

All the noncontextuality inequalities for arbitrary prepare-and-measure experiments with respect to any fixed set of operational equivalences

David Schmid,^{*} Robert W. Spekkens, and Elie Wolfe*Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, Ontario, Canada N2L 2Y5*

(Received 23 October 2017; published 5 June 2018)

Within the framework of generalized noncontextuality, we introduce a general technique for systematically deriving noncontextuality inequalities for any experiment involving finitely many preparations and finitely many measurements, each of which has a finite number of outcomes. Given any fixed sets of operational equivalences among the preparations and among the measurements as input, the algorithm returns a set of noncontextuality inequalities whose satisfaction is necessary and sufficient for a set of operational data to admit of a noncontextual model. Additionally, we show that the space of noncontextual data tables always defines a polytope. Finally, we provide a computationally efficient means for testing whether any set of numerical data admits of a noncontextual model, with respect to any fixed operational equivalences. Together, these techniques provide complete methods for characterizing arbitrary noncontextuality scenarios, both in theory and in practice. Because a *quantum* prepare-and-measure experiment admits of a noncontextual model if and only if it admits of a positive quasiprobability representation, our techniques also determine the necessary and sufficient conditions for the existence of such a representation.

DOI: [10.1103/PhysRevA.97.062103](https://doi.org/10.1103/PhysRevA.97.062103)

I. INTRODUCTION

Proofs of the failure of noncontextuality place strong constraints on our understanding of nature. For example, the argument by Kochen and Specker [1] showed that quantum theory cannot be explained by an ontological model in which every projective measurement has its outcome determined by hidden variables, independently of what other measurements it is implemented jointly with (that is, independently of context).

A generalized notion of noncontextuality was defined in Ref. [2]. Heuristically, it asserts that an ontological model of some operational theory is noncontextual if and only if laboratory operations which cannot be operationally distinguished are represented identically in the model. It was shown that the operational predictions of quantum theory are inconsistent with the existence of such a model. Furthermore, this generalized notion of noncontextuality (which we henceforth refer to simply as “noncontextuality”) applies to arbitrary operational theories—not just quantum theory—so one can ask whether *any* given set of data describing possible statistics for a prepare-and-measure experiment admits of a noncontextual model. We will refer to such sets of data as *data tables*, since one can organize the data into a table giving the probability for each pairing of a preparation with a measurement outcome.

Because the generalized notion of noncontextuality of Ref. [2] presumes neither the correctness of quantum theory nor that the preparations and measurements are noiseless, its operational consequences can be tested directly with real experimental data.¹ These tests take the form of *noncontext-*

tuality inequalities, whose violation by a data set witnesses the impossibility of explaining that data with a noncontextual ontological model. One such inequality has been violated with high confidence in a recent experiment, demonstrating that nature does not admit of a noncontextual model [7].

This article presents a systematic method for characterizing, for any prepare-and-measure experiment, all the data tables consistent with the principle of noncontextuality given any fixed set of operational equivalences among the preparations and any fixed set of operational equivalences among the measurements. The inputs to our algorithm are these two sets of operational equivalences, and the output is a set of inequalities on the elements of the data table. The satisfaction of these inequalities is necessary and sufficient for the data table to admit of a noncontextual model with respect to the specified operational equivalences. Our method shows that the space of noncontextual data tables for any prepare-and-measure scenario defines a polytope, which we term the *generalized-noncontextual polytope*, in analogy with the *local polytope* of a Bell scenario [9].

In addition, this article provides a systematic and efficient method for deciding if any given data table, specified numerically, admits of a noncontextual model, given any fixed set of operational equivalences.

Foundationally, the failure of noncontextuality plays a key role as a notion of nonclassicality which has broad applicability, and which subsumes other notions such as the negativity of quasiprobability representations [10], the generation of anomalous weak values [11], and violations of local causality

operational theories and thereby to derive experimental tests of this notion. These proposals have been contrasted with those based on generalized noncontextuality and criticized in Refs. [6–8].

^{*}dschmid@perimeterinstitute.ca¹Other articles, e.g., Refs. [3–5], have instead sought to make the Kochen-Specker notion of noncontextuality applicable to arbitrary

[2]. As such, our method for deriving inequalities allows one to quantitatively identify the boundary between the classical and the nonclassical in arbitrary prepare-and-measure scenarios. Furthermore, it allows one to directly generate noise-robust noncontextuality inequalities from arbitrary proofs of the failure of noncontextuality in quantum theory [12], thereby allowing one to find necessary and sufficient conditions for noncontextuality in scenarios where only necessary conditions were previously known [7,8,13–15].

Practically, quantifying the boundary between classical and nonclassical is important for identifying tasks which admit of a quantum advantage. For example, the failure of noncontextuality has been shown to be a resource in various tasks, providing advantages for cryptography [13,16–18] and computation [19–21]. If a given protocol for quantum cryptography or communication, or a given quantum circuit, can be cast into the form of a prepare-and-measure experiment, then one can apply our methods to determine whether the given protocol or circuit admits of a noncontextual model. By this sort of study, one might identify new quantum information-processing tasks for which contextuality is a resource. Our methods should also prove useful for the benchmarking of real quantum devices, just as Bell inequality violations allow one to certify randomness [22] and achieve key distribution [23,24] in a device-independent fashion.

Identifying operational equivalences among laboratory procedures is the first step in deriving observable consequences of generalized noncontextuality. The particular operational equivalences considered, in turn, determine the set of noncontextuality inequalities which follow. As such, it is important to understand where the operational equivalences come from in the first place. Historically, most operational equivalences for generalized noncontextuality have originated in quantum no-go theorems for noncontextuality [7,8,13,14,25,26] or in a frustrated network that can be used to build a quantum no-go theorem [27,28]. These quantum no-go theorems are typically designed to be as simple as possible or to highlight specific structures within quantum theory. One then extrapolates to an operational setting by identifying a set of (imperfect, noisy) operational procedures that satisfy the precise operational equivalence relations that held among the idealized (perfect, noiseless) quantum procedures. By considering procedures that are convex mixtures of the ones actually realized in an experiment, it is possible to satisfy these particular operational equivalences exactly [7]. To approximate what one would expect quantumly, it is important that the experiment be engineered to target the particular operational equivalences in question. However, there is a different approach that one can adopt: rather than trying to engineer one’s experiment to accommodate some specific sets of operational equivalences, one can instead take a set of data from an arbitrary experiment, and then simply determine the operational equivalences that it naturally exhibits.² The only requirement that such experiments must fulfill is that the set of preparations and the set of measurements are tomographically complete (as

²A procedure for finding complete sets of operational equivalences was demonstrated in a special case in Ref. [29]. A more general method is forthcoming.

discussed in depth in Refs. [7,29]). Operational equivalences extracted directly from experimental data will not generally not be as simple in structure as those exhibited in quantum no-go theorems (for instance, the latter often involve *uniform* mixtures and are often symmetric under certain permutations of the procedures). However, the techniques we provide in this paper apply just as well to them. Further, because one is not confined to testing a predetermined set of operational equivalences, one can test for failures of noncontextuality in generic experiments, rather than just those which target laboratory procedures motivated by some specific no-go theorem.

Our method for finding the generalized-noncontextual polytope comprises two distinct computational tasks. The first task is to catalog all extremal solutions which satisfy some initial set of linear constraints; i.e., it is an instance of the vertex enumeration problem. That catalog allows us to formulate a condition for membership in the generalized-noncontextual polytope in terms of an existential quantifier. The second task, then, is quantifier elimination and requires eliminating variables for a system of linear inequalities. As elaborated herein, a variety of standard algorithms are readily available for efficiently solving both of these tasks. Moreover, vertex enumeration and quantifier elimination algorithms are already widely used in quantum foundations.³

Finally, we note that a prepare-and-measure experiment that is consistent with quantum theory (i.e., a set of quantum preparations and a set of quantum measurements) admits of a positive quasiprobability representation if and only if it admits of a noncontextual ontological model [10]. We can therefore apply our technique to solve the problem of determining whether or not a given set of quantum preparations and measurements admits of a positive quasiprobability representation [10,41,42]. It suffices to determine all of the operational equivalences holding among the preparations and among the measurements (see footnote 2) and then determine whether the data table defined by the prepare-and-measure experiment violates any noncontextuality inequality that is implied by these operational equivalences.

