

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 [10]—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 linear logic proof-nets, where types ensure well-behaved composition of systems.

2012 ACM Subject Classification Theory of computation

Keywords and phrases Quantum Causal Models, Bayesian Networks, Proof-Nets, Linear Logic

1 Introduction


This extended abstract reports on recent work on Quantum Bayesian networks. Details of definitions and results can be found in the online preprint [3].

Pearl’s Bayesian networks [9, 10] 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 graphical models for classical probabilistic reasoning and inference, and as causal models, precisising the connections between observed data and causal relations.

When reasoning on quantum systems, the classical framework is not general enough to account for entanglement and the non-local correlations observed in the Bell experiments. 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, advancing along various axes. Its motivations span from foundational questions and non-locality, to enabling device-independent cryptographic protocols, and to facilitating data-driven discovery.

We focus on Quantum Bayesian networks, a direct generalization of Pearl’s networks introduced in foundational work by Henson, Lal, and Pusey [6]. They provide a mathematical framework to describe causal relations, to analyse correlations, and to predict the probabilities of measurement outcomes, in systems involving both classical and quantum data. The formalism builds on previous work by Leifer and Spekkens [7], where the perspective is that of *quantum theory as a theory of inference*. Quantum theory is indeed fundamentally probabilistic at its core, as it is concerned with predicting the probabilities of measurement outcomes on a physical system. In this sense, the prediction task can be framed as a problem of *probabilistic inference over models involving both classical and quantum data*, as Example 1 illustrates. Probabilistic inference then offers mathematical and logical tools to comprehend, predict, and control quantum phenomena, essential to both the theoretical understanding and to quantum technologies.

► **Example 1 (Alice & Bob: the Bell set-up).** The directed acyclic graph in Fig. 1 describes the well-known set-up for the Bell experiment. Alice and Bob—who stand in widely separated laboratories—are each able to perform two possible measurements on a qubit (for example, measuring it with respect to two different bases). Their colleague Quentin prepares a pair of qubits,



■ **Figure 1** Bell set-up (from [6])

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

sending one to Alice and the other to Bob. When Alice receives her qubit q_1 , she chooses to randomly perform one of the two possible measurements, by flipping a coin X . When Bob receives his qubit q_2 , he also performs a measurement, by flipping a coin Y . The result of the experiment is

$$\Pr(a, b \mid x, y) = \Pr(a, b, x, y) / \Pr(x, y) \quad (1)$$

i.e. the probability that the (classical) outcomes of Alice and Bob measurements are respectively a and b , given outcomes x for X and y for Y .

Bayesian Networks. In the framework of Bayesian networks, the causal structure is encoded by a *directed acyclic graph* (DAG), where nodes represent random variables and edges express conditional dependencies. The strength of the dependencies (or the degree of knowledge) is quantified by conditional probability tables. The structure enables a compact representation of large probability distributions, and efficient inference of marginals or conditionals through factorization. The strength of Bayesian networks is to provide efficient inference algorithms (both exact and approximate) that can answer queries about the underlying distribution without explicitly constructing it in full. Critically, the DAG includes both *observable* variables of interest and *hidden* (unobserved) ones; we adopt the common convention of denoting classical hidden variables by shaded nodes. In the quantum setting—as in the Bell’s experiments—hidden variables play a central role, and they may correspond to quantum systems.

Semantics and Inference. The semantics of a Bayesian network is the probability distribution it defines. More accurately, what one seeks is the *marginal distribution over variables of interest*. Exact inference computes it precisely. This involves two key operations:

product (*composing*) + summing-out (*hiding*) irrelevant variables.

The formalization and theory of inference rely on a class of functions, known as *factors*, which is an abstraction of conditional probability distributions. Tractability and efficiency rely on their properties, and specifically on two key aspects: 1. the product of factors *inherently shares variables*, and 2. product and sum *distribute* under suitable conditions—pushing the sum on smaller components reduces the size of computations.

Quantum Bayesian Networks, issues. Quantum Bayesian networks are still an emerging field, not as developed as their classical counterparts. A crucial missing feature is the ability to compute the semantics of the model (the desired marginal distribution) through intermediate, partial computations, without ever computing the full joint distribution. Put differently, what is lacking is *compositionality*, which would enable computing a model’s denotation as a function of its subparts, alongside modular reasoning. A closely related question concerns *modularity*: when can causal descriptions of systems as subparts be used to construct larger models? A well-established tool to ensure modularity are *types*, that specify precise contracts (*e.g.* input/output behaviors) for the components of a system.

