", there are no underlying variables. But if we had a natural analogue of weak relative independence, we could see this as a clue, or a constraint on the underlying theory: weak relative independence expresses some sort of statistical conditional independence relation that an agent's information about the underlying ontic states should satisfy.