II. OPERATIONAL AND ONTOLOGICAL PRELIMINARIES

An operational theory specifies a set of laboratory procedures, such as preparations and measurements, as well as a prescription for finding the probability distribution $p(m|M, P)$ over outcomes m of any given measurement M when implemented following any given preparation P .

³Familiar applications of vertex enumeration and quantifier elimination—albeit not always referred to by these names—include deriving standard Bell inequalities [30–33], deriving entropic inequalities for generalized correlation scenarios [34–36], and many others [8,25,37–40]. For instance, Bell inequalities are defined as the convex hull of all local strategies; i.e., they precisely characterize the region in probability space spanned by deterministic strategies [9]. This relates to vertex enumeration, because the convex hull problem (finding inequalities given extreme points) is *equivalent* to vertex enumeration (finding extreme points given inequalities) from an algorithmic perspective (see, e.g., Appendix A in Ref. [25]).

Two preparations P_1 and P_2 are termed *operationally equivalent* if they generate the same statistics for all possible measurements:

$$\forall M : p(m|M, P_1) = p(m|M, P_2). \quad (1)$$

We denote this operational equivalence relation by $P_1 \simeq P_2$.

Similarly, two measurement procedures M_1 and M_2 are termed *operationally equivalent* if they generate the same probability distribution over their outcomes (denoted m_1 and m_2) for all possible preparations:

$$\forall P : p(m_1|M_1, P) = p(m_2|M_2, P). \quad (2)$$

We denote this operational equivalence relation by $M_1 \simeq M_2$.

We will refer to the event of a measurement M yielding an outcome m as a *measurement effect*, denoted $[m|M]$.

If one samples the choice of laboratory procedure O_i from some probability distribution p_i and then forgets the value of i , we introduce the shorthand notation $\sum_i p_i O_i$ for the effective procedure so defined. The procedures here could be preparations or measurement effects.

An ontological model attempts to explain the probabilities $p(m|M, P)$ in an operational theory via a set Λ of ontic states. An ontic state $\lambda \in \Lambda$ specifies all the physical properties of the system and causally mediates correlations between the preparation and the measurement. For every laboratory preparation P , the model specifies a probability distribution $\mu_P(\lambda)$, where

$$\forall \lambda : \mu_P(\lambda) \geq 0, \quad (3)$$

$$\int_{\Lambda} d\lambda \mu_P(\lambda) = 1. \quad (4)$$

Whenever preparation P is implemented, the ontic state λ is sampled from an associated probability distribution $\mu_P(\lambda)$. Every measurement M generates an outcome m as a probabilistic function of the ontic state according to some fixed *response function* $\xi_{m|M}(\lambda)$, where

$$\forall \lambda, m : \xi_{m|M}(\lambda) \geq 0, \quad (5)$$

$$\forall \lambda : \sum_m \xi_{m|M}(\lambda) = 1. \quad (6)$$

For the ontological model to reproduce an operational theory's empirical predictions, one requires that

$$\forall m, M, P : p(m|M, P) = \int_{\Lambda} \xi_{m|M}(\lambda) \mu_P(\lambda) d\lambda. \quad (7)$$

Finally, note that an effective laboratory procedure $\sum_i p_i O_i$ is represented in an ontological model by the corresponding convex mixture of the ontological representations of the individual operations O_i [see Eq. (7) of Ref. [10] and the surrounding discussion].

An ontological model which respects the principle of *preparation noncontextuality* [2] is one in which two operationally equivalent preparations are represented by the same distribution over ontic states; that is,

$$P_1 \simeq P_2 \text{ implies that} \\ \forall \lambda : \mu_{P_1}(\lambda) = \mu_{P_2}(\lambda). \quad (8)$$

An ontological model which respects the principle of *measurement noncontextuality* [2] is one in which two operationally equivalent measurement effects are represented by the same response function; that is,

$$[m_1|M_1] \simeq [m_2|M_2] \text{ implies that} \\ \forall \lambda : \xi_{m_1|M_1}(\lambda) = \xi_{m_2|M_2}(\lambda). \quad (9)$$

In this paper, the term noncontextuality refers to *universal noncontextuality* [2], which posits noncontextuality for all procedures, including preparations and measurements.

III. PROBLEM SETUP

The scenario we are considering has a set $\{P_1, P_2, \dots, P_g\}$ of g preparations, a set $\{M_1, M_2, \dots, M_l\}$ of l measurements, a set of d outcomes $\{1, 2, \dots, d\}$ for each measurement, a set of operational equivalences among the preparations, denoted $\mathcal{O}\mathcal{E}_P$, and a set of operational equivalences among the measurements, denoted $\mathcal{O}\mathcal{E}_M$. There are no restrictions on any of these sets, beyond the fact that they must be finite, as they are in any real experiment.⁴ Furthermore, we have not presumed that one knows anything about the preparations and measurements, beyond the fact that they can be performed repeatedly so as to gather statistics. Without loss of generality, we treat each measurement as having exactly d outcomes for some sufficiently large value of d (since any measurement with $d^* < d$ outcomes can be redefined to have d outcomes, $d - d^*$ of which never occur).

The input to our algorithm is a specification of the operational equivalences $\mathcal{O}\mathcal{E}_P$ and $\mathcal{O}\mathcal{E}_M$, and the desired output is a set of inequalities such that a data table $\{p(m|M_i, P_j)\}_{i,j,m}$ admits a universally noncontextual model if and only if all the inequalities are satisfied. Although it is not obvious from the definition of universal noncontextuality, we will find that the final inequalities will be linear in the probabilities.

Generically, each of the operational equivalences $s \in \mathcal{O}\mathcal{E}_P$ is of the form

$$\sum_j \alpha_{P_j}^s P_j \simeq \sum_{j'} \beta_{P_{j'}}^s P_{j'} \quad (10)$$

for some sets of convex weights $\{\alpha_{P_j}^s\}_j$ and $\{\beta_{P_{j'}}^s\}_{j'}$ (a set of convex weights is a list of nonnegative real numbers which sum to one). Hence, the principle of preparation noncontextuality, Eq. (8), implies that the same functional relationships must hold among the ontological representations of the preparations; in other words,

$$\forall \lambda : \sum_j \alpha_{P_j}^s \mu_{P_j}(\lambda) = \sum_{j'} \beta_{P_{j'}}^s \mu_{P_{j'}}(\lambda). \quad (11)$$

⁴In any real experiment with continuous variable systems, one must coarse-grain the outcomes to a finite set to obtain nonzero probabilities of any given event. There is also a nuance concerning experiments with a finite number of preparations, measurements and outcomes: the full set of operational equivalences among these might be infinite, but one can always find a finite generating set of operational equivalences whose implications for noncontextual data tables are equivalent to the implications of the full infinite set. See footnote 2.

Similarly, each of the operational equivalences $r \in \mathcal{OE}_M$ is of the form

$$\sum_{i,m} \alpha_{m|M_i}^r [m|M_i] \simeq \sum_{i',m'} \beta_{m'|M_{i'}}^r [m'|M_{i'}] \quad (12)$$

for some sets of convex weights $\{\alpha_{m|M_i}^r\}_{i,m}$ and $\{\beta_{m'|M_{i'}}^r\}_{i',m'}$. Hence, the principle of measurement noncontextuality, Eq. (9), implies that the same functional relationships must hold among the ontological representations of the effects; in other words,

$$\forall \lambda: \sum_{i,m} \alpha_{m|M_i}^r \xi_{m|M_i}(\lambda) = \sum_{i',m'} \beta_{m'|M_{i'}}^r \xi_{m'|M_{i'}}(\lambda). \quad (13)$$

The question of whether a data table admits a universally noncontextual model, then, may be compactly summarized as follows:

Formulation F1 of the existence of a universally noncontextual model. A universally noncontextual model for a data table $\{p(m|M_i, P_j)\}_{i,j,m}$ exists (with respect to the sets of operational equivalences \mathcal{OE}_P and \mathcal{OE}_M) if and only if $\exists \Lambda, \exists \{\mu_{P_j}(\lambda)\}_{j,\lambda}, \{\xi_{m|M_i}(\lambda)\}_{i,m,\lambda}$ such that

$$\forall \lambda, i, m: \xi_{m|M_i}(\lambda) \geq 0, \quad (14a)$$

$$\forall \lambda, i: \sum_m \xi_{m|M_i}(\lambda) = 1, \quad (14b)$$

$$\forall \lambda, r: \sum_{i,m} (\alpha_{m|M_i}^r - \beta_{m|M_i}^r) \xi_{m|M_i}(\lambda) = 0, \quad (14c)$$

$$\forall \lambda, j: \mu_{P_j}(\lambda) \geq 0, \quad (14d)$$

$$\forall j: \int_{\Lambda} \mu_{P_j}(\lambda) = 1, \quad (14e)$$

$$\forall \lambda, s: \sum_j (\alpha_{P_j}^s - \beta_{P_j}^s) \mu_{P_j}(\lambda) = 0, \quad (14f)$$

$$\forall i, j, m: \int_{\Lambda} \xi_{m|M_i}(\lambda) \mu_{P_j}(\lambda) d\lambda = p(m|M_i, P_j). \quad (14g)$$

Equations (14a)–(14g) represent, respectively, positivity of the response functions Eq. (5); normalization of the response functions Eq. (6); the consequences of noncontextuality implied by the operational equivalences in \mathcal{OE}_M Eq. (13); positivity of the distributions associated with the preparations Eq. (3); normalization of the distributions associated with the preparations Eq. (4); the consequences of noncontextuality implied by the operational equivalences in \mathcal{OE}_P Eq. (11); and the expression for the probabilities in the data table in terms of the ontological model Eq. (7).

There are two key obstacles to deriving constraints directly on $\{p(m|M_i, P_j)\}_{i,j,m}$ from the implicit constraints imposed by Eqs. (14a)–(14g). First, the ontic state space Λ is unknown and possibly of unbounded cardinality, so that it is not obvious *a priori* whether there is an algorithm to solve the problem. Second, even if the number of ontic states were known to be finite, so that the problem could in principle be solved by quantifier elimination methods, the probabilities in Eq. (14g) are nonlinear in the unknown parameters $\{\xi_{m|M_i}(\lambda)\}_{i,m,\lambda}$ and $\{\mu_{P_j}(\lambda)\}_{j,\lambda}$ appearing in the quantifiers; hence, the problem would be one of nonlinear quantifier elimination, which is computationally difficult. We overcome both of these problems by leveraging the convex structure of the space of response

functions: we find the finite set of convexly extremal noncontextual assignments to the measurements, identify the set of ontic states with it, and then parametrize the distributions corresponding to the preparations in terms of their probability assignments to these ontic states. Thus, the unknown parameters form a finite set, and furthermore the operational probabilities are linear in these parameters.

Finally, we perform linear quantifier elimination to obtain constraints on the operational probabilities alone.

IV. CHARACTERIZING THE GENERALIZED-NONCONTEXTUAL POLYTOPE

A. Enumerating the convexly extremal noncontextual measurement assignments

No matter what the form or size of the ontic state space Λ , a measurement-noncontextual assignment of probabilities to all d outcomes of all l measurements, for a particular ontic state λ^* , is an (ld) -component vector $\xi(\lambda^*) \equiv (\xi_{1|M_1}(\lambda^*), \dots, \xi_{d|M_1}(\lambda^*), \xi_{1|M_2}(\lambda^*), \dots, \xi_{d|M_l}(\lambda^*))$ subject to the constraints of Eqs. (14a)–(14c). We call such an (ld) -component vector a *noncontextual measurement assignment*. The set of all such assignments defines a polytope:

Characterization P1 of the noncontextual measurement-assignment polytope. The (ld) -component vector $\xi(\lambda^*)$ lies inside the noncontextual measurement-assignment polytope if and only if

$$\forall i, m: \xi_{m|M_i}(\lambda^*) \geq 0, \quad (15a)$$

$$\forall i: \sum_m \xi_{m|M_i}(\lambda^*) = 1, \quad (15b)$$

$$\forall r: \sum_{i,m} (\alpha_{m|M_i}^r - \beta_{m|M_i}^r) \xi_{m|M_i}(\lambda^*) = 0. \quad (15c)$$

In what follows, it is critical to characterize this polytope by its vertices rather than its facets. The vertices are the convexly extremal noncontextual measurement assignments. (Note that if there are *no* operational equivalences among the measurements, then these extremal assignments are deterministic, that is, all of the elements of the vector have value 0 or 1.) In general, to find the vertices of a polytope that is given in terms of its facet inequalities, one must solve the *vertex enumeration problem* [43–47]. Many excellent software packages are freely available for vertex enumeration.⁵

We introduce the notation κ as a discrete variable ranging over the vertices, and we indicate the (explicit) noncontextual measurement assignment of vertex

⁵Dedicated software for performing vertex enumeration includes TRAF from PORTA (<http://porta.zib.de>), SKELETON64F from skeleton (<http://www.uic.unn.ru/zny/skeleton/>), and LCDD_GMP from CCDLIB (https://www.inf.ethz.ch/personal/fukudak/cdd_home), the last notably being readily available on Linux (<https://packages.ubuntu.com/zesty/amd64/libcdd-tools/filelist>) and MacOS (<https://github.com/Homebrew/homebrew-science>). An especially versatile computational geometry suite for Linux is POLYMAKE (<https://polymake.org/doku.php/howto/install#distributions>).

κ^* by the (ld) -component vector $\tilde{\xi}(\kappa^*) \equiv (\tilde{\xi}_{1|M_1}(\kappa^*), \dots, \tilde{\xi}_{d|M_1}(\kappa^*), \tilde{\xi}_{1|M_2}(\kappa^*), \dots, \tilde{\xi}_{d|M_2}(\kappa^*))$.

Now, since any point in a polytope can be written as a convex mixture of the vertices, the noncontextual measurement-assignment polytope can be defined alternatively but equivalently as the convex hull of its vertices:

Characterization P2 of the noncontextual measurement-assignment polytope. The (ld) -component vector $\xi(\lambda^*)$ lies inside the noncontextual measurement-assignment polytope if and only if there exist some convex weights $\{w(\kappa|\lambda^*)\}_\kappa$ such that

$$\forall i, m: \xi_{m|M_i}(\lambda^*) = \sum_{\kappa} w(\kappa|\lambda^*) \tilde{\xi}_{m|M_i}(\kappa), \quad (16)$$

where κ ranges over the vertices found by performing vertex enumeration on the linear constraints of characterization P1.

Below, we presuppose that one has indeed characterized the noncontextual measurement-assignment polytope by finding its vertices explicitly.

B. Constructing a noncontextual model with known ontic states and linearly constrained parameters

Suppose one has a universally noncontextual model of the experiment in the sense of formulation F1, where the ontic state space need not be of finite cardinality. The results of the previous subsection imply that it is always possible to infer from this model another universally noncontextual model wherein the ontic state space is of finite cardinality, as follows.

By substituting Eq. (16) into Eq. (7), each operational probability can be written in terms of a finite sum:

$$p(m|M_i, P_j) = \int_{\Lambda} \xi_{m|M_i}(\lambda) \mu_{P_j}(\lambda) d\lambda \quad (17)$$

$$= \int_{\Lambda} \left[\sum_{\kappa} \tilde{\xi}_{m|M_i}(\kappa) w(\kappa|\lambda) \right] \mu_{P_j}(\lambda) d\lambda \quad (18)$$

$$= \sum_{\kappa} \tilde{\xi}_{m|M_i}(\kappa) \left[\int_{\Lambda} w(\kappa|\lambda) \mu_{P_j}(\lambda) d\lambda \right] \quad (19)$$

$$= \sum_{\kappa} \tilde{\xi}_{m|M_i}(\kappa) v_{P_j}(\kappa), \quad (20)$$

where we have defined

$$v_{P_j}(\kappa) \equiv \int_{\Lambda} w(\kappa|\lambda) \mu_{P_j}(\lambda) d\lambda. \quad (21)$$

Because $v_{P_j}(\kappa^*)$ for a given vertex κ^* is a convex combination of the values of $\mu_{P_j}(\lambda)$, we can infer that each $v_{P_j}(\kappa)$ is a valid probability distribution, and furthermore that the set $\{v_{P_j}(\kappa)\}$ respects noncontextuality with respect to the operational equivalences in \mathcal{OE}_P .