Our goal is twofold.

- To address the lack of compositionality and modularity in the setting of Quantum Bayesian networks by introducing methods and concepts from denotational semantics and proof theory, thereby enabling compositional principles and a typing discipline.

- To develop a framework that is fully compatible with Bayesian networks and Bayesian inference, thus paving the way for the application of the techniques developed in that context.

Compositionality and modularity facilitate reasoning about complex systems and their properties, ensuring that the meaning of the system can be derived systematically and in a principled way from the meanings of its parts, and allowing components to be analyzed, replaced, and reused independently. Compositionality also enables modular reasoning about inference.

Types serve as abstract characterizations of systems behaviors, constraining and guiding their formation. Types discipline statically guarantees fundamental semantic properties such as termination, consistency, and compositional correctness. Well-typed programs exhibit well-behaved execution and preserve probabilistic and quantitative invariants throughout evaluation.

Encompassing the Bayesian networks semantics. As seen in Example 1 (the Bell set-up), even for systems involving quantum sources of causality, only the classical outcomes of measurements can be observed. Thus, what a model ultimately defines is a probability distribution over classical variables—such as Eq. (1). From the inference perspective, it is desirable to have a framework enabling also (when relevant) the extensive set of inference techniques and algorithms developed in the Bayesian networks domain. To make this possible, our semantics integrates and generalizes a notion which is key to exact inference algorithms—that of *factor*.

Contributions and challenges. Our first contribution is to develop a *compositional semantics*, which allows for the interpretation and modular combination of components. We adapt Selinger’s semantics in [11] to take into account the factor-based approach of Bayesian networks. The main technical challenge is to conciliate two very different behaviors:

- *Classical variables share their values*, and they do so in an efficient way: the mathematical setting underlying Bayesian networks integrates this feature in the very definition of the product of factors, and exploits it to obtain compact representations and efficient calculations.
- *Quantum data cannot be shared*: a defining feature of quantum computing is the *No-Cloning Theorem*, implying that qubits cannot be duplicated nor broadcast to multiple receivers.


We satisfy both requirements by introducing *quantum factors* whose mathematical development is our main and most technical result. Remarkably, when all causes are classical, our framework exactly coincides with the standard *factors-based* semantics of Bayesian networks, while in the purely quantum case, it behaves like tensor networks. We rely on quantum factors to reformulate the definition of quantum Bayesian networks. Our factor-based formalism is equivalent to that of [6], however our semantics enables a compositional interpretation by sub-components, unlike the original definition (see [6, page 12]). A crucial aspect to achieve compositionality is that quantum factors are closed under both product and sum-out (*i.e.*, marginalization of unobserved variables). Finally, we explore a *typed graphical formalism* based on linear logic proof-nets, where types ensure well-behaved compositions of systems.

Motivational examples: compositionality and modularity


Two examples can illustrate the desiderata and issues with compositionality and modularity.

Compositionality. Consider the DAG in Fig. 3. The nodes X, Y produce a classical output, while Q_1, Q_2 have a quantum nature. *Can we compute the semantics of the model in terms of sub-components*—as for example the highlighted sub-graphs? The approach in [6] does not adapt well to an interpretation by component, because it has a global nature (more details in [3]).

Do parts compose well? When considering components (possibly with open inputs), a natural question somehow dual to the previous one is whenever independently defined components compose well. Consider the three valid DAGs in Fig. 4. The leftmost, \mathcal{N}_0 , waits for an input A



■ **Figure 3** Compositionality



■ **Figure 4** Modularity

and outputs C . The others, \mathcal{N}_1 and \mathcal{N}_2 , wait for the same input C and both output A and D . One easily checks that the graph obtained by plugging together \mathcal{N}_0 and \mathcal{N}_1 (matching inputs and outputs) is a DAG, whereas the graph which plugs together \mathcal{N}_0 and \mathcal{N}_2 has a directed cycle. By moving to *typed graphs*, we guarantee that composing graphs of compatible types produces a DAG. Rather than defining yet-another-syntax, we encode quantum Bayesian networks into the graph syntax of linear logic: proof-nets. We do so by building on a recent line of work connecting (classical) Bayesian networks with proof-nets [4, 5, 2]. Each proof-net is typed by a sequent in multiplicative linear logic. As we will see (Fig. 6), the DAGs in Fig. 4 admit the following typing:

$$\mathcal{N}_0 \vdash A \multimap C, \quad \mathcal{N}_1 \vdash (A \multimap C) \multimap D, \quad \mathcal{N}_2 \vdash (A \otimes C^-) \multimap D$$

In particular, \mathcal{N}_0 and \mathcal{N}_1 compose together, producing a DAG of output D , while \mathcal{N}_2 cannot be given any type that matches the one of \mathcal{N}_0 .

