

General Bayesian theories and the emergence of the exclusivity principle

Giulio Chiribella ^{1,2,3,*} Adán Cabello ^{4,5,†} Matthias Kleinmann ^{6,7,‡} and Markus P. Müller ^{8,3,§}

¹*QICI Quantum Information and Computation Initiative, Department of Computer Science, The University of Hong Kong, Pokfulam Road, Hong Kong*

²*Department of Computer Science, University of Oxford, Parks Road, Oxford OX1 3QD, United Kingdom*

³*Perimeter Institute for Theoretical Physics, Waterloo, Ontario, Canada N2L 2Y5*

⁴*Departamento de Física Aplicada II, Universidad de Sevilla, E-41012 Sevilla, Spain*

⁵*Instituto Carlos I de Física Teórica y Computacional, Universidad de Sevilla, E-41012 Sevilla, Spain*

⁶*Department of Theoretical Physics, University of the Basque Country UPV/EHU, P.O. Box 644, E-48080 Bilbao, Spain*

⁷*Naturwissenschaftlich–Technische Fakultät, Universität Siegen, Walter-Flex-Straße 3, D-57068 Siegen, Germany*

⁸*Institute for Quantum Optics and Quantum Information, Austrian Academy of Sciences, Boltzmanngasse 3, A-1090 Vienna, Austria*



(Received 22 April 2019; revised 8 March 2020; accepted 2 September 2020; published 5 October 2020)

We address the problem of reconstructing quantum theory from the perspective of an agent who makes bets about the outcomes of possible experiments. We build a general Bayesian framework that can be used to organize the agent's beliefs and update them when new information becomes available. Our framework includes as special cases classical and quantum probability theory, as well as other forms of probabilistic reasoning that may arise in future physical theories. Building on this framework, we develop a notion of an ideal experiment, which in quantum theory coincides with the notion of projective measurement. We then prove that, in every general Bayesian theory, ideal experiments must satisfy the exclusivity principle, a property of projective measurements that plays a central role in the characterization of quantum correlations. Our result suggests that the set of quantum correlations may be completely characterized in terms of Bayesian consistency conditions.

DOI: [10.1103/PhysRevResearch.2.042001](https://doi.org/10.1103/PhysRevResearch.2.042001)

Introduction. Quantum theory portrays a world where the outcomes of individual measurements cannot be predicted with certainty, and yet the quantum predictions are strikingly accurate and explain an astonishingly broad range of phenomena. The reason for this broad applicability still remains controversial. Does the quantum formalism describe how Nature works at the fundamental level? Or it is just a mathematical tool for guessing the outcomes of our experiments?

Albeit with a variety of nuances, different interpretations of quantum theory tend to favor either one or the other view. For example, Everett's interpretation [1] holds that the quantum framework refers to a multitude of universes interfering with each other. On the other hand of the spectrum, QBism, which originally stood for quantum Bayesianism [2], views quantum theory as a set of rules that constrain how agents should make bets about the outcomes of their experiments.

Different interpretations are reflected into different ways to understand quantum correlations. Since Bell [3], it has been known that quantum correlations are incompatible with the

intuitive worldview known as local realism. But intriguingly, the quantum violations of Bell's inequalities are not maximal: More general theories compatible with relativistic causality could in principle lead to larger violations [4–6]. Following up on this observation, various physical principles have been proposed to explain the quantum bounds on correlations [7–11]. Behind this approach lies the idea that the quantum bounds should be explained in terms of principles constraining how Nature behaves. However, this is not the only option. Instead, one could search for principles constraining how agents assign probabilities to the outcomes of their experiments. This approach has remained mostly unexplored so far, partly due to the lack of a suitable framework, and partly due to a widespread belief that quantum correlations require new physical principles. Even within QBism, the Born rule is regarded as “an empirical *addition* to the laws of Bayesian probability” [2], rather than a consequence of Bayesian probability itself.

