Linear Program for Testing Nonclassicality and an Open-Source Implementation

John H. Selby,[1,*] Elie Wolfe,[2,†] David Schmid,[1,‡] Ana Belén Sainz,[1,§] and Vinicius P. Rossi[1,∥]
[1]*International Centre for Theory of Quantum Technologies, University of Gdańsk, 80-309 Gdańsk, Poland*
[2]*Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, Ontario N2L 2Y5, Canada*

(Received 16 November 2022; revised 12 November 2023; accepted 22 December 2023; published 30 January 2024)

A well-motivated method for demonstrating that an experiment resists any classical explanation is to show that its statistics violate generalized noncontextuality. We here formulate this problem as a linear program and provide an open-source implementation of it which tests whether or not any given prepare-measure experiment is classically explainable in this sense. The input to the program is simply an arbitrary set of quantum states and an arbitrary set of quantum effects; the program then determines if the Born rule statistics generated by all pairs of these can be explained by a classical (noncontextual) model. If a classical model exists, it provides an explicit model. If it does not, then it computes the minimal amount of noise that must be added such that a model does exist, and then provides this model. We generalize all these results to arbitrary generalized probabilistic theories (and accessible fragments thereof) as well; indeed, our linear program is a test of simplex embeddability as introduced in Schmid *et al.* [PRX Quantum **2**, 010331 (2021).] and generalized in Selby *et al.* [Phys. Rev. A **107**, 062203 (2023).].

DOI: 10.1103/PhysRevLett.132.050202

A rigorous method for demonstrating that a theory or a set of data resists any classical explanation is to prove that it cannot be reproduced in any generalized noncontextual model [1]. Generalized noncontextuality was first introduced as an improvement on Kochen-Specker's assumption of noncontextuality [2], making it more operationally accessible and providing stronger motivations for it, as a form of Leibniz's principle [3]. Since its inception, the list of motivations for taking it as one's notion of classicality has grown greatly. Notably, the existence of a generalized-noncontextual ontological model for an operational theory coincides with two independent notions of classicality: one that arises in the study of generalized probabilistic theories [4–6], and another that arises in quantum optics [4,6,7]. Generalized noncontextuality has been used as an indicator of classicality in the quantum Darwinist program [8], and any sufficiently noisy theory satisfies generalized noncontextuality [9,10]. Furthermore, violations of local causality [11], violations of Kochen-Specker noncontextuality [9,12], and some observations of anomalous weak values [13,14], are also instances of generalized contextuality. Finally, generalized contextuality is a resource for information processing [15–19], computation [20], state discrimination [21–24], cloning [25], and metrology [26]. Herein, we use the term *noncontextuality* to refer to the concept of generalized noncontextuality.

How, then, does one determine in practice whether a given theory or a given set of experimental data admits a classical explanation of this sort? We here provide the most direct algorithm to date for answering this question in arbitrary prepare-and-measure experiments, and we provide open-access *Mathematica* code for answering it in practice. One need only give a finite set of quantum states and a finite set of quantum POVM elements as input, and the code determines if the statistics these generate by the Born rule can be explained classically—i.e., by a noncontextual ontological model for the operational scenario. It furthermore returns an explicit noncontextual model, if one exists. If there is no such model, the code determines an operational measure of nonclassicality, namely, the minimum amount of noise which would be required until a noncontextual model would become possible.

In the Supplemental Material [27], we generalize these ideas beyond quantum theory to the case of arbitrary generalized probabilistic theories (GPTs) [46,47] or fragments thereof, leveraging the fact that an operational scenario admits a noncontextual model if and only if the corresponding GPT admits a simplex embedding [4]. Indeed, the linear program we derive is simply a test of whether any valid simplex embedding (of any dimension) can be found, answering the challenge first posed in Ref. [4]. We furthermore prove an upper bound on the number of ontic states needed in any such classical explanation, namely, the square of the GPT dimension.

The Supplemental Material [27] also explains how our open-source code implements the linear program we develop herein.

A large number of previous works have studied the question of when a set of data admits a generalized noncontextual model [4–6,48–55]. Most closely related to our work are Refs. [5,49,50,52]. We elaborate on the relationships between these works in our conclusion and in Ref. [27].

For now, we simply note that the linear program (and dimension bound) that we derive here is closely related to an optimization problem introduced in Ref. [52]. However, Ref. [52] focuses on a proposed modification of generalized noncontextuality (which we criticize in the Supplemental Material [27]), and so the two approaches do not always return the same result.

Our Letter aims to be accessible and self-contained, in order to provide a tool for the quantum information and foundations communities to directly test for nonclassicality in their own research problems.

