From statistical proofs of the Kochen-Specker theorem to noise-robust noncontextuality inequalities

Ravi Kunjwal[*] and Robert W. Spekkens[†]

Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, Ontario, Canada N2L 2Y5

The Kochen-Specker theorem rules out models of quantum theory wherein projective measurements are assigned outcomes deterministically and independently of context. This notion of noncontextuality is not applicable to experimental measurements because these are never free of noise and thus never truly projective. For nonprojective measurements, therefore, one must drop the requirement that an outcome be assigned deterministically in the model and merely require that it be assigned a *distribution over outcomes* in a manner that is context-independent. By demanding context independence in the representation of preparations as well, one obtains a generalized principle of noncontextuality that also supports a quantum no-go theorem. Several recent works have shown how to derive inequalities on experimental data which, if violated, demonstrate the impossibility of finding a generalized-noncontextual model of this data. That is, these inequalities do not presume quantum theory and, in particular, they make sense without requiring an operational analog of the quantum notion of projectiveness. We here describe a technique for deriving such inequalities starting from arbitrary proofs of the Kochen-Specker theorem. It extends significantly previous techniques that worked only for logical proofs, which are based on sets of projective measurements that fail to admit of any deterministic noncontextual assignment, to the case of statistical proofs, which are based on sets of projective measurements that *do* admit of some deterministic noncontextual assignments, but not enough to explain the quantum statistics.

I. INTRODUCTION

In quantum theory, a given sharp measurement (i.e., one associated with a projector-valued measure) can be compatible with each of two other sharp measurements that are incompatible with one another. In this case, the latter pair of measurements define two distinct compatibility contexts for the first measurement. The Kochen-Specker (KS) theorem [1] rules out a particular kind of explanation of the operational predictions of quantum theory, namely, one wherein sharp measurements are assigned outcomes deterministically and independently of context. We will term this sort of explanation a *KS-noncontextual model* of the statistics.

A particular proof of the KS theorem is said to be *logical* if it appeals to the existence of a set of sharp quantum measurements that admit of *no* deterministic noncontextual assignments. (For the case where all measurement outcomes correspond to rank-1 projectors, i.e., projectors onto rays of Hilbert space, a deterministic noncontextual assignment is one that assigns, for every basis of orthogonal rays, the value 1 to precisely one such ray and the value 0 to the others, termed a *KS coloring* of the rays. In this case, a proof is said to be logical if it appeals to the existence of a set of rays in a Hilbert space that admit of *no KS colorings* [1–4].) A proof is said to be *statistical* if it appeals to a set that *does* admit of some deterministic noncontextual assignments, but these assignments are insufficient to explain the statistics of the measurements on one or more quantum states. Explaining the statistics of the measurements here means recovering them as

a convex mixture of deterministic noncontextual assignments to these measurements.

As an aside, we note that in most of the literature on the subject (see, e.g., [7]), a proof of the KS theorem is said to be *state-independent* if the proof works for an arbitrary choice of quantum state on which the measurements are implemented, and it is said to be *state-dependent* if the proof appeals to a special quantum state. In our terminology, *logical* proofs of the KS theorem are always state-independent. The reason is that logical proofs by definition appeal to sets of sharp measurements that admit of *no* deterministic noncontextual assignments (for example, those that appeal to sets of rays that admit of no KS colorings [1,4]), and if the set of deterministic noncontextual assignments is empty, then the set of mixtures of deterministic noncontextual assignments is also empty, and we cannot explain measurement statistics in terms of such a mixture regardless of the quantum state. *Statistical* proofs of the KS theorem, on the other hand, are those that appeal to sets of sharp measurements that admit of one or more deterministic noncontextual assignments (for example, sets of rays that admit of one or more KS colorings). These can be state-dependent, such as the proof in Ref. [5], or state-independent, such as the proof in Refs. [6,7].