In this Rapid Communication we demonstrate that a surprisingly large portion of the set of quantum correlations follows directly from elementary Bayesian conditions. We first build a general Bayesian framework, describing the ways in which an agent can update its beliefs when a new piece of information becomes available. The framework assumes only basic Bayesian laws, such as the validity of the rule of conditional probability, and the consistency of beliefs at different moments of time. Surprisingly, we find that these rather minimalistic assumptions imply the validity of the exclusivity principle, a feature of quantum theory that characterizes a large portion of the set of quantum correlations [12–19].

*Corresponding author: giulio@cs.hku.hk

†adan@us.es

‡matthias.kleinmann@uni-siegen.de

§markus.mueller@oeaw.ac.at

The exact statement of the exclusivity principle will be given later in this Rapid Communication. For the moment, the crucial observation is that the exclusivity principle *does not hold for arbitrary experiments*. In quantum theory, it holds for projective measurements, but it fails to hold for certain nonprojective measurements [17,20]. This means that, in order to formulate the exclusivity principle in a general physical theory, one has first to extend the notion of projective measurement beyond the quantum framework. Such an extension is far from straightforward, because there are multiple inequivalent notions that generalize the notion of projective measurement in quantum theory [20]. On top of that, once a choice is made, the exclusivity principle may or may not hold, depending on the theory under consideration. For example, Ref. [12] showed that, for a certain generalization of the notion of projective measurement, the exclusivity principle is violated by theories that predict superstrong correlations such as Popescu-Rohrlich boxes [5].

Theories that satisfy the exclusivity principle have a remarkable property: Under the natural assumption that two statistically independent experiments can be performed in parallel, the correlations arising in every Bell and Kochen-Specker contextuality scenario must be contained in the quantum set [19]. This means that, under a mild assumption, the quantum set for these contextuality scenarios can be characterized as the largest set of correlations compatible with the exclusivity principle.

In this Rapid Communication we derive the exclusivity principle from Bayesian consistency conditions. We first introduce a general framework for Bayesian theories. We then formulate a general notion of an ideal experiment, which in the special case of quantum theory coincides with the notion of projective measurement. Our first result is that ideal experiments exist in every general Bayesian theory (Theorem 1). Our second result is that the outcome probabilities arising from ideal experiments satisfy the exclusivity principle (Theorem 2). Complemented with the results of Ref. [19], our results imply that the quantum set of correlations for any Bell or Kochen-Specker contextuality scenario can be derived just from elementary consistency conditions on the agent's probability assignment. This conclusion surpasses the expectations of QBism: Rather than being an empirical addition [2], the Born rule appears as a consequence of Bayesian consistency conditions alone.

Beliefs and probabilities. Consider the situation of an agent who makes bets about the outcomes of experiments on a given physical system. We start from a basic class of experiments, which we call *principal experiments*. The outcomes of the principal experiments form a single sample space X , equipped with a σ -algebra of events Σ , namely a collection of subsets of X satisfying the properties (i) $X \in \Sigma$, (ii) $E \in \Sigma$ implies $(X \setminus E) \in \Sigma$, and (iii) $(E_i)_i \subset \Sigma$ implies $\bigcup_i E_i \in \Sigma$. In typical cases, X is a finite set and Σ is the power set of X . A principal experiment corresponds to a partition $\mathbb{E} = (E_i)_i$ of the sample space X into disjoint events. For brevity, we will identify the experiment and the corresponding partition.

In making a bet, the agent will rely on its beliefs, including beliefs on the laws of physics, or beliefs on the prior history of the physical system involved in the bet. We denote by \mathbf{B} the set of all possible beliefs. For a given belief β ,

the agent assigns a probability distribution $p : E \mapsto p(E|\beta)$ satisfying the usual conditions (i) $p(E|\beta) \geq 0$ for all events E , (ii) $p(\bigcup_i E_i|\beta) = \sum_i p(E_i|\beta)$ whenever all E_i are mutually disjoint, and (iii) $p(X|\beta) = 1$. Here, we make the standard assumption that the probability of an event E is independent of the specific experiment \mathbb{E} in which E arises.

