Quantum Bayesian Networks: Compositionality and Typing via Linear Logic

Rémi Di Guardia, Thomas Ehrhard, Claudia Faggian
Université Paris Cité, CNRS, Inria, IRIF, F-75013, Paris, France

Abstract

Quantum Bayesian networks [6] provide a mathematical formalism to describe causal relations, to analyse correlations, and to predict the probabilities of measurement outcomes, in systems involving both *classical and quantum* data. They generalize Pearl's Bayesian networks [9]—prominent graphical models for classical probabilistic reasoning and inference.

This extended abstract reports on work [3] which brings compositional principles and a typing discipline into this setting. A key feature of our compositional semantics is that when all causes are classical, it coincides with the standard factor-based semantics of Bayesian networks, while in the purely quantum case it reduces to tensor networks. We then propose a typed formalism based on proof-nets from linear logic, where types ensure well-behaved composition of systems, and which we prove sound and complete with respect to quantum Bayesian networks.

• Extended Abstract

This extended abstract reports on recent work on quantum Bayesian networks. Details of definitions and results can be found in the preprint [3] (to be published in FSCD 2026).

Pearl's *Bayesian networks* [8, 9] provide a framework for reasoning under conditions of uncertainty and partial knowledge, with a wide range of applications from statistics to epidemiology, economics and computer science. Bayesian networks have a *dual nature*, serving both as probabilistic models for inference, and as causal models precising the connections between observed data. When reasoning on *quantum systems*, this framework is not general enough to account for entanglement and non-local correlations. The development of quantum causal models—see *e.g.* [1] and references therein—is an active research area across quantum information and the foundations of quantum theory. Here, we focus on *quantum Bayesian networks*, a generalization of Bayesian networks introduced by Henson, Lal, and Pusey [6]. They provide a mathematical framework to describe causal relations and to predict the probabilities of measurement outcomes in systems with both classical and quantum data. The formalism builds on work by Leifer and Spekkens [7], with as perspective *quantum theory as a theory of inference*. Quantum theory is indeed probabilistic, as it is concerned with predicting the probabilities of measurement outcomes on a physical system. Hence, it can be framed as a *probabilistic inference over models with both classical and quantum data*.

▶ **Example 1.** The directed acyclic graph in Fig. 1 describes the well-known Bell experiment. Quentin prepares a pair of (possibly entangled) qubits, sending one to Alice and the other to Bob. Alice randomly performs one of two possible measurements, by flipping a coin X. Bob also performs a measurement, by flipping a coin Y. The result of the experiment is $\Pr(a, b \mid x, y) = \Pr(a, b, x, y)/\Pr(x, y)$ *i.e.* the probability that the (classical) outcomes of Alice's and Bob's measurements are a and b, given outcomes x for X and y for Y.

Bayesian Networks. In the theory of Bayesian networks, the causal structure is encoded by a *directed acyclic graph* (DAG) where nodes represent random variables and edges conditional dependencies. The semantics of a Bayesian network is a probability distribution. Bayesian networks provide a compact representation of large probability distributions, and efficient inference algorithms to answer queries about these distributions without constructing them in full. These rely on *factors*, a generalization of conditional probability distributions, whose operations *inherently share variables*, allowing for tractability and efficiency.

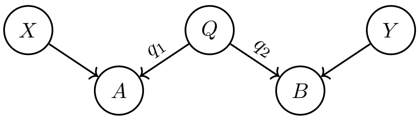

Figure 1 Bell set-up (from [6])

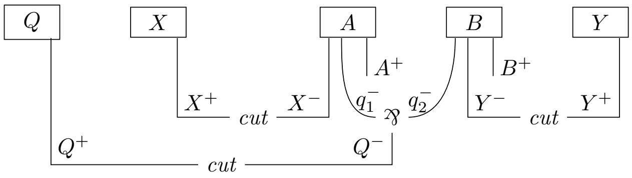

Figure 2 Bell set-up as a quantum proof-net

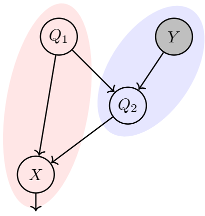

Figure 3 Compositionality

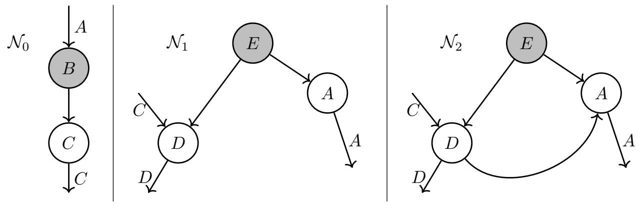

Figure 4 Modularity

Quantum Bayesian Networks. Quantum Bayesian networks are still an emerging field, not as developed as their classical counterparts. A crucial missing feature is *compositionality*, the ability to compute the semantics through intermediate, partial computations—*i.e.* computing a model's denotation as a function of its subparts. Consider the DAG in Fig. 3: nodes X and Y produce a classical output, while nodes Q_1 and Q_2 have a quantum nature. A natural question is: *Can we compute the semantics of the model in terms of sub-components—e.g.* the highlighted sub-graphs? The approach in [6] does not adapt well (see [6, page 12]).