Thus, if any noncontextual ontological model exists, then there must also exist a noncontextual model with an ontic state space of finite cardinality. The latter model is constructed by identifying one ontic state with each extremal noncontextual measurement assignment, and then imagining every preparation as a probability distribution over those ontic states, as done in Eq. (21). In other words,

Formulation F2 of the existence of a universally noncontextual model. For a data table $\{p(m|M_i, P_j)\}_{i,j,m}$, an ontological model that is universally noncontextual with respect to the operational equivalences in \mathcal{OE}_P and \mathcal{OE}_M exists if and only if $\exists \{v_{P_j}(\kappa)\}_{j,\kappa}$ such that

$$\forall \kappa, j: v_{P_j}(\kappa) \geq 0, \quad (22a)$$

$$\forall j: \sum_{\kappa} v_{P_j}(\kappa) = 1, \quad (22b)$$

$$\forall \kappa, s: \sum_j (\alpha_{P_j}^s - \beta_{P_j}^s) v_{P_j}(\kappa) = 0, \quad (22c)$$

$$\forall i, j, m: \sum_{\kappa} \tilde{\xi}_{m|M_i}(\kappa) v_{P_j}(\kappa) = p(m|M_i, P_j), \quad (22d)$$

where κ ranges over the discrete set of vertices of the polytope defined by Eqs. (15a)–(15c) or (16).

In this formulation, each operational probability $p(m|M_i, P_j)$ is given as a linear function of a finite set of unknown parameters. This is because the only unknown parameters on the left-hand side of Eq. (22d) are $\{v_{P_j}(\kappa)\}_\kappa$, while the $\{\tilde{\xi}_{m|M_i}(\kappa)\}_\kappa$ are specified numerically—they are the solution of the vertex enumeration problem described in the previous section. Achieving linearity in all the constraints is a critical intermediate step towards finding a final quantifier-free formulation, as we do in the next section, and is critical for the numerical methods we introduce in Sec. V.

C. The inequalities formulation of the generalized-noncontextual polytope

To obtain constraints that refer only to operational probabilities, we eliminate the unobserved $\{v_{P_j}(\kappa)\}_{j,\kappa}$ from the system of equations (22a)–(22d), obtaining a system of linear inequalities over the $\{p(m|M_i, P_j)\}_{i,j,m}$ alone. The linearity of the final inequalities follows from the linearity of the inequalities and equalities in Eqs. (22a)–(22d). This establishes that the space of noncontextual data tables defines a polytope.

The standard method for solving this problem of *linear quantifier elimination* is the Chernikov-refined Fourier-Motzkin algorithm [48–52], which is implemented in a variety of software packages.⁶

⁶Dedicated software for eliminating variables from a set of linear inequalities via the Fourier-Motzkin algorithm includes FMEL from PORTA (<http://porta.zib.de>), FME from qskeleton (<http://sbastrakov.github.io/qskeleton/>), and FOURIER from Irs (<http://cgm.cs.mcgill.ca/avis/C/IRSLIB/USERGUIDE.html>). From a geometric perspective, each variable in a linear system is an axis of some high-dimensional coordinate system; consequently, eliminating a variable is equivalent to projecting the polytope onto a hyperplane orthogonal to that particular axis. As such, polytope projection is the titular topic of most of the relevant literature on linear quantifier elimination [48–52]. Furthermore, polytope projection and vertex enumeration are intimately related: One can define the vertex enumeration problem as a task of linear quantifier elimination, and therefore any polytope projection algorithm can be used to perform vertex enumeration, albeit less efficiently than specialized algorithms [43–47]. Conversely, a brute-force technique for polytope projection

Denoting the quantifier-free list of linear facet inequalities of the generalized-noncontextual polytope by $\{h^1, h^2, \dots, h^n\} \equiv \mathcal{H}$ (for “half spaces”) and letting the coefficients of a specific facet inequality h be given by $\gamma_{i,j,m}^h$ while γ_0^h indicates the constant term in that inequality, we find that:

Formulation F3 of the existence of a universally noncontextual model. For a data table $\{p(m|M_i, P_j)\}_{i,j,m}$, an ontological model that is universally noncontextual with respect to the operational equivalences in $\mathcal{O}\mathcal{E}_P$ and $\mathcal{O}\mathcal{E}_M$ exists if and only if

$$\forall h \in \mathcal{H}: \sum_{i,j,m} \gamma_{i,j,m}^h p(m|M_i, P_j) + \gamma_0^h \geq 0, \quad (23)$$

where \mathcal{H} is the set of n inequalities resulting from eliminating all free parameters $\{v_{P_j}(\kappa)\}_{j,\kappa}$ in the formulation of Eqs. (22a)–(22d).

V. DOES A GIVEN NUMERICAL DATA TABLE ADMIT OF A NONCONTEXTUAL MODEL?

To date, experimental tests of generalized noncontextuality [7,13] have targeted the *specific* preparations, measurements, and operational equivalences of some particular quantum no-go theorem [7,8,13,14,25,26]. By abstracting away the quantum-specific elements of the proof and describing the experiment in entirely operational terms, one can identify operational features of a set of preparations and measurements, such that any theory exhibiting those features fails to admit of a noncontextual ontological model. To test these operational features, previous experiments have used postprocessed data to enforce specific operational equivalences appearing in the quantum no-go argument. (See the “secondary procedures” technique described in Ref. [7].)

We here introduce a much more general analysis technique, in which one need not target any specific preparations, measurements, and operational equivalences. Rather, arbitrary numerical data tables can be directly analyzed. With respect to *whatever* operational equivalences happen to be manifest in the data (see footnote 2), one can use the methods we present below to efficiently test whether the numerical data table admits of a noncontextual model or not. Because answering this yes-no question does not require deriving the full set of noncontextuality inequalities for the scenario under study, it is computationally very efficient.⁷ Furthermore, analyzing data in this manner always allows for larger inequality violations,

is to first enumerate the polytope’s vertices, manually discard the to-be-eliminated coordinates from each vertex, and then revert back to inequalities using a convex hull algorithm. This roundabout method of performing polytope projection is generally suboptimal but can be used in practice.

⁷An analogous pair of problems with widely differing computational difficulties has long been appreciated in the study of Bell nonlocality. Obtaining *all* the Bell inequalities which characterize some nonlocality scenario can be quite difficult, but ascertaining if a particular correlation admits a local model or not can be resolved with the application of a single linear program [53,54]. The same (efficient) linear program can be used to return a *single* Bell inequality which certifies the nonlocality of the given correlation [55].

since the postprocessing required in the secondary procedures technique of Ref. [7] always introduces additional noise.

To test whether a numerically specified data table $\{p^*(m|M_i, P_j)\}_{i,j,m}$ admits of a noncontextual model, we leverage the formulation in Eqs. (22a)–(22d). All the equality constraints of Eqs. (22b), (22c), and (22d) can be encoded in a single matrix equality constraint,

$$\mathbf{M} \cdot \mathbf{x} = \mathbf{b}^*, \quad (24)$$

where \mathbf{M} contains the parameters $\alpha_{P_j}^s - \beta_{P_j}^s$ and the quantities $\{\xi_{m|M_i}(\kappa)\}_{i,m,\kappa}$, \mathbf{x} contains the unknown parameters $\{v_{P_j}(\kappa)\}_{j,\kappa}$, and \mathbf{b}^* contains the probabilities $\{p^*(m|M_i, P_j)\}_{i,j,m}$, as well as zeros and ones corresponding to the right-hand sides of Eqs. (22b) and (22c). Equation (22a) becomes simply $\mathbf{x} \geq 0$. Hence, for a numerically specified $\{p^*(m|M_i, P_j)\}_{i,j,m}$, the formulation of Eqs. (22a)–(22d) defines a linear program (LP).⁸

The primal LP is the search for a solution to a linear system of equations, namely:

$$\begin{aligned} &\exists \mathbf{x} \text{ such that} \\ &\mathbf{M} \cdot \mathbf{x} = \mathbf{b}^*, \\ &\text{and } \mathbf{x} \geq 0. \end{aligned} \quad (25)$$

Because no objective function to maximize or minimize is specified in the LP defined by Eq. (25), this means the LP is just checking for the existence of an \mathbf{x} which satisfies the constraints and hence guarantees the existence of a noncontextual model, via Eqs. (22a)–(22d).