We stress that the belief β determines the probability assignment $p(E|\beta)$, but not vice versa: In general, a belief contains much more than just the outcome probabilities of principal experiments. For example, we will see that in quantum theory a belief is described by a density matrix, while the principal experiments are commuting projective measurements, corresponding to projectors that are diagonal with respect to a fixed standard basis. The probability assignment $p(E|\beta)$ depends only on the diagonal entries of the density matrix β , and therefore it is not sufficient to determine the whole matrix.

Bayesian updates. Suppose that the agent receives a guarantee that an event E is the case. As a consequence, the agent will update its old belief β into a new belief, which we denote as $\beta' = E\beta$. Again, we make the standard assumption that the new belief is independent of the specific experiment in which E arises.

The point of updating beliefs is to compute conditional probability distributions. We demand that the probability assignment for the updated belief $E\beta$ is given by the rule of conditional probabilities:

Axiom 1: Rule of conditional probabilities. For every initial belief β , and for every pair of events E and F with $p(E|\beta) \neq 0$, the updated belief $E\beta$ satisfies the rule of conditional probabilities

$$p(F|E\beta) = \frac{p(E \cap F|\beta)}{p(E|\beta)}. \quad (1)$$

The rule of conditional probabilities implies several properties of the update map $\beta \mapsto E\beta$. For example, it implies that, once the agent updates its belief based on the event E , the agent becomes certain of the event E . Indeed, one has $p(E|E\beta) = 1$, which follows from letting $E = F$ in Eq. (1).

We stress that the update map does not represent a physical process on the observed system, but rather an operation internal to the agent. In the following, we formulate two conditions that the update should satisfy in order for the agent to be consistent with its beliefs.

Forward consistency. Suppose that the agent is certain of the event E , namely $p(E|\beta) = 1$. In this case, a guarantee that E is the case does not add any new information, and should not lead to any update:

Axiom 2: Forward consistency. If the agent is certain of the event E , then the agent's belief does not change under the update for event E . Mathematically: For every $\beta \in \mathbf{B}$ and every $E \in \Sigma$, $p(E|\beta) = 1$ implies $E\beta = \beta$.

Forward consistency constrains the agent in how it updates the belief forward in time. A simple consequence of this constraint is that the total event $E = X$ does not lead to any update, namely $X\beta = \beta$, $\forall \beta \in \mathbf{B}$.

Actions. So far, we considered the situation where an agent bets directly on the occurrence of a certain event. More generally, the conditions under which a bet is made can be altered by some action. We use "action" broadly, including situations

in which the agent lets the system evolve under its natural dynamics.

We denote the set of all actions as Act , and we assume that it is a monoid, meaning that (i) actions can be composed with one another, (ii) the composition is associative, and (iii) there exists an identity action [21,22]. When an action \mathcal{A} is performed, the agent generally changes its belief to a new belief $\beta' = \mathcal{A}\beta$. We assume that the belief change satisfies the conditions (i) $(\mathcal{A}\mathcal{B})\beta = \mathcal{A}(\mathcal{B}\beta), \forall \beta \in \mathbf{B}, \forall \mathcal{A}, \mathcal{B} \in \text{Act}$, and (ii) $\mathcal{I}\beta = \beta, \forall \beta \in \mathbf{B}$, where \mathcal{I} is the identity action.

Backward consistency. When actions are included in the picture, a new class of *sequential experiments* arise. A sequential experiment consists of a sequence of actions interspersed by principal experiments. For example, $(\mathcal{A}, \mathbb{E}, \mathcal{B}, \mathbb{F})$ represents a sequence consisting of an action \mathcal{A} , followed by a principal experiment \mathbb{E} , by another action \mathcal{B} , and by another principal experiment \mathbb{F} . Crucially, the belief β supplies the agent with a joint probability distribution $P_{\beta, \mathcal{A}, \mathbb{E}, \mathcal{B}, \mathbb{F}}(E, F) := p(F|\mathcal{B}\mathbb{E}\mathcal{A}\beta) p(\mathbb{E}|\mathcal{A}\beta)$ if $p(\mathbb{E}|\mathcal{A}\beta) \neq 0$, and zero otherwise.

