Quantum Bayesian Networks: Compositionality and Typing via Linear Logic

Rémi Di Guardia, Thomas Ehrhard, Claudia Faggian

Journées PPS, 22 May 2026

To be published in FSCD

Quantum Bayesian Networks: Compositionality and Typing via Linear Logic

Rémi Di Guardia, Thomas Ehrhard, Claudia Faggian

Journées PPS, 22 May 2026

To be published in FSCD

Quantum – Quick & Dirty

Quantum for this talk:

- **Data/Object** = state in a *Hilbert Space* $\mathcal{H}$ respecting some properties
- **Operation/Morphism** = linear maps between Hilbert Spaces preserving the properties
- **Tensor** product $\otimes$
- **Measure** from a Hilbert Space to its dimension $\mathcal{H} \to X$, probabilistic & modify the input state

Main differences with classical:

- states cannot be cloned/duplicated/broadcast
- given states $q_1 \otimes q_2$, measuring q_1 may modify q_2

$$|0\rangle \langle 0| = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$$

$$\mathcal{E} : \rho \mapsto |0\rangle \langle 0| \rho |0\rangle \langle 0|$$

Quantum – Bell's Experiment

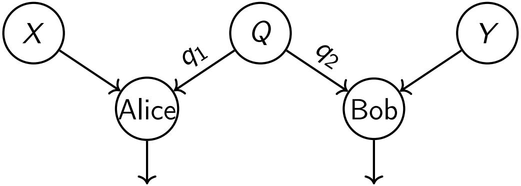

$$
\begin{aligned}
Q &= \textit{Preparation of the entangled } q_1 \textit{ and } q_2\\
X &= \textit{Flip a coin}\\
Alice &= \textit{Measure on } q_1 \textit{ parameterized by the value of } X
\end{aligned}
$$

Question: What is $Pr(Alice = 0, Bob = 0 \mid X = 1, Y = 0)$?

Quantum – Bell's Experiment

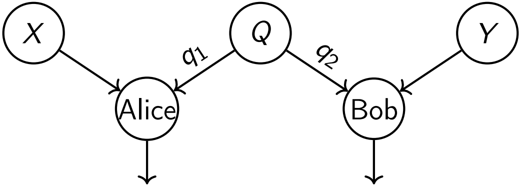

$$\begin{aligned} Q &= \textit{Preparation of the entangled } q_1 \textit{ and } q_2 \\ X &= \textit{Flip a coin} \\ Alice &= \textit{Measure on } q_1 \textit{ parameterized by the value of } X \end{aligned}$$

Question: What is $Pr(Alice = 0, Bob = 0 \mid X = 1, Y = 0)$?

General Problem: Considering a quantum system, what is the probability of obtaining a given measurement?

Bayesian Networks

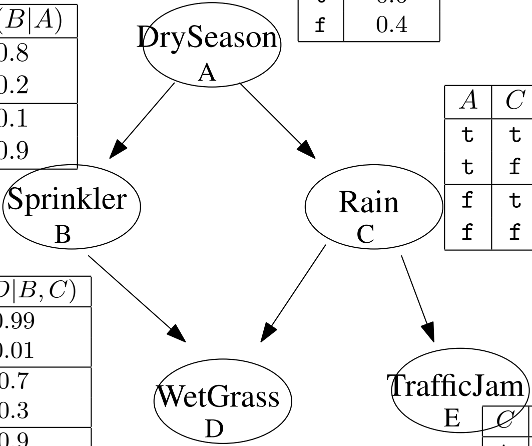

A	$Pr(A)$
t	0.6
f	0.4

A	B	$Pr(B\|A)$
t	t	0.8
t	f	0.2
f	t	0.1
f	f	0.9

A	C	$Pr(C\|A)$
t	t	0.2
t	f	0.8
f	t	0.75
f	f	0.25

C	B	D	$Pr(D\|B,C)$
t	t	t	0.99
t	t	f	0.01
t	f	t	0.7
t	f	f	0.3
f	t	t	0.9
f	t	f	0.1
f	f	t	0.01
f	f	f	0.99

C	E	$Pr(E\|C)$
t	t	0.7
t	f	0.3
f	t	0.1
f	f	0.9

Quantum Bayesian Networks

Definition: Quantum Bayesian Networks [HLP14]

DAG with a family of **quantum operations** per node such that composing these yields a probability distribution.