A linear program for deciding classicality.—We now set up the preliminaries required to state our linear program for testing whether the quantum statistics generated by given sets of quantum states and effects can be explained classically—i.e., by a noncontextual model for the operational scenario. The Supplemental Material [27] generalizes these ideas and results to arbitrary GPTs.

Consider any finite set of (possibly subnormalized [56]) quantum states Ω, and any finite set of quantum effects $\mathcal{E}$, living in the real vector space $\mathsf{Herm}[\mathcal{H}]$ of Hermitian operators on some finite dimensional Hilbert space $\mathcal{H}$. In general, neither the set of states nor the set of effects need span the full vector space $\mathsf{Herm}[\mathcal{H}]$, nor need the two sets span the same subspace of $\mathsf{Herm}[\mathcal{H}]$. Next, we introduce some useful mathematical objects related to Ω and $\mathcal{E}$.

Let us first focus on the case of states. We denote the subspace of $\mathsf{Herm}[\mathcal{H}]$ spanned by the states Ω by S_Ω. The inclusion map from S_Ω to $\mathsf{Herm}[\mathcal{H}]$ is denoted by I_Ω. In addition, we define the cone of positive operators that arises from Ω by

$$\mathsf{Cone}[\Omega] = \left\{\rho \middle| \rho = \sum_\alpha r_\alpha \rho_\alpha, \rho_\alpha \in \Omega, r_\alpha \in \mathbb{R}^+\right\} \subset S_\Omega. \quad (1)$$

This cone can also be characterized by its facet inequalities, indexed by $i = \{1, \ldots, n\}$, where n is necessarily finite as we start with a finite set of states (see, for example, McMullen's upper bound theorem [57]). These inequalities are specified by Hermitian operators $h_i^\Omega \in S_\Omega$ such that

$$\mathsf{tr}(h_i^\Omega v) \geq 0 \quad \forall\ i \Leftrightarrow v \in \mathsf{Cone}[\Omega]. \quad (2)$$

From these facet inequalities, one can define a linear map $H_\Omega\colon S_\Omega \to \mathbb{R}^n$, such that

$$H_\Omega(v) = (\mathsf{tr}(h_1^\Omega v), \ldots, \mathsf{tr}(h_n^\Omega v))^T \quad \forall\ v \in S_\Omega. \quad (3)$$

Note that the matrix elements of $H_\Omega(v)$ are all non-negative if and only if $v \in \mathsf{Cone}[\Omega]$. We denote entrywise non-negativity by $H_\Omega(v) \geq_e 0$ (to disambiguate from using ≥ 0 to represent positive semi-definiteness). Succinctly, we have

$$H_\Omega(v) \geq_e 0 \Leftrightarrow v \in \mathsf{Cone}[\Omega], \quad (4)$$

and so H_Ω is simply an equivalent characterization of the cone.

Consider now the set of effects $\mathcal{E}$. We denote the subspace of $\mathsf{Herm}[\mathcal{H}]$ spanned by $\mathcal{E}$ by $S_\mathcal{E}$, and the inclusion map from $S_\mathcal{E}$ to $\mathsf{Herm}[\mathcal{H}]$ by $I_\mathcal{E}$. In addition, we define the cone of positive operators that arises from $\mathcal{E}$ as

$$\mathsf{Cone}[\mathcal{E}] = \left\{\gamma \middle| \gamma = \sum_\beta r_\beta \gamma_\beta, \gamma_\beta \in \mathcal{E}, r_\beta \in \mathbb{R}^+\right\} \subset S_\mathcal{E}. \quad (5)$$

This cone can also be characterized by its facet inequalities, indexed by $j = \{1, \ldots, m\}$, where m is again finite, as we are considering a finite set of effects. These inequalities are specified by Hermitian operators $h_j^\mathcal{E}$ such that

$$\mathsf{tr}(h_j^\mathcal{E} w) \geq 0 \quad \forall\ j \Leftrightarrow w \in \mathsf{Cone}[\mathcal{E}]. \quad (6)$$

From these facet inequalities one can define a linear map $H_\mathcal{E}\colon S_\mathcal{E} \to \mathbb{R}^m$, such that

$$H_\mathcal{E}(w) = (\mathsf{tr}(w h_1^\mathcal{E}), \ldots, \mathsf{tr}(w h_m^\mathcal{E}))^T \quad \forall\ w \in S_\mathcal{E}. \quad (7)$$

This fully characterizes $\mathsf{Cone}[\mathcal{E}]$, since

$$H_\mathcal{E}(w) \geq_e 0 \Leftrightarrow w \in \mathsf{Cone}[\mathcal{E}]. \quad (8)$$

One can also pick an arbitrary orthonormal basis of Hermitian operators for each of the spaces $\mathsf{Herm}[\mathcal{H}]$, S_Ω, and $S_\mathcal{E}$, and represent I_Ω, $I_\mathcal{E}$, $H_\mathcal{E}$, and H_Ω as matrices with respect to these.