When the marginal probability $P_{\beta, \mathcal{A}, \mathbb{E}, \mathcal{B}, \mathbb{F}}(F) := \sum_{E \in \mathbb{E}} P_{\beta, \mathcal{A}, \mathbb{E}, \mathcal{B}, \mathbb{F}}(E, F)$ is nonzero, one can define the conditional probability $P_{\beta, \mathcal{A}, \mathbb{E}, \mathcal{B}, \mathbb{F}}(E|F) := P_{\beta, \mathcal{A}, \mathbb{E}, \mathcal{B}, \mathbb{F}}(E, F)/P_{\beta, \mathcal{A}, \mathbb{E}, \mathcal{B}, \mathbb{F}}(F)$, which quantifies the agent's confidence in *retrodicting* that the earlier outcome must have been E , given that the later outcome is F . We say that the event F is *unaffected by the experiment* \mathbb{E} if the probability of F in the sequential experiment $(\mathcal{A}, \mathbb{E}, \mathcal{B}, \mathbb{F})$ coincides with the probability of F in the sequential experiment $(\mathcal{A}, \mathcal{B}, \mathbb{F})$, where the principal experiment \mathbb{E} has been removed. In formula, $P_{\beta, \mathcal{A}, \mathbb{E}, \mathcal{B}, \mathbb{F}}(F) = P_{\beta, \mathcal{A}, \mathcal{B}, \mathbb{F}}(F) := p(F|\mathcal{B}\mathcal{A}\beta)$.

If $P_{\beta, \mathcal{A}, \mathbb{E}, \mathcal{B}, \mathbb{F}}(E|F) = 1$ and F is unaffected by the experiment \mathbb{E} , we say that the event F *implies* the event E . In this case, it is natural to require that the update for event E is already included in the update for event F . This idea motivates the following axiom:

Axiom 3: Backward consistency. If the event $F \in \mathbb{F}$ implies the event $E \in \mathbb{E}$ in a sequential experiment $(\mathcal{A}, \mathbb{E}, \mathcal{B}, \mathbb{F})$, then the update for event E can be omitted in the final belief updated for event F . Mathematically: If $P_{\beta, \mathcal{A}, \mathbb{E}, \mathcal{B}, \mathbb{F}}(E|F) = 1$ and $P_{\beta, \mathcal{A}, \mathbb{E}, \mathcal{B}, \mathbb{F}}(F) = P_{\beta, \mathcal{A}, \mathcal{B}, \mathbb{F}}(F)$, then $F\mathcal{B}\mathbb{E}\mathcal{A}\beta = F\mathcal{B}\mathcal{A}\beta$.

Axioms 1–3 define a set of theories, which we call *general Bayesian theories* (GBTs). In the Supplemental Material [23] we show that quantum theory is an example of a GBT, with the principal experiments corresponding to projective measurements on a fixed standard basis, and the updates following Lüders' rule.

Ideal experiments. We now show that every GBT contains a special class of ideal experiments that leave the agent with the option of gathering more refined information in the future.

We say that an experiment $(\mathcal{B}, \mathbb{F})$ with partition $\mathbb{F} = (F_{i,l})_{i,l}$ is a *refinement* of another experiment $(\mathcal{A}, \mathbb{E})$ with partition $\mathbb{E} = (E_i)_i$ if

$$\sum_l p(F_{i,l}|\mathcal{B}\beta) = p(E_i|\mathcal{A}\beta) \quad \forall i, \forall \beta \in \mathbf{B}. \quad (2)$$

The experiment $(\mathcal{A}, \mathbb{E})$ is *sequentially refinable* if there exists an action \mathcal{A}' such that, for every refinement $(\mathcal{B}, \mathbb{F})$ and for every initial belief β , the probability of the event $F_{i,l}$ in the experiment $(\mathcal{B}, \mathbb{F})$ is equal to the joint probability of the events $(E_i, F_{i,l})$ in the sequential experiment $(\mathcal{A}, \mathbb{E}, \mathcal{A}', \mathcal{B}, \mathbb{F})$.