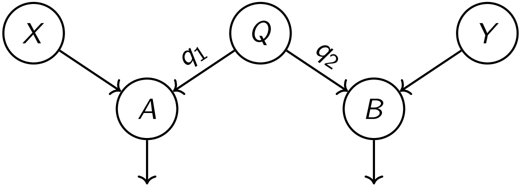

Q	:	$\mathbb{C} \to \mathcal{H}_1 \otimes \mathcal{H}_2$
	$=$	Preparation of the entangled q_1 and q_2
X	:	a map $f_x : \mathbb{C} \to \mathbb{C}$ for each value x of X (with $\sum_x f_x(1) = 1$)
	$=$	Flip a coin
A	:	a map $g_{x,a} : \mathcal{H}_1 \to \mathbb{C}$ for each value x of X and a of A (with $\sum_a g_{x,a}(q_1) = 1$)
	$=$	Measure on q_1 parameterized by the value of X

Limits

Problems: two wished results we do not have

- **Compositionality:** Given a decomposition of a network, is the data of the full network obtained from the data of each part?
  $\longrightarrow$ a big advantage of a graphical syntax, and a main property of Bayesian networks

- **Modularity:** Given two parts, can they compose to give a network?
  $\longrightarrow$ the result must be a DAG

Our contributions:
Solutions for these two problems, by giving another presentation of Quantum Bayesian Networks

Compositionality Limit

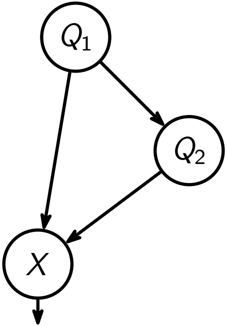

$$\phi^{Q_1} : \mathbb{C} \to \mathcal{H}_1 \otimes \mathcal{H}_2$$
$$\phi^{Q_2} : \mathcal{H}_2 \to \mathcal{H}_3$$
$$\phi^{X}(x) : \mathcal{H}_1 \otimes \mathcal{H}_3 \to \mathbb{C}$$

Whole graph:

$$\begin{aligned}\phi^{Q_1,Q_2,X}(x) &: \mathbb{C} \to \mathbb{C} \\ &= \phi^{X}(x) \circ (id_{\mathcal{H}_1} \otimes \phi^{Q_2}) \circ \phi^{Q_1}\end{aligned}$$

Compositionality Limit

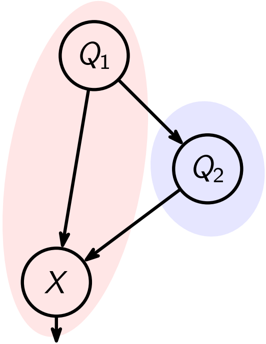

$$\phi^{Q_1} : \mathbb{C} \to \mathcal{H}_1 \otimes \mathcal{H}_2$$
$$\phi^{Q_2} : \mathcal{H}_2 \to \mathcal{H}_3$$
$$\phi^{X}(x) : \mathcal{H}_1 \otimes \mathcal{H}_3 \to \mathbb{C}$$

Whole graph:

$$\phi^{Q_1,Q_2,X}(x) : \mathbb{C} \to \mathbb{C}$$
$$= \phi^{X}(x) \circ (id_{\mathcal{H}_1} \otimes \phi^{Q_2}) \circ \phi^{Q_1}$$

$\phi^{Q_1,X}(x)$: $\mathcal{H}_3 \to \mathcal{H}_2$ $\qquad$ ϕ^{Q_2}: $\mathcal{H}_2 \to \mathcal{H}_3$