Whenever the primal LP is infeasible—that is, no solution can be found—one can obtain a *certificate of primal infeasibility*, also known as the *Farkas dual* [56,57]. The certificate of primal infeasibility is obtained by solving the complementary⁹ linear system

$$\begin{aligned} &\min_{\mathbf{y}} \mathbf{y} \cdot \mathbf{b}^* \text{ such that} \\ &\mathbf{1} \geq \mathbf{y} \cdot \mathbf{M} \geq 0. \end{aligned} \quad (26)$$

Farkas’s lemma states that either the primal LP is feasible, or else the certificate \mathbf{y} resulting from Eq. (26) satisfies the strict inequality $\mathbf{y} \cdot \mathbf{b}^* < 0$.

Farkas’s lemma is easily proven: Plainly, if there exists such a \mathbf{y} (i.e., not only $\mathbf{y} \cdot \mathbf{M} \geq 0$ but also $\mathbf{y} \cdot \mathbf{b}^* < 0$), then there *cannot* exist an \mathbf{x} which satisfies the primal LP of Eq. (25),

⁸Linear programming is used across many fields; specialized LP software packages include MOSEK (<https://www.mosek.com/resources/downloads>), GUROBI (<http://www.gurobi.com/products/gurobi-optimizer>), and CPLEX (<https://www-01.ibm.com/software/commerce/optimization/linear-programming/>).

⁹The complementary LP defined in Eq. (26) is meant to explain how infeasibility certificates are generated in practice. Note, however, that the Farkas dual of an LP is not the same as the LP’s dual formulation, although the concepts are related. See Refs. [56,57], as well as Theorem 1 and Supplemental Materials of Ref. [39].

since the inequalities

$$\mathbf{x} \geq \mathbf{0}, \quad (27a)$$

$$\mathbf{y} \cdot \mathbf{M} \geq \mathbf{0}, \quad (27b)$$

$$\text{and } \mathbf{y} \cdot \mathbf{b}^* = \mathbf{y} \cdot \mathbf{M} \cdot \mathbf{x} < 0 \quad (27c)$$

cannot all be satisfied simultaneously.

Of relevance to this work is that we may interpret the certificate \mathbf{y} resulting from Eq. (26) as a noncontextuality inequality, since Farkas's lemma ensures that $\mathbf{y} \cdot \mathbf{b} \geq 0$ for every \mathbf{b} for which the primal LP is feasible. The extent to which $\mathbf{y} \cdot \mathbf{b}^*$ is negative is identically the amount by which the corresponding noncontextuality inequality is violated by the (contextual) $\{p^*(m|M_i, P_j)\}_{i,j,m}$. When interpreting certificates of primal infeasibility as noncontextuality inequalities, one deduces the constant term from those elements of \mathbf{b}^* which do not depend on $\{p^*(m|M_i, P_j)\}_{i,j,m}$. In practice, therefore, the constant term is the sum of those elements in \mathbf{y} which correspond to the normalization conditions of Eq. (22b).

This technique allows an experimenter to optimally certify the contextuality of a numerical data table $\{p^*(m|M_i, P_j)\}_{i,j,m}$ without first performing the computationally expensive task of finding all the noncontextuality inequalities, i.e., without doing any work to transform formulation F2 into formulation F3. Instead, by seeking a certificate of primal infeasibility—a single query to a linear program—one obtains *the* noncontextuality inequality which best witnesses the contextuality of the data table.

VI. APPLICATIONS

A. Generalized-noncontextual polytope in the simplest nontrivial case

As argued in Ref. [29], the simplest possible scenario in which the principle of noncontextuality implies nontrivial constraints on operational probabilities involves four preparations and two binary-outcome measurements.¹⁰ We imagine for simplicity that the preparations satisfy the operational equivalence

$$\frac{1}{2}P_1 + \frac{1}{2}P_2 \simeq \frac{1}{2}P_3 + \frac{1}{2}P_4 \quad (28)$$

and that there are no operational equivalences among the measurements.

We denote the operational probability $p(0|M_i, P_j)$ by p_{ij} . [By normalization, probability $p(1|M_i, P_j)$ is then $1 - p_{ij}$.] Vertex enumeration finds four vertices for the non-contextual measurement-assignment polytope, corresponding to the four deterministic assignments $(\xi_{0|M_1}(\lambda), \xi_{0|M_2}(\lambda)) \in \{(0,0), (0,1), (1,0), (1,1)\}$. Each of the four preparations defines a probability distribution over these four ontic states, so there are 16 free parameters to be eliminated. Linear quantifier elimination finds the polytope of data tables consistent with the principle of noncontextuality and the operational equivalence

of Eq. (28) to be

$$\forall i, j: \quad 0 \leq p_{ij} \leq 1, \quad (29a)$$

$$p_{12} + p_{22} - p_{23} - p_{14} \leq 1, \quad (29b)$$

$$p_{12} + p_{22} - p_{13} - p_{24} \leq 1, \quad (29c)$$

$$p_{22} + p_{13} - p_{12} - p_{24} \leq 1, \quad (29d)$$

$$p_{12} + p_{23} - p_{22} - p_{14} \leq 1, \quad (29e)$$

$$p_{22} + p_{14} - p_{12} - p_{23} \leq 1, \quad (29f)$$

$$p_{23} + p_{14} - p_{21} - p_{22} \leq 1, \quad (29g)$$

$$p_{12} + p_{24} - p_{22} - p_{13} \leq 1, \quad (29h)$$

$$p_{13} + p_{24} - p_{12} - p_{22} \leq 1. \quad (29i)$$

Note that the two probabilities which do not appear in Eqs. (29b)–(29i), p_{11} and p_{21} , are fixed by the operational equivalence relation, Eq. (28):

$$p_{11} = p_{13} + p_{14} - p_{12}, \quad (30a)$$

$$p_{21} = p_{23} + p_{24} - p_{22}. \quad (30b)$$

Inequalities (29) and Eqs. (30) tightly define the generalized-noncontextual polytope for this scenario. Furthermore, all of Ineqs. (29) are equivalent under relabeling. That is, any one of the inequalities can generate all eight by applying relabelings which respect the operational equivalences: $M_1 \leftrightarrow M_2$, $P_1 \leftrightarrow P_2$, and $(P_1, P_2) \leftrightarrow (P_3, P_4)$. Similarly, each of Eqs. (30) is equivalent to the other under the same relabelings.

As an illustration of how a noncontextuality inequality can be derived from a numerically specified (contextual) data table, consider the following example:

$$\begin{aligned} p_{11} = 1, \quad p_{12} = 0, \quad p_{13} = 1, \quad p_{14} = 0, \\ p_{21} = 1, \quad p_{22} = 0, \quad p_{23} = 0, \quad p_{24} = 1, \end{aligned} \quad (31)$$

which respects the operational equivalence relation of Eq. (28), but maximally violates Ineq. (29i), since it has $p_{13} + p_{24} - p_{12} - p_{22} = 2 \not\leq 1$. Indeed, when we construct the primal linear program per Eq. (25), we find it to be infeasible, and we find that the certificate of infeasibility returned by our numerical solver corresponds to Ineq. (29i).

1. Relevance to parity-oblivious multiplexing

In the communication task of “parity-oblivious multiplexing,” an agent Alice wishes to communicate two bits to an agent Bob, in such a way that Bob can extract information about either of the two bits but cannot extract any information about their parity [13]. This task involves four preparations (associated to the four possibilities for the values of the two bits) and two measurements (corresponding to which bit Bob wishes to learn about), and the parity-obliviousness condition implies an operational equivalence relation among the preparations, namely, that of Eq. (28). Consequently, this task fits precisely the operational scenario considered in this section.

In Ref. [13], it was shown that contextuality provides an advantage for the task of parity-oblivious multiplexing: the maximum probability of succeeding at this task in a noncontextual model is $3/4$, so that any higher probability

¹⁰Reference [29] also assumes that these two measurements are tomographically complete, but we do not make this assumption here. See Refs. [7,29] for details on the issue of tomographic completeness.

of success requires contextuality. In a quantum world, for instance, one can succeed with probability $\frac{1}{4}(2 + \sqrt{2}) \approx 0.85$.

If one identifies our preparations P_1, P_2, P_3 , and P_4 with preparations P_{00}, P_{11}, P_{01} , and P_{10} , respectively, of Ref. [13], then the inequality in Ref. [13] is our facet Ineq. (29i). This is the same inequality which witnesses the failure of noncontextuality for the data table defined in Eq. (31), which is to be expected, as this particular data table describes a set of probabilities which achieve the maximum logically possible probability of success at parity-oblivious multiplexing.