FIG. 1. Exclusivity principle in an example. Three experiments have possible outcomes $\mathcal{O}_1 = \{a, d, f\}$, $\mathcal{O}_2 = \{d, b, e\}$, and $\mathcal{O}_3 = \{a, b, c\}$, respectively. Each outcome x has a probability $p(x)$ assigned to it, and the probabilities of the outcomes within each experiment sum up to 1, corresponding to the normalization condition $p(a) + p(d) + p(f) = p(d) + p(b) + p(e) = p(a) + p(b) + p(c) = 1$. The outcomes $\{a, b, d\}$ satisfy the conditions $\{a, b\} \subset \mathcal{O}_3$, $\{a, d\} \subset \mathcal{O}_1$, and $\{b, d\} \subset \mathcal{O}_2$, and therefore are called pairwise exclusive. The exclusivity principle demands $p(a) + p(b) + p(d) \leq 1$.

In formula,

$$p(F_{i,l}|\mathcal{B}\beta) = p(F_{i,l}|\mathcal{B}\mathcal{A}'E_i\mathcal{A}\beta) p(E_i|\mathcal{A}\beta) \quad (3)$$

for every refinement $(\mathcal{B}, \mathbb{F})$, for every event $F_{i,l} \in \mathbb{F}$, and for every belief $\beta \in \mathbf{B}$ with $p(E_i|\mathcal{A}\beta) \neq 0$. In other words, the coarse-grained experiment $(\mathcal{A}, \mathbb{E})$ does not alter the probability assignment for the fine-grained experiment $(\mathcal{B}, \mathbb{F})$, provided that the agent performs the action \mathcal{A}' between them. Intuitively, if an experiment does not disturb any of its refinements then this experiment alters the belief as little as possible, this being a central property of projective measurements in quantum theory [20,24,25].

Consider now the family of experiments of the form $(\mathcal{A}, \mathbb{E})$ where the action \mathcal{A} is fixed and the partition \mathbb{E} is variable. If there exists an action \mathcal{A}' such that condition (3) holds for all partitions \mathbb{E} , then we call each experiment $(\mathcal{A}, \mathbb{E})$ *ideal*. Such experiments exist in every theory:

Theorem 1. Ideal experiments exist in every GBT. In particular, every principal experiment \mathbb{E} is ideal.

The theorem is proved in the Supplemental Material [23], where we also prove that the set of ideal experiments in quantum theory coincides with the set of projective measurements.

The emergence of the exclusivity principle. The exclusivity principle, as originally introduced in the literature [12–18], refers to scenarios where multiple alternative experiments share some of their outcomes, as in the example of Fig. 1. Let \mathcal{S} be the set of experiments under consideration, and, for each experiment $\mathcal{E} \in \mathcal{S}$, let $\mathcal{O}_{\mathcal{E}}$ be the set of its possible outcomes. The outcomes of all experiments are assigned probabilities, with the constraint that the probabilities should sum up to 1 for every experiment, namely $\sum_{x \in \mathcal{O}_{\mathcal{E}}} p(x) = 1, \forall \mathcal{E} \in \mathcal{S}$. A set of outcomes \mathcal{O} , possibly belonging to different experiments, is called *pairwise exclusive* if, for every pair of outcomes $\{x_1, x_2\} \subseteq \mathcal{O}$, there exists an experiment $\mathcal{E} \in \mathcal{S}$ such that both x_1 and x_2 belong to $\mathcal{O}_{\mathcal{E}}$. A set of experiments \mathcal{S} satisfies the exclusivity principle if the condition $\sum_{x \in \mathcal{O}} p(x) \leq 1$ holds for every set \mathcal{O} of pairwise exclusive outcomes. This

is a nontrivial property, and is not satisfied by all probability assignments. In the example of Fig. 1, the assignment $p(a) = 1/2$, $p(b) = p(d) = 3/8$, $p(c) = p(f) = 1/8$, $p(e) = 1/4$ satisfies the condition $p(a) + p(d) + p(f) = p(d) + p(b) + p(e) = p(a) + p(b) + p(c) = 1$ but violates the exclusivity principle.