With these defined, we can now present the linear program which tests for classical explainability (i.e., simplex embeddability) of any set of quantum states and any set of quantum effects in terms of the matrices I_Ω, $I_\mathcal{E}$, H_Ω, and $H_\mathcal{E}$, defined above and computed from the set of states and set of effects.

Linear program 1.—The Born rule statistics obtained by composing any state-effect pair from Ω and $\mathcal{E}$ is classically explainable if and only if the following linear program is satisfiable:

$$\exists\, \sigma \geq_e 0, \qquad \text{an } m \times n \text{ matrix such that} \quad (9a)$$

$$I_\mathcal{E}^T \cdot I_\Omega = H_\mathcal{E}^T \cdot \sigma \cdot H_\Omega. \quad (9b)$$

Note that if Ω and $\mathcal{E}$ span the full vector space of Hermitian operators, then the linear program simplifies somewhat, as the l.h.s. of Eq. (9b) reduces to the identity map on $\mathsf{Herm}[\mathcal{H}]$. Note that satisfiability is only a function of the cones defined by Ω and by $\mathcal{E}$, and so no other features of the states and effects are relevant to their nonclassicality, as was also shown in Refs. [48,53]. A useful consequence of this fact is that Ω and $\mathcal{E}$ are classically explainable if and only if their convex hulls are also classically explainable.

Testing for the existence of such a σ is a linear program. In the repository [58], we give open-source *Mathematica* code for computing the relevant preliminaries and solving this linear program. The input to the code is simply a set of density matrices and a set of POVM elements (or, more generally, GPT state and effect vectors). In practice the code runs in a few seconds for values of n and m up to around 20.

In the case that a classical explanation does exist, the code will output a specification of an ontological model which represents the operational scenario in a noncontextual manner. This model can be computed from the matrix σ, as described in the Supplemental Material [27]. In particular, every density matrix in $\rho \in \Omega$ is represented in the ontological model by a probability distribution μ_ρ over some set of ontic states Λ, while every POVM element in $\varepsilon \in \mathcal{E}$ is represented by a response function ξ_ε—that is, a [0, 1]-valued function over Λ. Specifically, we compute a particular non-negative factorization $\sigma = \beta \cdot \alpha$, where $\alpha\colon \mathbb{R}^n \to \mathbb{R}^\Lambda \geq_e 0$ and $\beta\colon \mathbb{R}^\Lambda \to \mathbb{R}^m \geq_e 0$, and then construct linear maps $\tau_\Omega \coloneqq \alpha \cdot H_\Omega$ and $\tau_\mathcal{E} \coloneqq \beta^T \cdot H_\mathcal{E}$, and use these to define the epistemic states and response functions via

$$\mu_\rho(\lambda) \coloneqq [\tau_\Omega(\rho)]_\lambda \quad \text{and} \quad \xi_\varepsilon(\lambda) \coloneqq [\tau_\mathcal{E}(\varepsilon)]_\lambda \tag{10}$$

for all $\lambda \in \Lambda$. That these functions are all non-negative follows from the definition of H_Ω and $H_\mathcal{E}$ together with element-wise non-negativity of α and β; that they are suitably normalized follows from the manner in which the decomposition into α and β is chosen. In particular, the decomposition is constructed by taking $\beta = \sigma \cdot R$ and $\alpha = R^{-1}$, where R is a diagonal rescaling matrix which ensures that $\xi_\mathbb{1}(\lambda) = 1$ for all $\lambda \in \Lambda$ (see Supplemental Material [27], Sec. C.I for details). Note that other choices for the decomposition of $\sigma = \beta \cdot \alpha$ are possible, and that this nonuniqueness translates into a nonuniqueness of the ontological model.