$\implies \phi^{Q_1,X,Q_2}(x)$ cannot simply be the **composition** of $\phi^{Q_1,X}(x)$ and ϕ^{Q_2}!
(gives either $\mathcal{H}_3 \to \mathcal{H}_3$ or $\mathcal{H}_2 \to \mathcal{H}_2$)

Idea: Functions do not work well at graph level, matrices would be better (also in Bayesian networks: factors instead of conditional probabilities)

Modularity Limit

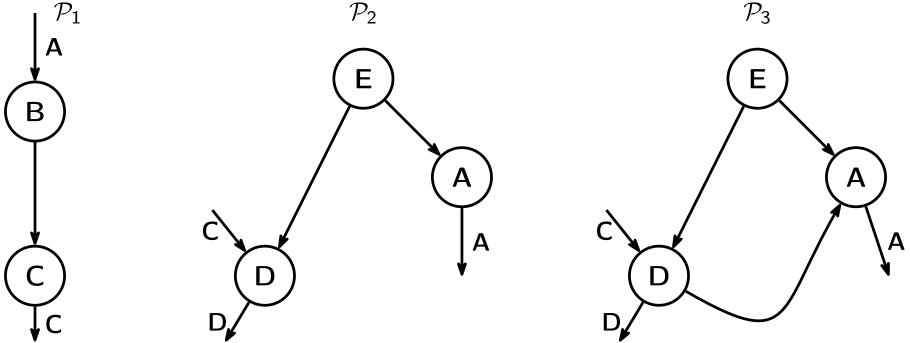

- Part $\mathcal{P}_1$ waits for input A and outputs C
- Parts $\mathcal{P}_2$ and $\mathcal{P}_3$ wait for input C and output A and D

Question: Is it "legal" to branch $\mathcal{P}_1$ to $\mathcal{P}_2$? $\mathcal{P}_1$ to $\mathcal{P}_3$?

Modularity Limit

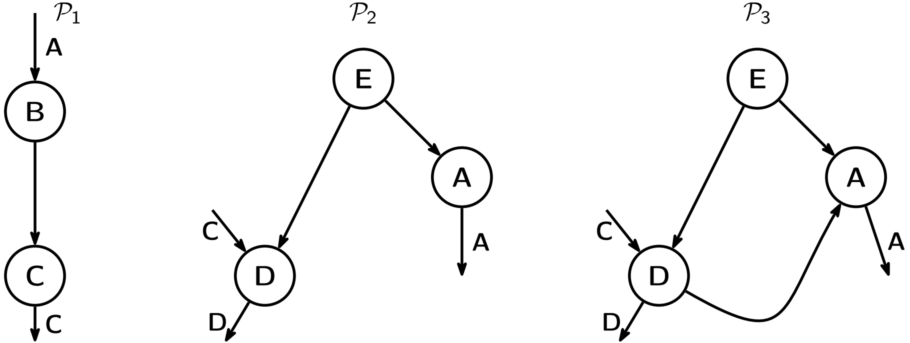

- Part $\mathcal{P}_1$ waits for input A and outputs C
- Parts $\mathcal{P}_2$ and $\mathcal{P}_3$ wait for input C and output A and D

Question: Is it "legal" to branch $\mathcal{P}_1$ to $\mathcal{P}_2$? $\mathcal{P}_1$ to $\mathcal{P}_3$?
$\longrightarrow$ $\mathcal{P}_1 \cup \mathcal{P}_2$ is a QBN but $\mathcal{P}_1 \cup \mathcal{P}_3$ is *not* (cycle $A \to B \to C \to D \to A$)

Idea: Inputs & outputs are insufficient, we need a **type**

Plan

- Compositionality by Quantum Factors

- Modularity by Typing

Quantum Factors

Instead of associating quantum operations to a node, we associate a:

Definition: Quantum Factor

Consider random variables $(X_1, \ldots, X_n)$ and Hilbert spaces $(\mathcal{H}_1, \ldots, \mathcal{H}_m)$. A **Quantum Factor** on $(X_1, \ldots, X_n, \mathcal{H}_1, \ldots, \mathcal{H}_m)$ is a function ϕ from $\prod_{i=1}^{n} Val(X_i)$ to *positive* matrices in $\bigotimes_{j=1}^{m} \mathcal{H}_j$.