We now show how to translate the exclusivity principle in the language of GBTs. The principle applies to outcomes of ideal experiments. By an outcome of an ideal experiment $\mathcal{E} = (\mathcal{A}, \mathbb{E})$ we mean a pair $x = (\mathcal{A}, E)$ with $E \in \mathbb{E}$. The set of outcomes of the experiment \mathcal{E} is then $O_{\mathcal{E}} := \{(\mathcal{A}, E) \mid E \in \mathbb{E}\}$.

To say that two experiments share an outcome, as in Fig. 1, we need a criterion to identify outcomes of different experiments. For two experiments $\mathcal{E} = (\mathcal{A}, \mathbb{E})$ and $\mathcal{E}' = (\mathcal{A}', \mathbb{E}')$, we say that two outcomes $(\mathcal{A}, E) \in O_{\mathcal{E}}$ and $(\mathcal{A}', E') \in O_{\mathcal{E}'}$ are *equivalent*, denoted $(\mathcal{A}, E) \simeq (\mathcal{A}', E')$, if they have the same probability for every possible belief, namely $p(E|\mathcal{A}\beta) = p(E'|\mathcal{A}'\beta)$, $\forall \beta \in \mathbf{B}$. Diagrams such as the one in Fig. 1 arise when equivalent outcomes are identified.

Two outcomes $(\mathcal{A}, E) \in O_{\mathcal{E}}$ and $(\mathcal{A}', E') \in O_{\mathcal{E}'}$, possibly arising from different experiments $\mathcal{E} \neq \mathcal{E}'$, are *mutually exclusive* if (\mathcal{A}, E) and (\mathcal{A}', E') are equivalent to two distinct outcomes of a single experiment \mathcal{F} , namely $(\mathcal{A}, E) \simeq x$ and $(\mathcal{A}', E') \simeq x'$ for $\{x, x'\} \subseteq O_{\mathcal{F}}$ and $x \neq x'$. A set of outcomes O (of generally different experiments) is pairwise exclusive if, for every pair of outcomes $\{x, x'\} \subset O$, x and x' are mutually exclusive.

We are now ready to translate the exclusivity principle in the framework of general Bayesian theories. The principle states that for every set $O = \{(\mathcal{A}_n, E_n)\}$ of pairwise exclusive outcomes of ideal experiments, the condition

$$\sum_n p(E_n|\mathcal{A}_n\beta) \leq 1 \quad (4)$$

must be satisfied for every belief $\beta \in \mathbf{B}$.

The central result of our paper is that the bound (4) holds in every GBT:

Theorem 2. In every GBT, the outcomes of ideal experiments satisfy the exclusivity principle.

The proof is based on the refinability property of ideal experiments, which allows us to construct a sequence of binary experiments with the property that a subset of outcomes is equivalent to the original set of pairwise exclusive outcomes. Since all the outcomes in the pairwise exclusive set can be turned into outcomes of a single measurement, the bound (4) then follows from the normalization of the probability distribution for that measurement. The details of the proof are provided in the Supplemental Material [23].

Combining our derivation of the exclusivity principle with the result of Ref. [19], we obtain that the set of quantum correlations can be characterized completely in terms of Bayesian consistency conditions:

Corollary 1. In every GBT where every two statistically independent ideal experiments can be performed in parallel, the largest set of correlations arising from ideal experiments coincides with the quantum set for every Bell or Kochen-Specker contextuality scenario.

The proof of Corollary 1 follows from a proof in Ref. [19], by replacing Lemma 1 in Ref. [19] with our Theorem 2 and Assumption 2 in Ref. [19] with the assumption that any two experiments can be performed in parallel with a statistically independent joint distribution.