In the case that no solution exists, one can ask how much depolarizing noise must be added to one's experiment until a solution becomes possible. This constitutes an operational measure of nonclassicality which we refer to as the *robustness of nonclassicality*. Finding the minimal amount r of noise is also a linear program:

Linear program 2.—Let r be the minimum depolarising noise that must be added in order for the statistics obtained by composing any state-effect pair from Ω and $\mathcal{E}$ to be classically explainable. It can be computed by the linear program:

minimize r such that

$$\exists\, \sigma \geq_e 0, \text{ an } m \times n \text{ matrix such that} \tag{11a}$$

$$r I_\mathcal{E}^T \cdot D \cdot I_\Omega + (1-r) I_\mathcal{E}^T \cdot I_\Omega = H_\mathcal{E}^T \cdot \sigma \cdot H_\Omega, \tag{11b}$$

where D is the completely depolarizing channel for the quantum system.

Again, the corresponding ontological model can be straightforwardly computed from the matrix σ found for the minimal value of r, and we give open-source code that returns both the value of r and the associated model.

We also discuss in the Supplemental Material [27] how one can easily adapt one's definition of robustness and the linear program for it to an arbitrary noise model.

Examples.—Here we present three examples of sets of states and effects, and we assess the classical-explainability of their statistics using our linear program. In the case where the statistics are not classical, we also compute the noise robustness. A fully detailed analysis of these examples (including the explicit calculation of the matrices H_Ω, $H_\mathcal{E}$, I_Ω, and $I_\mathcal{E}$), is given in the Supplemental Material [27]. These specific examples are chosen to illustrate particular features of our approach, as we discuss therein.

Example 1.—Consider the set of four quantum states

$$\Omega = \{|0\rangle\langle 0|, |1\rangle\langle 1|, |+\rangle\langle +|, |-\rangle\langle -|\} \tag{12}$$

on a qubit. In addition, consider the set of six effects

$$\mathcal{E} = \{|0\rangle\langle 0|, |1\rangle\langle 1|, |+\rangle\langle +|, |-\rangle\langle -|, \mathbb{1}_2, 0\}. \tag{13}$$

Next, consider the observable statistics—that is, the data that can be generated from any measurement constructed with these effects, when applied to any of these states.

Our linear program finds that these statistics admit a classical explanation. This is to be expected, as this scenario is a subtheory of the noncontextual toy theory of Ref. [59] (namely, that given by restricting to the real plane). Indeed, this is the model which our code returns, and is depicted in Fig. 1.

Example 2.—Consider the set of four quantum states

$$\Omega = \{|0\rangle\langle 0|, |1\rangle\langle 1|, |2\rangle\langle 2|, |3\rangle\langle 3|\} \tag{14}$$

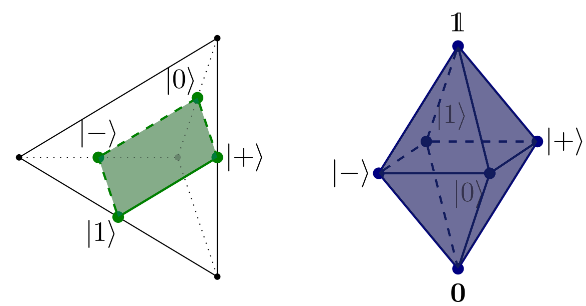

(a) Embedding of states (b) Embedding of effects

FIG. 1. Classical explanation for example 1. (a) Depiction of the states in Ω (green dots), embedded in a three-dimensional slice of a four-dimensional simplex. (b) Depiction of the effects in $\mathcal{E}$ (blue dots), embedded in a 3D slice of the 4D hypercube that is dual to the simplex in (a). Note that the convex hull of the effects happens to cover the entire hypercube in this particular slice. The simplex in (a) can be viewed as the set of probability distributions over a 4-element set Λ of ontic states (black dots), while the hypercube in (b) can be viewed as the set of logically possible response functions for Λ. Hence, this simplex embedding corresponds to a noncontextual ontological model for—and hence [4] a classical explanation of—the operational scenario.

on a four-dimensional quantum system. In addition, consider the set of six effects

$$\mathcal{E} = \{|0\rangle\langle 0| + |1\rangle\langle 1|, |1\rangle\langle 1| + |2\rangle\langle 2|, |2\rangle\langle 2| + |3\rangle\langle 3|, |3\rangle\langle 3| + |0\rangle\langle 0|, \mathbb{1}_4, 0\}. \tag{15}$$

Notably, the states and effects in this example do not span the same vector space. Still, our linear program also finds that the statistical data that arise from this admits a classical explanation. This is a useful sanity check, since all the states and effects are diagonal in the same basis. We provide a depiction of the classical model which our code returns for this scenario in Fig. 2.