Additionally, one could apply our algorithm to generalized types of parity-oblivious multiplexing. For instance, Ref. [13] derived a bound on the probability of success in a noncontextual model of n -bit parity-oblivious multiplexing, in which Alice wishes to communicate n bits to Bob under the constraint that Bob can learn no information about the parity between any two of the bits. These bounds are tight, and can be saturated by a naïve classical strategy. However, one could further use our techniques to learn whether or not they constitute facet inequalities of the generalized-noncontextual polytope, as well as to find the full generalized-noncontextual polytope.

For the still more general case where the different n -bit strings which Alice wishes to send do not have equal *a priori* probabilities, our method can also find the generalized-noncontextual polytope, from which one can immediately infer the maximum success probability for the task.

B. Generalized-noncontextual polytopes for scenarios relevant to state discrimination

One way in which one can generalize the simplest operational scenario described above is to increase the number of binary-outcome measurements from two to three while still not assuming any operational equivalences among them, so that the only operational equivalence relation remains the one between the preparations [Eq. (28)]. This operational scenario can also be related to an information-theoretic task, namely, the task of minimum error state discrimination, as noted by two of the present authors in Ref. [58].

In the quantum version of this task, an agent wishes to guess which of two pure quantum states a system was prepared in given a single sample of the system, where the identity of the two quantum states is known. Quantum theory prescribes a particular trade-off relation between the probability of success and the nonorthogonality of the two quantum states, and Ref. [58] showed that this trade-off contradicts the principle of generalized noncontextuality. The ideal quantum realization of minimum error state discrimination fits the operational scenario described above: the two pure quantum states define two of the preparation procedures, while their orthogonal complements in the two-dimensional (2D) subspace that they span define the other two. The fact that the equal mixture of any two orthogonal pure states in a 2D subspace is independent of the basis implies the operational equivalence of Eq. (28). Finally, the degree of nonorthogonality has an operational interpretation as the probability of one state passing a test for the other (termed the *confusability*). Therefore, the measurements of each of the two bases, together with the discriminating measurement, provide the three binary-outcome measurements in the scenario.

The facet noncontextuality inequalities for this operational scenario are given in Appendix D of Ref. [58],¹¹ and these are seen to imply a nontrivial upper bound on the probability of successful discrimination for a given confusability. Hence, contextuality provides an advantage for minimum error state discrimination. The quantum probability of successful discrimination for a given confusability is higher than that allowed in a noncontextual model and hence partakes in this contextual advantage.

Using our technique, one can also immediately derive the generalized-noncontextual polytope for more general minimum error state discrimination scenarios, such as those in which the quantum states (preparations) being discriminated are sampled with unequal probabilities, or in which there are more than two quantum states (preparations).

Because state discrimination in various forms is a primitive for many other quantum information-processing tasks, such analyses should be valuable for identifying the circumstances in which contextuality constitutes a resource.

C. Generalized-noncontextual polytopes for a scenario involving both preparation and measurement noncontextuality

So far, our examples have involved operational equivalences only among preparations. In this section, we revisit the scenario considered in the recent experimental test of noncontextuality in Ref. [7], which involves operational equivalences among the preparations and also among the measurements. Specifically, we imagine a set of six preparations and three binary-outcome measurements, where the preparations satisfy the operational equivalences

$$\frac{1}{2}P_1 + \frac{1}{2}P_2 \simeq \frac{1}{2}P_3 + \frac{1}{2}P_4 \simeq \frac{1}{2}P_5 + \frac{1}{2}P_6, \quad (32)$$

and the measurement effects satisfy the operational equivalence:

$$\begin{aligned} & \frac{1}{3}[0|M_1] + \frac{1}{3}[0|M_2] + \frac{1}{3}[0|M_3] \\ & \simeq \frac{1}{3}[1|M_1] + \frac{1}{3}[1|M_2] + \frac{1}{3}[1|M_3]. \end{aligned} \quad (33)$$

(See Ref. [7] for a discussion of the significance of these operational equivalences.)

We denote the operational probability $p(0|M_i, P_j)$ by p_{ij} and $p(1|M_i, P_j)$ by \bar{p}_{ij} . Vertex enumeration finds six vertices for the noncontextual measurement-assignment polytope, corresponding to the four *indeterministic* assignments defined by $(\xi_{0|M_1}(\lambda), \xi_{0|M_2}(\lambda), \xi_{0|M_3}(\lambda)) \in \{(0, \frac{1}{2}, 1), (\frac{1}{2}, 0, 1), (1, 0, \frac{1}{2}), (1, \frac{1}{2}, 0), (0, 1, \frac{1}{2}), (\frac{1}{2}, 1, 0)\}$. Each of the six preparations defines a probability distribution over these six ontic states, so there are 36 free parameters to be eliminated. Linear quantifier elimination finds the polytope of data tables consistent with the principle of noncontextuality and with the operational equivalences of Eqs. (32) and (33). We find that this polytope has 1596 facet inequalities.

Plainly, 1596 inequalities is far too many to list explicitly. However, by considering the physical symmetries of this

¹¹Actually, the polytope given therein is the intersection of the generalized-noncontextual polytope with two additional inequalities, which are implied by making sensible labeling choices.

scenario, we can significantly simplify our description of these facets. Since the scenario is invariant under various relabelings of measurements [Eqs. (34a)–(34b)], outcomes [Eq. (34c)], and preparations [Eqs. (34d)–(34f)]—i.e., those relabelings which respect the operational equivalences—we know *a priori* that the generalized-noncontextual polytope will possess significant internal symmetry. The symmetry group which leaves our polytope invariant is generated by the six relabelings:

$$M_1 \leftrightarrow M_2, \tag{34a}$$

$$M_1 \leftrightarrow M_3, \tag{34b}$$

$$([0|M_1],[0|M_2],[0|M_3]) \leftrightarrow ([1|M_1],[1|M_2],[1|M_3]), \tag{34c}$$

$$P_1 \leftrightarrow P_2, \tag{34d}$$

$$(P_1, P_2) \leftrightarrow (P_3, P_4), \tag{34e}$$

$$(P_1, P_2) \leftrightarrow (P_5, P_6). \tag{34f}$$

We use parentheses to indicate a coherent relabeling: e.g., the outcomes of three measurements can be flipped per Eq. (34c), but only if all three measurements have their outcomes relabeled simultaneously. This is in contrast to an exchange like $P_1 \leftrightarrow P_2$ per Eq. (34d), which can be performed in isolation. The total order of this symmetry group is 576.

Under this group, we find that the 1596 facet inequalities admit classification into seven symmetry classes. We therefore explicitly list a single representative inequality from each class:

Inequality terms	Upper bound	Orbit size
p_{11}	≤ 1	35
$p_{11} + p_{23} + p_{35}$	≤ 2.5	48
$p_{11} + p_{22} + p_{35}$	≤ 2.5	72
$p_{11} - p_{14} - 2 p_{15} - 2 p_{22} + 2 p_{23} + 2 p_{35}$	≤ 3	576
$2 p_{11} - p_{22} + 2 p_{23}$	≤ 3	144
$p_{11} - p_{15} + p_{22} + p_{23} + 2 p_{35}$	≤ 4	576
$p_{11} - p_{15} + 2 p_{22} + 2 p_{35}$	≤ 4	144

(35)

The number of inequalities in each symmetry class is given by the ‘‘Orbit size’’ in Ineqs. (35). The generalized-noncontextual polytope for this scenario is defined by the 1596 facet inequalities, as well as by equalities which hold for *any* data table (contextual or noncontextual) admitting the operational equivalence relations per Eqs. (32) and (33). These equalities fall into three distinct symmetry classes, represented by

$$p_{11} + p_{14} = p_{12} + p_{15}, \tag{36a}$$

$$p_{11} + p_{21} + p_{31} = 3/2, \tag{36b}$$

$$p_{11} + \bar{p}_{11} = 1. \tag{36c}$$

The first two equalities are enforced by the operational equivalence relations, Eqs. (32) and (33), respectively, while the third equality is guaranteed by normalization of measurements, Eq. (6).

To test if a given data table lies inside this generalized-noncontextual polytope, one could reconstruct all 1596 inequalities from the seven given in Ineqs. (35), but it is likely much easier to instead artificially generate equivalent-up-to-symmetries data tables from one’s actual data table, and then to test each of *those* against the seven canonical inequalities.¹² One only needs to consider at most 576 data table variants (per the group order), although in practice there will be fewer variants to consider if the data table one wishes to investigate possesses any internal symmetry of its own.