Conclusion. No interpretation of quantum theory can be considered satisfactory unless one can derive the formalism of the theory from the key elements of the interpretation [26]. This problem is especially pressing for QBism, which views quantum theory as a set of rules that constrain how agents make bets about the outcomes of their experiments. Can the quantum formalism be obtained from Bayesian consistency conditions alone?

To address this question, we have introduced a framework for general Bayesian theories. The core of our framework are three consistency conditions on how the agent updates its beliefs: the law of conditional probabilities, forward consistency, and backward consistency. These conditions are sufficient to prove the existence of a set of ideal experiments, which in quantum theory coincide with the projective measurements. We showed that the correlations arising from ideal experiments must satisfy the exclusivity principle, which implies that the set of correlations for every Bell or Kochen-Specker contextuality scenario is equal to the corresponding set in quantum theory, provided that any two statistically independent ideal experiments admit a joint realization.

Our result shows that the set of quantum correlations in Bell and Kochen-Specker contextuality scenarios can be derived from elementary consistency rules on how an agent should bet on the outcomes of future experiments. The problem of whether the whole quantum formalism can be recovered solely from Bayesian consistency conditions is still open. Unlike earlier reconstructions of quantum theory [27–33] that adopted an interpretation agnostic approach, a fully Bayesian reconstruction would have the potential to shed new light on the interpretation of quantum theory.

Acknowledgments. We thank Debbie Leung, Hans J. Briegel, Časlav Brukner, Christopher A. Fuchs, Rüdiger Schack, and Matthew Pusey for discussions during the development of this work. This work is supported by the Foundational Questions Institute through Grant No. FQXi-RFP3-1325, the National Natural Science Foundation of China through Grant No. 11675136, the Hong Kong Research Grant Council through Grant No. 17300317, the Spanish Ministry of Science, Innovation and Universities (MICINN), the Spanish Ministry of Economy, Industry and Competitiveness (MINECO), the European Regional Development Fund FEDER through Grants No. FIS2017-89609-P and No. FIS2015-67161-P, the ERC (Starting Grant No. 258647/GEDENTQOPT and Consolidator Grant No. 683107/TempoQ), the Basque Government (Grant No. IT986-16), and the Foundational Questions Institute and Fetzer Franklin Fund, a donor advised fund of Silicon Valley Community Foundation (FQXi FFF Grant No. FQXi-RFP-1815). This research was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Research, Innovation and Science.