Example 3.—Our third example is obtained from the first example by rotating all of the effects by an angle of $(\pi/4)$ about the σ_y axis. (This is the set of states and effects relevant for parity-oblivious multiplexing [15].) In this case, our linear program finds that there is no classical explanation of the observable statistics. Moreover, it finds that the depolarizing-noise robustness for these states and effects is $r = 1 - (1/\sqrt{2}) \sim 0.3$. In Fig. 3 we depict the classical model for the case of depolarization at this noise threshold.

Related linear programs.—We reiterate that the core of our linear program is closely related to the linear program introduced in Sec. 4.2 of Ref. [52] as specialized to the polytopic case (that we consider here) in Sec. 4.3 of Ref. [52]. However, the approach of Ref. [52] differs from ours in a critical preprocessing step, and so its assessment of classicality differs from ours in some examples. Indeed, their proposal deems example 2 nonclassical, while our approach deems it classical. But, the "nonclassical" verdict is clearly mistaken, since all the states and effects in that example are simultaneously diagonalizable. Still, we emphasize that the mathematical tools of Ref. [52] are quite useful and applicable to our notion of classicality, and indeed even extend some results to non-polytopic GPTs (although in this case, testing for nonclassicality is likely not a linear program) via inner and outer polytopic approximations as discussed in Sec. 4.4 of Ref. [52].

Reference [50] also presented a linear programming approach which could determine if a prepare-measure scenario admits a noncontextual model or not. In that work, however, the input to the linear program required the specification of a set of operational equivalences for the states and another set for the effects. In contrast, in the current work, the input to the algorithm is simply a set of quantum (or GPT) states and effects. The full set of operational equivalences that hold among states and among effects are derivable from this input; however, one need not consider them explicitly. The linear program we present here determines if there is a noncontextual model with respect to *all* of the operational equivalences that hold in quantum theory (or within the given GPT).

Reference [49] provided another linear programming approach to testing noncontextuality in the context of a particular class of prepare-measure scenarios; namely, those wherein all operational equivalences arise from different ensembles of preparation procedures, all of which define the same average state. Using the flag-convexification technique of Refs. [48,53], we suspect that *all* prepare-measure

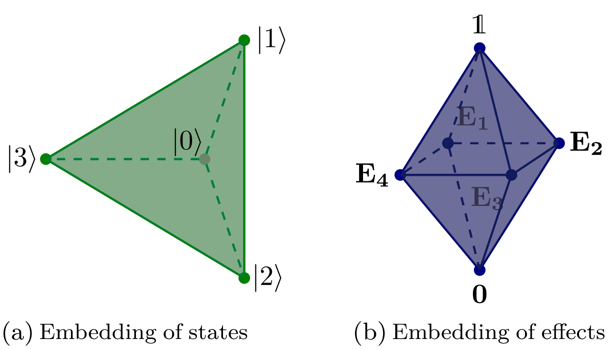

FIG. 2. Classical explanation for example 2. (a) Depiction of the states in Ω (green dots), embedded in a 3D slice of a 4D simplex. (b) Depiction of the effects in $\mathcal{E}$ (blue dots), embedded in a 3D slice of the 4D hypercube that is dual to the simplex in (a). Note that the convex hull of the states (effects) happens to cover the entire simplex (hypercube) in this particular slice. Exactly as in the last example, this simplex embedding corresponds to a noncontextual ontological model for—and hence [4] a classical explanation of—the operational scenario.

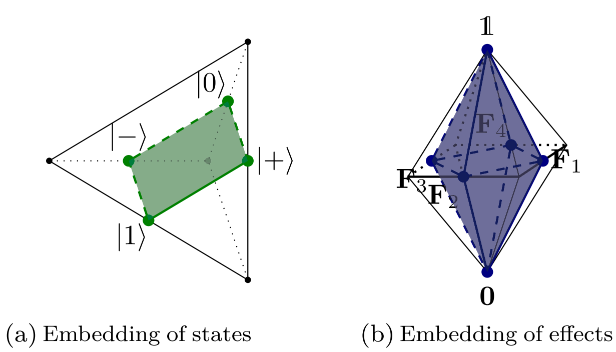

FIG. 3. Classical explanation for example 3, when depolarized by $r = 1 - (1/\sqrt{2})$. (a) Depiction of the states in Ω (green dots), embedded in a 2D slice of a 4D simplex. (b) Depiction of the effects in $\mathcal{E}$ (blue dots), embedded in a 3D slice of the 4D hypercube that is dual to the simplex in (a). Exactly as in the last example, this simplex embedding corresponds to a noncontextual ontological model for—and hence [4] a classical explanation of—the depolarized operational scenario. If the depolarization was less strong, then such a noncontextual ontological model would not exist. Visually, we can get some intuition for this by observing that if we grow either the green square or the blue octahedron, then we would end up with the states or effects lying outside of the simplex or hypercube.

scenarios can be transformed into prepare-measure scenarios of this particular type, in which case the linear program from Ref. [49] would be as general as the approach we have discussed herein. However, this remains to be proven.