Noting that¹³

$$p_{11} + p_{23} + p_{35} = \bar{p}_{12} + \bar{p}_{24} + \bar{p}_{36}, \tag{37}$$

the single inequality derived (and experimentally violated) in Eq. (6) of Ref. [7] is recognized as a facet of the generalized-noncontextual polytope; namely, it is precisely the second inequality in Ineqs. (35):

$$2(p_{11} + p_{23} + p_{35}) \leq 5. \tag{38}$$

One can also derive Ineq. (38) directly (and efficiently!) by using the linear program presented in Sec. V. Namely, it is the inequality corresponding to the certificate of infeasibility returned by our numerical solver, when we construct the primal linear program using Eq. (25) together with the ideal (contextual) quantum data table.

VII. CONNECTIONS TO BELL POLYTOPES AND TO POLYTOPES OF GENERAL PROBABILISTIC MODELS IN KOCHEN-SPECKER CONTEXTUALITY SCENARIOS

The central object of interest in prepare-and-measure contextuality scenarios are data tables, in particular the polytope of noncontextual data tables discussed herein. Along the way to finding this polytope, we have found it useful to first compute the noncontextual measurement-assignment polytope. It turns out that any Bell polytope is directly connected to our generalized-noncontextual polytope for a corresponding prepare-and-measure scenario, and that the polytope of generalized probabilistic models of Ref. [59] (and instances thereof, such as the polytope of ‘no-disturbance correlations’ defined in Ref. [60]) are special cases of our noncontextual measurement-assignment polytope.

In the literature on Bell inequalities, the set of allowed classical correlations is known as the *Bell polytope* (or equivalently the local polytope) [61,62]. As discussed in Ref. [27], any Bell scenario can be analyzed as a prepare-and-measure contextuality experiment. In brief, one can always reimagine the measurements of one of the parties in the Bell scenario as a remote preparation of the distant party’s state (via quantum steering). Under this mapping, the no-signaling condition in the Bell scenario implies an operational equivalence in the contextuality scenario: that the average state of the steered system is the same for any choice of steering measurement.

¹²The first inequality in Ineqs. (35) is inviolable, even by contextual data tables, so technically one only needs to test against the remaining six inequalities.

¹³Equation (37) is a nontrivial but readily verifiable consequence of the inviolable equalities in Eqs. (36).

Further, under this mapping the assumption of local causality implies the assumption of preparation noncontextuality for this operational equivalence relation. Then one finds that our generalized-noncontextual polytope is directly connected to the Bell polytope: specifically, the latter is generated from the former by multiplying each prepare-and-measure probability by the probability of steering to the corresponding remote preparation.

For example, our contextuality scenario in Sec. VIA is isomorphic to the Clauser-Horne-Shimony-Holt Bell scenario [63] (with two parties, each with two binary outcome measurements), as discussed in Ref. [13]. Additionally, our contextuality scenario in Sec. VIB is isomorphic to a Bell scenario (with two parties, one with two binary measurements and another with two ternary measurements), as discussed in Ref. [58].

In this way, all Bell scenarios are special cases of our framework. However, our framework is much more general. As stated in Ref. [58], the only prepare-and-measure contextuality scenarios which are equivalent to a related Bell scenario are those which do not leverage any operational equivalences among the measurements, and which leverage only those operational equivalences among the preparations which arise from the various ensemble decompositions of a single mixed preparation.

There is also a connection between our work and that of Ref. [59], which gives an operational framework for studying Kochen-Specker contextuality scenarios. Reference [59] considers an arbitrary set of measurements, some of which have one or more measurement effects in common, and characterizes various classes of probability assignments to those measurement effects (i.e., various classes of what we have here called “measurement assignments”). They term these classes of assignments “classical,” “quantum,” “probabilistic,” etc. Note, however, that the class of assignments termed “classical” in Ref. [59] does *not* correspond to the noncontextual measurement-assignment polytope in our approach, since Ref. [59] assumes outcome determinism, while in our approach *indeterministic* probability assignments are allowed. Rather, it is the polytope of *general* probabilistic models in Ref. [59] that corresponds to our noncontextual measurement-assignment polytope because the definition of a general probabilistic model for the scenarios considered in Ref. [59] is precisely our definition of a measurement-noncontextual assignment (which, we recall, can be indeterministic).

The two polytopes coincide exactly in every scenario to which the framework of Ref. [59] applies, since in that case the operational equivalences among the measurements are

precisely those given by the occurrences of some effect in more than one measurement. However, our noncontextual measurement-assignment polytope is also defined for scenarios with other types of operational equivalences among the measurements, where the framework of Ref. [59] does not apply.

VIII. CONCLUSIONS

For arbitrary prepare-and-measure experiments, we have presented a method for finding necessary and sufficient conditions for a data table to admit of a noncontextual model, subject to any fixed sets of operational equivalences among preparations and among measurements. We have also presented an efficient method for determining whether a *numerical* data table is noncontextual, in this same setting.

We have provided worked examples of each of these methods, in the process deriving necessary and sufficient conditions for operational scenarios in which only necessary conditions were previously known. Equivalently, we have derived the full generalized-noncontextual polytopes for scenarios in which only a single facet inequality was previously known. The operational scenario studied in Sec. VIA is of relevance to parity-oblivious multiplexing [13], while the operational scenario studied in Sec. VIC originates in a recent experimental test of contextuality [7].

A precursor to the current work can be found in Ref. [37]. A distinct method, introduced in Ref. [25], also allows one to derive all of the facet inequalities for many operational scenarios. However, the method therein is not fully general, as it applies only to scenarios in which one special equivalence class of preparations is singled out (see Sec. III B of Ref. [25] for details). It would be interesting to compare the two approaches, e.g., in terms of computational efficiency, and to modify the approach in Ref. [25] to make it as general as the approach described in this article.

ACKNOWLEDGMENTS

D.S. thanks Ravi Kunjwal for useful discussions. This research was supported by a Discovery grant of the Natural Sciences and Engineering Research Council of Canada and by the Perimeter Institute for Theoretical Physics. Research at the 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] S. Kochen and E. Specker, The problem of hidden variables in quantum mechanics, *J. Math. Mech.* **17**, 59 (1967), also available from Indiana Univ. Math. J. (<http://www.iuj.indiana.edu/IUMJ/fulltext.php?year=1968&volume=17&artid=17004>).
 - [2] R. W. Spekkens, Contextuality for preparations, transformations, and unsharp measurements, *Phys. Rev. A* **71**, 052108 (2005).
 - [3] A. Cabello and G. García-Alcaine, Proposed Experimental Tests of the Bell-Kochen-Specker Theorem, *Phys. Rev. Lett.* **80**, 1797 (1998).
 - [4] A. Cabello, Experimentally Testable State-Independent Quantum Contextuality, *Phys. Rev. Lett.* **101**, 210401 (2008).
 - [5] A. A. Klyachko, M. A. Can, S. Binicioğlu, and A. S. Shumovsky, Simple Test for Hidden Variables in Spin-1 Systems, *Phys. Rev. Lett.* **101**, 020403 (2008).
 - [6] R. W. Spekkens, The status of determinism in proofs of the impossibility of a noncontextual model of quantum theory, *Found. Phys.* **44**, 1125 (2014).
 - [7] 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**, 11780 (2016).