- [1] H. Everett III, “Relative state” formulation of quantum mechanics, *Rev. Mod. Phys.* **29**, 454 (1957).
- [2] C. A. Fuchs and R. Schack, Quantum-Bayesian coherence, *Rev. Mod. Phys.* **85**, 1693 (2013).
- [3] J. S. Bell, On the Einstein Podolsky Rosen paradox, *Physics* **1**, 195 (1964).
- [4] L. A. Khalfin and B. S. Tsirelson, Quantum/classical correspondence in the light of Bell’s inequalities, *Found. Phys.* **22**, 879 (1992).
- [5] S. Popescu and D. Rohrlich, Quantum nonlocality as an axiom, *Found. Phys.* **24**, 379 (1994).
- [6] P. Rastall, Locality, Bell’s theorem, and quantum mechanics, *Found. Phys.* **15**, 963 (1985).
- [7] G. Brassard, H. Buhrman, N. Linden, A. A. Méthot, A. Tapp, and F. Unger, Limit on Nonlocality in Any World in which Communication Complexity is Not Trivial, *Phys. Rev. Lett.* **96**, 250401 (2006).
- [8] N. Linden, S. Popescu, A. J. Short, and A. Winter, Quantum Nonlocality and Beyond: Limits from Nonlocal Computation, *Phys. Rev. Lett.* **99**, 180502 (2007).
- [9] M. Pawłowski, T. Paterek, D. Kaszlikowski, V. Scarani, A. Winter, and M. Żukowski, Information causality as a physical principle, *Nature (London)* **461**, 1101 (2009).
- [10] M. Navascués and H. Wunderlich, A glance beyond the quantum model, *Proc. R. Soc. London, Ser. A* **466**, 881 (2009).
- [11] T. Fritz, A. B. Sainz, R. Augusiak, J. Bohr Brask, R. Chaves, A. Leverrier, and A. Acín, Local orthogonality as a multipartite principle for quantum correlations, *Nat. Commun.* **4**, 2263 (2013).
- [12] A. Cabello, S. Severini, and A. Winter, (Non-)contextuality of physical theories as an axiom, Report No. 8, 2010/2011, fall, Institut Mittag-Leffler, 2010.
- [13] A. Cabello, Simple Explanation of the Quantum Violation of a Fundamental Inequality, *Phys. Rev. Lett.* **110**, 060402 (2013).
- [14] B. Yan, Quantum Correlations are Tightly Bound by the Exclusivity Principle, *Phys. Rev. Lett.* **110**, 260406 (2013).
- [15] B. Amaral, M. Terra Cunha, and A. Cabello, Exclusivity principle forbids sets of correlations larger than the quantum set, *Phys. Rev. A* **89**, 030101(R) (2014).
- [16] J. Henson, Bounding Quantum Contextuality with Lack of Third-Order Interference, *Phys. Rev. Lett.* **114**, 220403 (2015).
- [17] A. Acín, T. Fritz, A. Leverrier, and A. B. Sainz, A combinatorial approach to nonlocality and contextuality, *Commun. Math. Phys.* **334**, 533 (2015).
- [18] A. Cabello, Simple Explanation of the Quantum Limits of Genuine n -Body Nonlocality, *Phys. Rev. Lett.* **114**, 220402 (2015).
- [19] A. Cabello, Quantum correlations from simple assumptions, *Phys. Rev. A* **100**, 032120 (2019).
- [20] G. Chiribella and X. Yuan, Bridging the gap between general probabilistic theories and the device-independent framework for nonlocality and contextuality, *Inf. Comput.* **250**, 15 (2016).
- [21] L. Krämer and L. del Rio, Operational locality in global theories, *Philos. Trans. R. Soc. A* **376**, 20170321 (2018).
- [22] G. Chiribella, Agents, subsystems, and the conservation of information, *Entropy* **20**, 358 (2018).
- [23] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevResearch.2.042001> for a formulation of quantum theory in the GBT framework, for a proof of Theorem 1, a proof that the set of ideal measurements in quantum theory coincides with the set of projective measurements, and a proof of Theorem 2.
- [24] G. Chiribella and X. Yuan, Measurement sharpness cuts nonlocality and contextuality in every physical theory, [arXiv:1404.3348](https://arxiv.org/abs/1404.3348).
- [25] M. Kleinmann, Sequences of projective measurements in generalized probabilistic models, *J. Phys. A* **47**, 455304 (2014).
- [26] C. A. Fuchs, Quantum mechanics as quantum information, mostly, *J. Mod. Opt.* **50**, 987 (2003).
- [27] L. Hardy, Quantum theory from five reasonable axioms, [arXiv:quant-ph/0101012](https://arxiv.org/abs/quant-ph/0101012).
- [28] G. Chiribella, G. M. D’Ariano, and P. Perinotti, Informational derivation of quantum theory, *Phys. Rev. A* **84**, 012311 (2011).
- [29] B. Dakić and Č. Brukner, Quantum theory and beyond: Is entanglement special? in *Deep Beauty. Understanding the Quantum World through Mathematical Innovation*, edited by H. Halvorson (Cambridge University Press, New York, 2011), p. 365.
- [30] L. Masanes and M. P. Müller, A derivation of quantum theory from physical requirements, *New J. Phys.* **13**, 063001 (2011).
- [31] L. Hardy, Reformulating and reconstructing quantum theory, [arXiv:1104.2066](https://arxiv.org/abs/1104.2066).
- [32] A. Wilce, A royal road to quantum theory (or thereabouts), *Entropy* **20**, 227 (2018).
- [33] P. A. Höhn and C. S. P. Wever, Quantum theory from questions, *Phys. Rev. A* **95**, 012102 (2017).