Quantum Factors

Instead of associating quantum operations to a node, we associate a:

Definition: Quantum Factor

Consider random variables $(X_1, \ldots, X_n)$ and Hilbert spaces $(\mathcal{H}_1, \ldots, \mathcal{H}_m)$. A **Quantum Factor** on $(X_1, \ldots, X_n, \mathcal{H}_1, \ldots, \mathcal{H}_m)$ is a function ϕ from $\prod_{i=1}^n Val(X_i)$ to *positive* matrices in $\bigotimes_{j=1}^m \mathcal{H}_j$.

Equipped with a **product** $\odot$ and with a **sum** $\sum$, such that for ϕ_1 and ϕ_2 respectively on $(\mathbb{X}_1, \mathbb{H}_1)$ and $(\mathbb{X}_2, \mathbb{H}_2)$, $\phi_1 \odot \phi_2$ is on $(\mathbb{X}_1 \cup \mathbb{X}_2, \mathbb{H}_1 \,\Delta\, \mathbb{H}_2)$.

Only quantum: get $\otimes$-networks, $\odot$ = their product and $\sum$ = partial trace.
Only classical: get factors from Bayesian networks, $\odot$ = their product and $\sum$ = their sum.

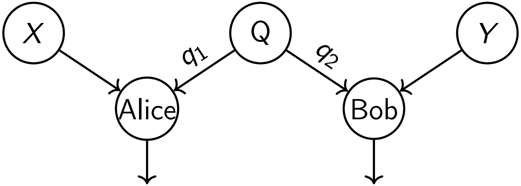

Compositionality with Quantum Factors

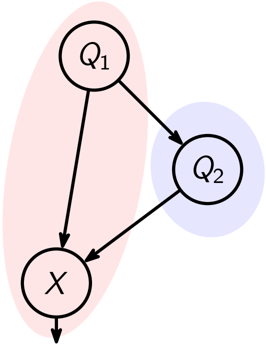

$$\phi^{Q_1} : \mathbb{C} \to \mathcal{H}_1 \otimes \mathcal{H}_2$$
$$\phi^{Q_2} : \mathcal{H}_2 \to \mathcal{H}_3$$
$$\phi^{X}(x) : \mathcal{H}_1 \otimes \mathcal{H}_3 \to \mathbb{C}$$

Whole graph:

$$\phi^{Q_1,Q_2,X}(x) : \mathbb{C} \to \mathbb{C}$$
$$= \phi^{X}(x) \circ (id_{\mathcal{H}_1} \otimes \phi^{Q_2}) \circ \phi^{Q_1}$$

$\phi^{Q_1,X}(x)$: $\mathcal{H}_3 \to \mathcal{H}_2$ ϕ^{Q_2}: $\mathcal{H}_2 \to \mathcal{H}_3$

$\implies \phi^{Q_1,X,Q_2}(x)$ cannot simply be the **composition** of $\phi^{Q_1,X}(x)$ and ϕ^{Q_2}!

Compositionality with Quantum Factors

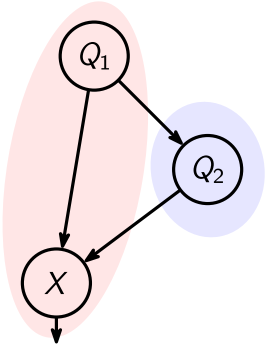

$$\phi^{Q_1} : \mathcal{H}_1 \otimes \mathcal{H}_2$$
$$\phi^{Q_2} : \mathcal{H}_2 \otimes \mathcal{H}_3$$
$$\phi^{X} : Val(X) \to \mathcal{H}_1 \otimes \mathcal{H}_3$$

Whole graph:

$$\phi^{Q_1,Q_2,X} : Val(X) \to \mathbb{C}$$
$$= \phi^{X} \odot \phi^{Q_2} \odot \phi^{Q_1}$$

$\phi^{Q_1,X} : Val(X) \to \mathcal{H}_3 \otimes \mathcal{H}_2$ $\quad$ $\phi^{Q_2} : \mathcal{H}_2 \otimes \mathcal{H}_3$

$\implies \phi^{Q_1,X,Q_2}(x) = \phi^{Q_1,X} \odot \phi^{Q_2}$

More generally, can compute for any order on the nodes!