A closely related question concerns *modularity*: *When can causal descriptions of systems as subparts be used to construct larger models?* Consider the three DAGs in Fig. 4: $\mathcal{N}_0$ awaits an input A and outputs C, $\mathcal{N}_1$ and $\mathcal{N}_2$ both await an input C and output A and D. The graph obtained by plugging together $\mathcal{N}_0$ and $\mathcal{N}_1$ (matching inputs and outputs) is a DAG, while the graph that plugs together $\mathcal{N}_0$ and $\mathcal{N}_2$ has a directed cycle.

Contributions & Challenges. We address this lack of compositionality and modularity by introducing methods and concepts from denotational semantics and proof theory. We do so in a way that is fully compatible with Bayesian networks and Bayesian inference.

Our first contribution is a *compositional semantics*, allowing the interpretation and modular combination of components while still being equivalent to [6]. We adapt Selinger's semantics in [10] to the factors of Bayesian networks. The technical challenge is to conciliate two very different behaviors: classical variables share their values, and they do so in an efficient way, whereas quantum data cannot be shared. We satisfy both requirements by introducing *quantum factors*. Remarkably, when all causes are classical they coincide with the standard factors, while in the purely quantum case they behave like tensor networks.

Our second contribution is a *typed graphical formalism*. Rather than defining yet-another-syntax, we encode quantum Bayesian networks into the graph syntax of linear logic: *proof-nets*, see Fig. 2. These typed graphs allow to guarantee that composing graphs of compatible types produces a DAG, *i.e.* types ensure well-behaved compositions of systems. This formalism is *sound and complete w.r.t.* quantum Bayesian networks: every quantum Bayesian network can be represented as a proof-net, and every closed proof-net corresponds to a quantum Bayesian network. This builds upon a recent line of work connecting (classical) Bayesian networks with proof-nets [4, 5, 2]. The DAGs in Fig. 4 admit the following typing: $\mathcal{N}_0 \vdash A \multimap C$, $\mathcal{N}_1 \vdash (A \multimap C) \multimap D$, $\mathcal{N}_2 \vdash C \multimap (A \otimes D)$. The DAGs $\mathcal{N}_0$ and $\mathcal{N}_1$ compose together, producing a DAG of output D, while $\mathcal{N}_2$ cannot be given any type matching the one of $\mathcal{N}_0$.

References

1 Jonathan Barrett, Robin Lorenz, and Ognyan Oreshkov. Quantum causal models, 2019. URL: `https://arxiv.org/abs/1906.10726`.

2 Rémi Di Guardia, Thomas Ehrhard, Jérôme Evrard, and Claudia Faggian. Bayesian networks and proof-nets: the proof-theory of Bayesian inference, 2026. URL: `https://arxiv.org/abs/2602.04045`.

3 Rémi Di Guardia, Thomas Ehrhard, and Claudia Faggian. Quantum Bayesian networks: Compositionality and typing via linear logic, 2026. URL: `https://arxiv.org/abs/2604.26059`.

4 Thomas Ehrhard, Claudia Faggian, and Michele Pagani. The sum-product algorithm for quantitative multiplicative linear logic. In *8th International Conference on Formal Structures for Computation and Deduction, FSCD 2023*, volume 260 of *LIPIcs*, pages 8:1–8:18, 2023. `doi:10.4230/LIPICS.FSCD.2023.8`.

5 Claudia Faggian, Daniele Pautasso, and Gabriele Vanoni. Higher order Bayesian networks, exactly. In *Proceedings of the 51st Annual Symposium on Principles of Programming Languages (POPL)*, volume 8, pages 2514–2546. ACM, January 2024. `doi:10.1145/3632926`.

6 Joe Henson, Raymond Lal, and Matthew F. Pusey. Theory-independent limits on correlations from generalised Bayesian networks. *New Journal of Physics*, 16(11):113043, 2014. URL: `https://arxiv.org/abs/1405.2572`, `doi:10.1088/1367-2630/16/11/113043`.

7 Matthew S. Leifer and Robert W. Spekkens. Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference. *Physical Review A*, 88(5):052130, Nov 2013. arXiv:1107.5849. `doi:10.1103/PhysRevA.88.052130`.

8 Judea Pearl. Fusion, propagation, and structuring in belief networks. *Artificial Intelligence*, 29(3):241–288, 1986. `doi:10.1016/0004-3702(86)90072-X`.

9 Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, UK, 2nd edition, 2009.

10 Peter Selinger. Toward a semantics for higher-order quantum computation. In *Proceedings of the 2nd International Workshop on Quantum Programming Languages*, volume 33, pages 127–143. TUCS General Publication, 2004.