An interesting open question is to determine the relative efficiency of these algorithms.

Closing remarks.—Our arguments in the Supplemental Material [27] demonstrate that if a noncontextual model exists for a scenario, then there also exists a model with $\mathsf{dim}[S_\Omega]\,\mathsf{dim}[S_\mathcal{E}] \leq \mathsf{dim}[\mathcal{H}]^2$ ontic states (or less), This bound was first proven in Ref. [52] by similar arguments. It is not yet clear if this bound is tight.

Additionally, our arguments in the Supplemental Material [27] hinge on the existence of a particular kind of decomposition of the identity channel. The arguments proving a structure theorem for noncontextual models in Ref. [6] hinged on a similar decomposition of the identity channel, and it would be interesting to investigate this connection further. We hope that a synthesis of the algorithmic techniques herein with the compositional techniques of Refs. [6,60] might lead to algorithms for testing nonclassicality in prepare-transform-measure scenarios and eventually in arbitrary circuits.

In Ref. [54], the definition of simplex embedding was generalized to embeddings into arbitrary GPTs. It would be interesting to investigate whether similar programs (albeit most likely not linear ones [61].) could be developed for testing for such embeddings.

Finally, we note that our linear program and open source implementation are ideally suited for proving nonclassicality in real experiments [62], especially when coupled with theory-agnostic tomography techniques [63,64].

We acknowledge useful feedback from Rob Spekkens. J. H. S. was supported by the National Science Centre, Poland (Opus project, Categorical Foundations of the Non-Classicality of Nature, Project No. 2021/41/B/ST2/03149). A. B. S., D. S., and V. P. R. were supported by the Foundation for Polish Science (IRAP project, ICTQT, Contract No. MAB/2018/5, co-financed by EU within Smart Growth Operational Programme). E. W. was supported by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported in part 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 Colleges and Universities. All diagrams were made using TikZiT.

*john.h.selby@gmail.com
†ewolfe@perimeterinstitute.ca
‡davidschmid10@gmail.com
§ana.sainz@ug.edu.pl
‖prettirossi.vinicius@gmail.com

[1] R. W. Spekkens, Contextuality for preparations, transformations, and unsharp measurements, Phys. Rev. A **71**, 052108 (2005).

[2] S. Kochen and E. Specker, The problem of hidden variables in quantum mechanics, J. Math. Mech. **17**, 59 (1967), also available from the Indiana University Mathematics Journal

[3] R. W. Spekkens, The ontological identity of empirical indiscernibles: Leibniz's methodological principle and its significance in the work of Einstein, arXiv:1909.04628.

[4] D. Schmid, J. H. Selby, E. Wolfe, R. Kunjwal, and R. W. Spekkens, Characterization of noncontextuality in the framework of generalized probabilistic theories, PRX Quantum **2**, 010331 (2021).

[5] F. Shahandeh, Contextuality of general probabilistic theories, PRX Quantum **2**, 010330 (2021).

[6] D. Schmid, J. H. Selby, M. F. Pusey, and R. W. Spekkens, A structure theorem for generalized-noncontextual ontological models, arXiv:2005.07161.

[7] R. W. Spekkens, Negativity and contextuality are equivalent notions of nonclassicality, Phys. Rev. Lett. **101**, 020401 (2008).

[8] R. D. Baldijão, R. Wagner, C. Duarte, B. Amaral, and M. T. Cunha, Emergence of noncontextuality under quantum Darwinism, PRX Quantum **2**, 030351 (2021).

[9] R. Kunjwal and R. W. Spekkens, From the Kochen-Specker theorem to noncontextuality inequalities without assuming determinism, Phys. Rev. Lett. **115**, 110403 (2015).

[10] I. Marvian, Inaccessible information in probabilistic models of quantum systems, non-contextuality inequalities and noise thresholds for contextuality, arXiv:2003.05984.

[11] J. S. Bell, On the Einstein Podolsky Rosen paradox, Physics **1**, 195 (1964).

[12] R. Kunjwal and R. W. Spekkens, From statistical proofs of the Kochen-Specker theorem to noise-robust noncontextuality inequalities, Phys. Rev. A **97**, 052110 (2018).

[13] M. F. Pusey, Anomalous weak values are proofs of contextuality, Phys. Rev. Lett. **113**, 200401 (2014).