Modularity

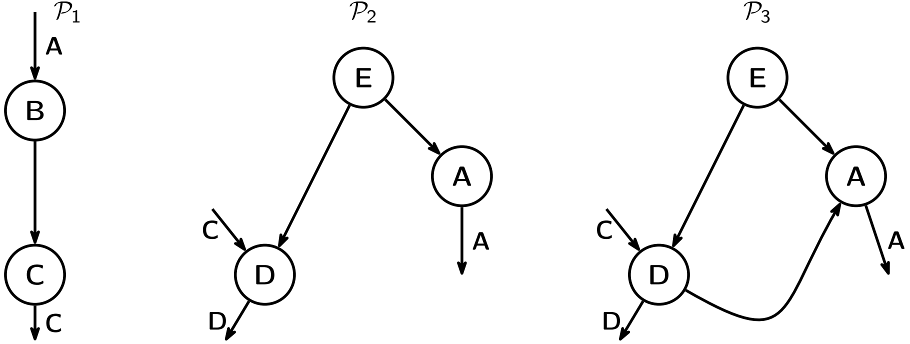

Observation: $\mathcal{P}_1 \cup \mathcal{P}_2$ is a QBN but not $\mathcal{P}_1 \cup \mathcal{P}_3$

How to ensure two parts always form a QBN?
We add a type = an interface
Could do so by patching QBN and getting yet another new syntax...
We prefer to use **proof-nets**, graphs from linear logic adapted to typing

Proof-Nets for Quantum Bayesian Networks

Definition: Proof-Net

A graph respecting some *graphical criterion* and built from:

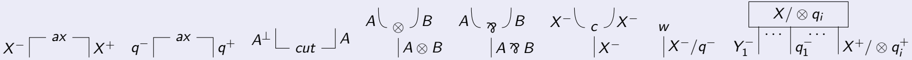

Example

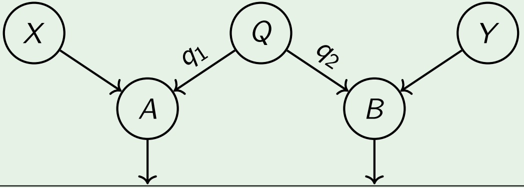

Definition: Proof-Net

A graph respecting some *graphical criterion* and built from:

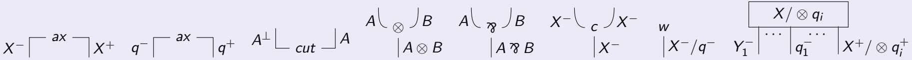

Example

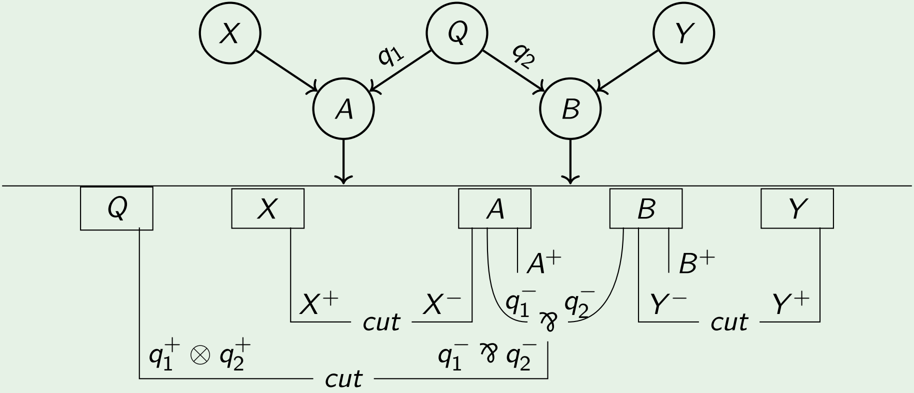

Modularity with types

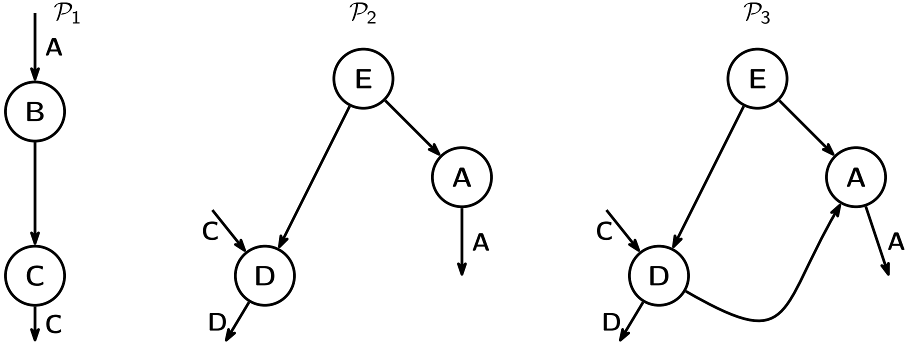

Observation: $\mathcal{P}_1 \cup \mathcal{P}_2$ is a QBN but not $\mathcal{P}_1 \cup \mathcal{P}_3$

Modularity with types

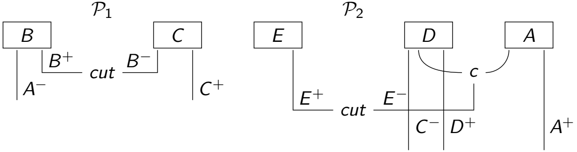

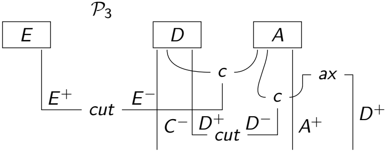

Observation: $\mathcal{P}_1 \cup \mathcal{P}_2$ is a QBN but not $\mathcal{P}_1 \cup \mathcal{P}_3$

Modularity with types

$\mathcal{P}_1$

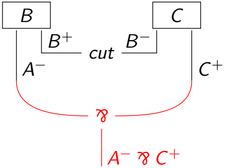

$\mathcal{P}_2$

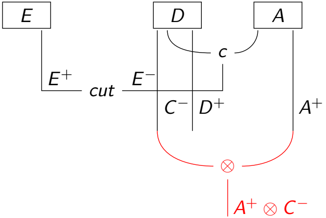

$\mathcal{P}_3$

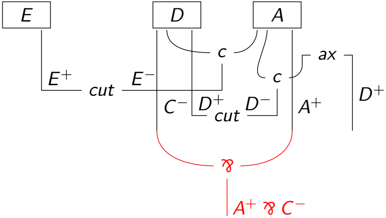

Observation: $\mathcal{P}_1 \cup \mathcal{P}_2$ is a QBN but not $\mathcal{P}_1 \cup \mathcal{P}_3$

Modularity with types

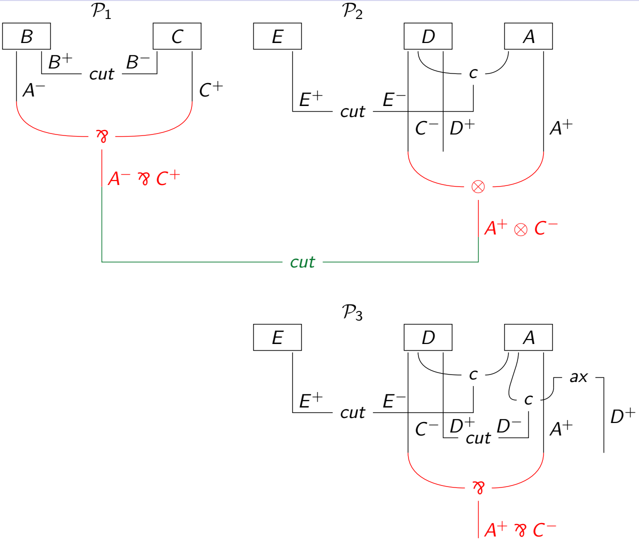

Observation: $\mathcal{P}_1 \cup \mathcal{P}_2$ is a QBN but not $\mathcal{P}_1 \cup \mathcal{P}_3$