2 Quantum Bayesian Networks

A quantum Bayesian network is a DAG with nodes a *set of r.v.s* \mathbf{X} and a *set of quantum systems* \mathbf{Q} . We will associate to each node a *function*—called quantum factor—from $\text{Val}(\mathbf{X})$ to *positive operators* [8] on the relevant Hilbert space. We will then interpret in the same way quantum Bayesian networks, as well as any sub-component.

We assume given a countable set \mathbb{Q} of indexed qubits q_0, q_1 , etc. with associated 2-dimensional Hilbert spaces $\mathcal{H}_{q_0}, \mathcal{H}_{q_1}$, etc. We call **quantum register** a *finite* set of qubits $Q \subseteq \mathbb{Q}$, whose associated Hilbert space is $\mathcal{H}_Q \stackrel{\text{def}}{=} \bigotimes_{q \in Q} \mathcal{H}_q$.

► **Definition 2** (Quantum Factor (Q-factor)). *Let \mathbf{X} be a set of random variables and Q a quantum register. A **quantum factor** (shortened as **Q-factor**) ϕ over (\mathbf{X}, Q) is a function*

$$\phi : \text{Val}(\mathbf{X}) \rightarrow \mathcal{L}(\mathcal{H}_Q)^+ \quad (2)$$

*mapping each tuple of values \mathbf{x} to a **positive** operator from \mathcal{H}_Q to \mathcal{H}_Q .*

We show in [3] that Q-factors admit and *are closed* under both *product* \odot and *summing-out* \sum operations. Remarkably, Q-factors encompasses both factors from Bayesian networks theory and tensor networks:

- For Q-factors ϕ and ϕ' over classical variables only, $\phi \odot \phi'$ is the product of factors and $\sum_X \phi$ is the sum-out of factors.
- For Q-factors ϕ and ϕ' over quantum registers only, $\phi \odot \phi'$ is the product of tensor networks.

Quantum Bayesian Networks. We formulate quantum Bayesian networks as DAGs with a Q-factor for each node. These Q-factors must respect some technical condition, indicating each of them corresponds to a quantum instrument [12].

► **Definition 3** (Quantum Bayesian Network (QBN)). *A **quantum Bayesian network (QBN)** over a set of r.v.s \mathbf{X} and a set of disjoint quantum registers \mathbf{Q} is a pair (\mathcal{G}, Φ) where:*

- \mathcal{G} is a directed acyclic graph over the set of nodes $\mathbf{X} \cup \mathbf{Q}$;
- each out-edge of a node $Q \in \mathbf{Q}$ is labelled by a non-empty register Q_i , with $\uplus_i Q_i = Q$;

References

- 1 Jonathan Barrett, Robin Lorenz, and Ognyan Oreshkov. Quantum causal models. Preprint, June 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. Preprint, 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. Preprint, 2026. URL: <https://hal.science/hal-05549936>.
- 4 Thomas Ehrhard, Claudia Faggian, and Michele Pagani. The sum-product algorithm for quantitative multiplicative linear logic. In Marco Gaboardi and Femke van Raamsdonk, editors, *8th International Conference on Formal Structures for Computation and Deduction, FSCD 2023, July 3-6, 2023, Rome, Italy*, volume 260 of *LIPICs*, pages 8:1–8:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPICs.FSCD.2023.8.
- 5 Claudia Faggian, Daniele Pautasso, and Gabriele Vanoni. Higher order bayesian networks, exactly. *Proc. ACM Program. Lang.*, 8(POPL):2514–2546, 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 Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, 10th anniversary edition edition, 2010.
- 9 Judea Pearl. Fusion, propagation, and structuring in belief networks. *Artif. Intell.*, 29(3):241–288, 1986. doi:10.1016/0004-3702(86)90072-X.
- 10 Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, UK, 2nd edition, 2009.
- 11 Peter Selinger. Toward a semantics for higher-order quantum computation. In Peter Selinger, editor, *Proceedings of the 2nd International Workshop on Quantum Programming Languages*, volume 33, pages 127–143. TUCS General Publication, June 2004.
- 12 John Watrous. *The Theory of Quantum Information*. Cambridge University Press, Cambridge, 2018.