[14] R. Kunjwal, M. Lostaglio, and M. F. Pusey, Anomalous weak values and contextuality: Robustness, tightness, and imaginary parts, Phys. Rev. A **100**, 042116 (2019).

[15] R. W. Spekkens, D. H. Buzacott, A. J. Keehn, B. Toner, and G. J. Pryde, Preparation contextuality powers parity-oblivious multiplexing, Phys. Rev. Lett. **102**, 010401 (2009).

[16] A. Chailloux, I. Kerenidis, S. Kundu, and J. Sikora, Optimal bounds for parity-oblivious random access codes, New J. Phys. **18**, 045003 (2016).

[17] A. Ambainis, M. Banik, A. Chaturvedi, D. Kravchenko, and A. Rai, Parity oblivious d-level random access codes and class of noncontextuality inequalities, Quantum Inf. Process. **18**, 111 (2019).

[18] D. Saha, P. Horodecki, and M. Pawłowski, State independent contextuality advances one-way communication, New J. Phys. **21**, 093057 (2019).

[19] S. A. Yadavalli and R. Kunjwal, Contextuality in entanglement-assisted one-shot classical communication, Quantum **6**, 839 (2022).

[20] D. Schmid, H. Du, J. H. Selby, and M. F. Pusey, Uniqueness of noncontextual models for stabilizer subtheories, Phys. Rev. Lett. **129**, 120403 (2022).