- [8] R. Kunjwal and R. W. Spekkens, From the Kochen-Specker Theorem to Noncontextuality Inequalities Without Assuming Determinism, *Phys. Rev. Lett.* **115**, 110403 (2015).
- [9] N. Brunner, D. Cavalcanti, S. Pironio, V. Scarani, and S. Wehner, Bell nonlocality, *Rev. Mod. Phys.* **86**, 419 (2014).
- [10] R. W. Spekkens, Negativity and Contextuality are Equivalent Notions of Nonclassicality, *Phys. Rev. Lett.* **101**, 020401 (2008).
- [11] M. F. Pusey, Anomalous Weak Values are Proofs of Contextuality, *Phys. Rev. Lett.* **113**, 200401 (2014).
- [12] N. D. Mermin, Hidden variables and the two theorems of John Bell, *Rev. Mod. Phys.* **65**, 803 (1993).
- [13] 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).
- [14] 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).
- [15] R. Kunjwal, Beyond the Cabello-Severini-Winter framework: Making sense of contextuality without sharpness of measurements, [arXiv:1709.01098](https://arxiv.org/abs/1709.01098) [quant-ph].
- [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, [arXiv:1607.05490](https://arxiv.org/abs/1607.05490).
- [18] D. Saha, P. Horodecki, and M. Pawłowski, State independent contextuality advances one-way communication, [arXiv:1708.04751](https://arxiv.org/abs/1708.04751) [quant-ph].
- [19] M. Howard, J. Wallman, V. Veitch, and J. Emerson, Contextuality supplies the ‘magic’ for quantum computation, *Nature (London)* **510**, 351 (2014).
- [20] R. Raussendorf, Contextuality in measurement-based quantum computation, *Phys. Rev. A* **88**, 022322 (2013).
- [21] J. Anders and D. E. Browne, Computational Power of Correlations, *Phys. Rev. Lett.* **102**, 050502 (2009).
- [22] A. Acín and L. Masanes, Certified randomness in quantum physics, *Nature (London)* **540**, 213 (2016).
- [23] J. Barrett, L. Hardy, and A. Kent, No Signaling and Quantum Key Distribution, *Phys. Rev. Lett.* **95**, 010503 (2005).
- [24] A. Acín, N. Gisin, and L. Masanes, From Bell’s Theorem to Secure Quantum Key Distribution, *Phys. Rev. Lett.* **97**, 120405 (2006).
- [25] A. Krishna, R. W. Spekkens, and E. Wolfe, Deriving robust non-contextuality inequalities from algebraic proofs of the Kochen-Specker theorem: The Peres-Mermin square, *New J. Phys.* **19**, 123031 (2017).
- [26] Z.-P. Xu, D. Saha, H.-Y. Su, M. Pawłowski, and J.-L. Chen, Reformulating noncontextuality inequalities in an operational approach, *Phys. Rev. A* **94**, 062103 (2016).
- [27] Y.-C. Liang, R. W. Spekkens, and H. M. Wiseman, Specker’s parable of the overprotective seer: A road to contextuality, nonlocality and complementarity, *Phys. Rep.* **506**, 1 (2011).
- [28] R. Kunjwal, Fine’s theorem, noncontextuality, and correlations in Specker’s scenario, *Phys. Rev. A* **91**, 022108 (2015).
- [29] M. F. Pusey, The robust noncontextuality inequalities in the simplest scenario, [arXiv:1506.04178](https://arxiv.org/abs/1506.04178).
- [30] D. Collins and N. Gisin, A relevant two qubit Bell inequality inequivalent to the chsh inequality, *J. Phys. A* **37**, 1775 (2004).
- [31] N. Brunner and N. Gisin, Partial list of bipartite Bell inequalities with four binary settings, *Phys. Lett. A* **372**, 3162 (2008).
- [32] J.-D. Bancal, N. Gisin, and S. Pironio, Looking for symmetric Bell inequalities, *J. Phys. A* **43**, 385303 (2010).
- [33] C. Budroni and A. Cabello, Bell inequalities from variable-elimination methods, *J. Phys. A* **45**, 385304 (2012).
- [34] T. Fritz and R. Chaves, Entropic inequalities and marginal problems, *IEEE Trans. Info. Theor.* **59**, 803 (2013).
- [35] R. Chaves, L. Luft, and D. Gross, Causal structures from entropic information: Geometry and novel scenarios, *New J. Phys.* **16**, 043001 (2014).
- [36] R. Chaves, C. Majenz, and D. Gross, Information-theoretic implications of quantum causal structures, *Nat. Commun.* **6**, 5766 (2015).
- [37] A. Krishna, Experimentally testable noncontextuality inequalities via Fourier-Motzkin elimination, Master’s thesis, University of Waterloo, 2015.
- [38] E. Wolfe, R. W. Spekkens, and T. Fritz, The inflation technique for causal inference with latent variables, [arXiv:1609.00672](https://arxiv.org/abs/1609.00672).
- [39] S. Abramsky, R. S. Barbosa, and S. Mansfield, Contextual Fraction as a Measure of Contextuality, *Phys. Rev. Lett.* **119**, 050504 (2017).
- [40] S. Pironio, J.-D. Bancal, and V. Scarani, Extremal correlations of the tripartite no-signaling polytope, *J. Phys. A* **44**, 065303 (2011).
- [41] C. Ferrie and J. Emerson, Frame representations of quantum mechanics and the necessity of negativity in quasi-probability representations, *J. Phys. A* **41**, 352001 (2008).
- [42] C. Ferrie, Quasi-probability representations of quantum theory with applications to quantum information science, *Rep. Prog. Phys.* **74**, 116001 (2011).
- [43] D. Avis, A revised implementation of the reverse search vertex enumeration algorithm, in *Polytopes—Combinatorics and Computation*, DMV Seminar (Birkhäuser, Basel, 2000), Vol. 29, pp. 177–198.
- [44] D. Bremner, M. Dutour Sikirić, and A. Schürmann, Polyhedral representation conversion up to symmetries, in *Polyhedral Computation*, CRM Proc. Lecture Notes (American Mathematical Society, Providence, RI, 2009), Vol. 48, pp. 45–71.
- [45] N. Y. Zolotykh, New modification of the double description method for constructing the skeleton of a polyhedral cone, *Commun. Math. Phys.* **52**, 146 (2012).
- [46] D. Avis, D. Bremner, and R. Seidel, How good are convex Hull algorithms? *Comp. Geom.* **7**, 265 (1997).
- [47] S. Lörwald and G. Reinelt, PANDA: A software for polyhedral transformations, *Eur. J. Comp. Optim.* **3**, 297 (2015).
- [48] G. B. Dantzig and B. C. Eaves, Fourier-Motzkin elimination and its dual, *J. Comb. Theory A* **14**, 288 (1973).
- [49] E. Balas, Projection with a minimal system of inequalities, *Comp. Optimiz. Appl.* **10**, 189 (1998).
- [50] C. Jones, E. C. Kerrigan, and J. Maciejowski, Equality set projection: A new algorithm for the projection of polytopes

- in halfspace representation, Tech. Rep., Cambridge University Engineering Department (unpublished).
- [51] D. V. Shapot and A. M. Lukatskii, Solution building for arbitrary system of linear inequalities in an explicit form, *Am. J. Comp. Math.* **2**, 1 (2012).
- [52] S. I. Bastrakov and N. Y. Zolotykh, Fast method for verifying Chernikov rules in Fourier-Motzkin elimination, *Commun. Math. Phys.* **55**, 160 (2015).
- [53] M. Zukowski, D. Kaszlikowski, A. Baturó, and J.-A. Larsson, Strengthening the Bell Theorem: Conditions to falsify local realism in an experiment, [arXiv:quant-ph/9910058](https://arxiv.org/abs/quant-ph/9910058).
- [54] R. M. Basoalto and I. C. Percival, A general computer program for the Bell detection loophole, *Phys. Lett. A* **280**, 1 (2001).
- [55] M. B. Elliott, A linear program for testing local realism, [arXiv:0905.2950](https://arxiv.org/abs/0905.2950).
- [56] E. D. Andersen, Certificates of primal or dual infeasibility in linear programming, *Comp. Optim. Appl.* **20**, 171 (2001).
- [57] D. G. Luenberger and Y. Ye, Duality and complementarity, in *Linear and Nonlinear Programming* (Springer, Cham, 2016), pp. 83–114.
- [58] D. Schmid and R. W. Spekkens, Contextual Advantage for State Discrimination, *Phys. Rev. X* **8**, 011015 (2018).
- [59] A. Acín, T. Fritz, A. Leverrier, and A. B. Sainz, A combinatorial approach to nonlocality and contextuality, *Commun. Math. Phys.* **334**, 533 (2015).
- [60] M. Araújo, M. T. Quintino, C. Budroni, M. T. Cunha, and A. Cabello, All noncontextuality inequalities for the n -cycle scenario, *Phys. Rev. A* **88**, 022118 (2013).
- [61] A. Fine, Hidden Variables, Joint Probability, and the Bell Inequalities, *Phys. Rev. Lett.* **48**, 291 (1982).
- [62] I. Pitowsky, *Quantum Probability Quantum Logic*, Lecture Notes in Physics (Springer, Berlin, 2014).
- [63] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, Proposed Experiment to Test Local Hidden-Variable Theories, *Phys. Rev. Lett.* **23**, 880 (1969).