Modularity with types

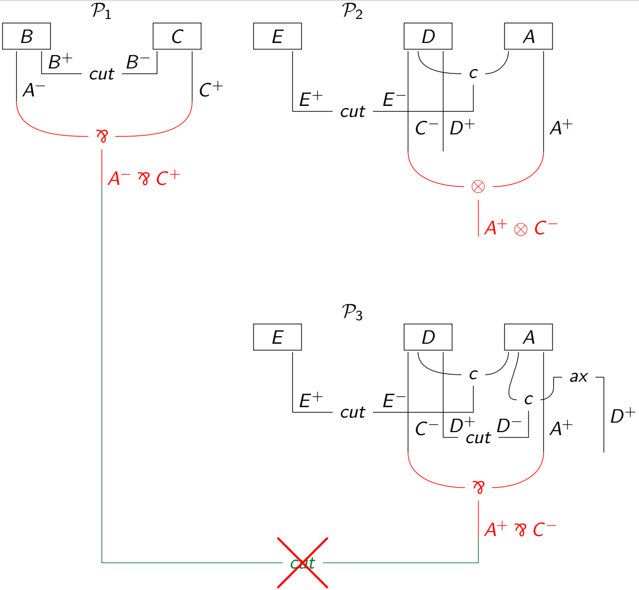

Observation: $\mathcal{P}_1 \cup \mathcal{P}_2$ is a QBN but not $\mathcal{P}_1 \cup \mathcal{P}_3$

Conclusion

Contributions

- Compositionality by modifying the semantic:
  from quantum instruments to (more general) *quantum factors*
- Conservative extension: our QBNs are exactly the usual Bayesian Networks in the absence of quantum nodes
- Modularity by typing in proof-nets:
  proof-theoretic approach adding an *interface*, parts with compatible interfaces are those giving a QBN

Conclusion

Contributions

- Compositionality by modifying the semantic:
  from quantum instruments to (more general) *quantum factors*
- Conservative extension: our QBNs are exactly the usual Bayesian Networks in the absence of quantum nodes
- Modularity by typing in proof-nets:
  proof-theoretic approach adding an *interface*, parts with compatible interfaces are those giving a QBN

Perspectives

- Using *terms* instead of proof-nets for typing
  $\longrightarrow$ work started with Claudia Faggian and Gabriele Tedeschi
- Compositionality can be used to study *conditional independence (no-signaling)* between random variables, even with quantum causes

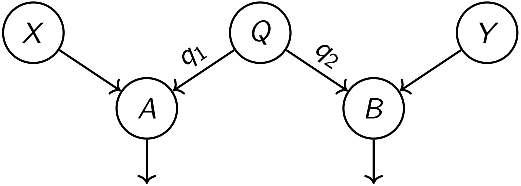

Thank you for your attention!

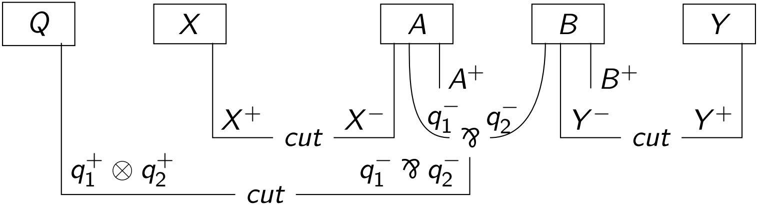

References I

[HLP14] Joe Henson, Raymond Lal, and Matthew F Pusey. "Theory-independent limits on correlations from generalized Bayesian networks". In: *New Journal of Physics* 16.11 (Nov. 2014). ISSN: 1367-2630. DOI: 10.1088/1367-2630/16/11/113043.