[21] D. Schmid and R. W. Spekkens, Contextual advantage for state discrimination, Phys. Rev. X **8**, 011015 (2018).
[22] K. Flatt, H. Lee, C. R. I. Carceller, J. B. Brask, and J. Bae, Contextual advantages and certification for maximum-confidence discrimination, PRX Quantum **3**, 030337 (2022).
[23] S. Mukherjee, S. Naonit, and A. K. Pan, Discriminating three mirror-symmetric states with a restricted contextual advantage, Phys. Rev. A **106**, 012216 (2022).
[24] J. Shin, D. Ha, and Y. Kwon, Quantum contextual advantage depending on nonzero prior probabilities in state discrimination of mixed qubit states, Entropy **23**, 1583 (2021).
[25] M. Lostaglio and G. Senno, Contextual advantage for state-dependent cloning, Quantum **4**, 258 (2020).
[26] M. Lostaglio, Certifying quantum signatures in thermodynamics and metrology via contextuality of quantum linear response, Phys. Rev. Lett. **125**, 230603 (2020).
[27] See the Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.132.050202 for the Appendices referred to along this Letter, which include Refs. [28–45].
[28] F. Shahandeh, M. J. W. Hall, and T. C. Ralph, Measurement-device-independent approach to entanglement measures, Phys. Rev. Lett. **118**, 150505 (2017).
[29] M. F. Pusey, L. Del Rio, and B. Meyer, Contextuality without access to a tomographically complete set, arXiv:1904.08699.
[30] R. W. Spekkens, Quasi-quantization: Classical statistical theories with an epistemic restriction, in *Quantum Theory: Informational Foundations and Foils*, edited by G. Chiribella and R. W. Spekkens (Springer Netherlands, Dordrecht, 2016), pp. 83–135.
[31] S. D. Bartlett, T. Rudolph, and R. W. Spekkens, Reconstruction of Gaussian quantum mechanics from Liouville mechanics with an epistemic restriction, Phys. Rev. A **86**, 012103 (2012).
[32] L. Catani, M. Leifer, D. Schmid, and R. W. Spekkens, Why interference phenomena do not capture the essence of quantum theory, Quantum **7**, 1119 (2023).
[33] M. Müller, Probabilistic theories and reconstructions of quantum theory, SciPost Phys. Lect. Notes 028 (2021).
[34] L. Lami, Non-classical correlations in quantum mechanics and beyond, arXiv:1803.02902.
[35] M. Plávala, General probabilistic theories: An introduction, Phys. Rep. **1033**, 1 (2023).
[36] G. Chiribella, G. M. D'Ariano, and P. Perinotti, Quantum from principles, in *Fundamental Theories of Physics* (Springer Netherlands, Dordrecht, 2015), pp. 171–221.
[37] L. Hardy, Reformulating and reconstructing quantum theory, arXiv:1104.2066.
[38] G. Chiribella, G. M. D'Ariano, and P. Perinotti, Probabilistic theories with purification, Phys. Rev. A **81**, 062348 (2010).
[39] B. Coecke, T. Fritz, and R. W. Spekkens, A mathematical theory of resources, Inf. Comput. **250**, 59 (2016).
[40] T. Gonda and R. W. Spekkens, Monotones in general resource theories, Compositionality **5**, 7 (2023).
[41] V. P. Rossi, D. Schmid, J. H. Selby, and A. B. Sainz, Contextuality with vanishing coherence and maximal robustness to dephasing, Phys. Rev. A **108**, 032213 (2023).
[42] J. Barrett, N. Linden, S. Massar, S. Pironio, S. Popescu, and D. Roberts, Nonlocal correlations as an information-theoretic resource, Phys. Rev. A **71**, 022101 (2005).
[43] G. Kimura, The Bloch vector for n-level systems, Phys. Lett. A **314**, 339 (2003).
[44] T. S. Motzkin, H. Raiffa, G. L. Thompson, and R. M. Thrall, The double description method, Contrib. Theory Games **2**, 51 (1953).
[45] D. Avis and K. Fukuda, A pivoting algorithm for convex hulls and vertex enumeration of arrangements and polyhedra, in *Proceedings of the Seventh Annual Symposium on Computational Geometry* (1991), pp. 98–104.
[46] L. Hardy, Quantum theory from five reasonable axioms, arXiv:quant-ph/0101012.
[47] J. Barrett, Information processing in generalized probabilistic theories, Phys. Rev. A **75**, 032304 (2007).
[48] J. H. Selby, D. Schmid, E. Wolfe, A. B. Sainz, R. Kunjwal, and R. W. Spekkens, Accessible fragments of generalized probabilistic theories, cone equivalence, and applications to witnessing nonclassicality, Phys. Rev. A **107**, 062203 (2023).
[49] A. Krishna, R. W. Spekkens, and E. Wolfe, Deriving robust noncontextuality inequalities from algebraic proofs of the Kochen–Specker theorem: The Peres-Mermin square, New J. Phys. **19**, 123031 (2017).
[50] D. Schmid, R. W. Spekkens, and E. Wolfe, All the noncontextuality inequalities for arbitrary prepare-and-measure experiments with respect to any fixed set of operational equivalences, Phys. Rev. A **97**, 062103 (2018).
[51] M. F. Pusey, Robust preparation noncontextuality inequalities in the simplest scenario, Phys. Rev. A **98**, 022112 (2018).
[52] V. Gitton and M. P. Woods, Solvable criterion for the contextuality of any prepare-and-measure scenario, Quantum **6**, 732 (2022).
[53] J. H. Selby, D. Schmid, E. Wolfe, A. B. Sainz, R. Kunjwal, and R. W. Spekkens, Contextuality without incompatibility, Phys. Rev. Lett. **130**, 230201 (2023).
[54] M. P. Müller and A. J. P. Garner, Testing quantum theory by generalizing noncontextuality, Phys. Rev. X **13**, 041001 (2023).
[55] A. Chaturvedi, M. Farkas, and V. J. Wright, Characterising and bounding the set of quantum behaviours in contextuality scenarios, Quantum **5**, 484 (2021).
[56] This allows one to describe states which are not prepared deterministically, as happens, e.g., when the preparation is part of a probabilistic source or is the result of remote steering.
[57] P. McMullen, The maximum numbers of faces of a convex polytope, Mathematika **17**, 179 (1970).
[58] Open-source Mathematica code for implementing the linear program introduced in this Letter (2022), https://github.com/eliewolfe/SimplexEmbedding. See also https://github.com/pjcavalcanti/SimplexEmbeddingGPT for a PYTHON implementation of the program.
[59] R. W. Spekkens, Evidence for the epistemic view of quantum states: A toy theory, Phys. Rev. A **75**, 032110 (2007).
[60] D. Schmid, J. H. Selby, and R. W. Spekkens, Unscrambling the omelette of causation and inference: The framework of causal-inferential theories, arXiv:2009.03297.

[61] E.g., in the quantum case these are typically SDP hierarchies [55].

[62] M. D. Mazurek, M. F. Pusey, R. Kunjwal, K. J. Resch, and R. W. Spekkens, An experimental test of noncontextuality without unphysical idealizations, Nat. Commun. **7**, 1 (2016).

[63] M. J. Grabowecky, C. A. J. Pollack, A. R. Cameron, R. W. Spekkens, and K. J. Resch, Experimentally bounding deviations from quantum theory for a photonic three-level system using theory-agnostic tomography, Phys. Rev. A **105**, 032204 (2022).

[64] M. D. Mazurek, M. F. Pusey, K. J. Resch, and R. W. Spekkens, Experimentally bounding deviations from quantum theory in the landscape of generalized probabilistic theories, PRX Quantum **2**, 020302 (2021).