In recent years, there has been much work deriving inequalities for operational statistics that follow directly from an assumption of noncontextuality, without assuming the validity of quantum theory. If these are violated experimentally, one can conclude that not just quantum theory, but any successor thereof must fail to admit of a noncontextual model.[1] One

---

*rkunjwal@perimeterinstitute.ca
†rspekkens@perimeterinstitute.ca

---

[1]The difference between a quantum no-go theorem for noncontextuality and a noncontextuality inequality is precisely analogous to

approach to deriving these inequalities is to seek inspiration from particular proofs of the KS theorem. References [10] and [11] do so for logical proofs. Here, we consider statistical proofs.

Section II reviews background material, in particular, the notions of operational theories, ontological models, measurement noncontextuality, and Kochen-Specker (KS) noncontextuality. For pedagogical clarity, we present our result as a generalization of the results obtained previously for logical proofs [10], but cast into a slightly different form. In Sec. III, therefore, we recast these earlier results. Along the way, we review the notion of preparation noncontextuality and how it can be used to infer that certain measurements are assigned outcomes deterministically in the ontological model if certain operational correlations hold. We also review how this inference allows for a proof of the failure of KS noncontextuality to be translated into a proof of the failure of preparation and measurement noncontextuality. Then, in Sec. IV, we generalize to the case of statistical proofs. Our main result is presented in Theorem 1 and we give some concrete examples of its application. We close with a discussion in Sec. V.

Xu *et al.* [12] have previously obtained noise-robust noncontextuality inequalities for some statistical proofs of the KS theorem. Our approach in this article is distinct from that of Ref. [12] and allows a consideration of arbitrary statistical proofs of the KS theorem, rather than particular examples of it. A discussion of the relation between Ref. [12] and this article is provided in Sec. V and Appendix E.

II. PRELIMINARIES

We conceptualize an experiment as a source followed by a measurement. A source is a procedure that samples a random variable and implements a preparation on the system conditioned on the value of this variable. The value of the variable is termed the *outcome* of the source. The event of obtaining an outcome s of a source S will be termed a *source event*, denoted $[s|S]$. For each source event $[s|S]$, one can associate a preparation procedure on the system; it is the one implemented by the source S conditional on outcome s occurring. The event of obtaining any outcome in some subset V of all possible outcomes of a source S will be denoted by $[V|S]$, and is associated with the preparation procedure wherein one implements S, coarse-grains over the outcomes in V, and then conditions on obtaining this coarse-grained outcome. We denote the full set of outcomes by $\top$, so that the preparation procedure corresponding to implementing S and *not* conditioning on obtaining any particular outcome is denoted $[\top|S]$.[2] A measurement M has outcome denoted by

m, and the event of obtaining outcome m of measurement M is termed a *measurement event*, denoted $[m|M]$. The measurement event $[V|M]$ for a subset V of the outcomes is defined in the obvious manner, similarly to the case of sources.

An *operational theory* specifies a rule for assigning a joint probability distribution $p(m,s|M,S)$, denoting the probability that in a prepare-and-measure experiment with source S and measurement M, the source event $[s|S]$ occurs followed by the measurement event $[m|M]$. For example, when the operational theory is quantum theory, the source event $[s|S]$ is represented by some density matrix, say $\rho_{[s|S]}$, the measurement event $[m|M]$ is represented by a positive operator, say $E_{[m|M]}$, and the rule for assigning the joint probability is $p(m,s|M,S) = p(s|S)\mathrm{Tr}(E_{[m|M]}\rho_{[s|S]})$, where $p(s|S)$ is the probability that the source S yields outcome s. Note that our analysis in this paper does not depend on the particular representation of preparations and measurements in an operational theory, nor on the particular probability rule associated with the theory; in this sense, we consider operational theories more general than quantum theory.

An *ontological model* of an operational theory posits that the causal influence of the source on the measurement is mediated by the ontic state, λ, of the system (a point in the underlying ontic state space Λ, which for our purposes can be taken to be discrete). For a source S with outcome s, the ontological model associates a conditional probability $\mu(\lambda,s|S)$ such that $\sum_s \sum_{\lambda \in \Lambda} \mu(\lambda,s|S) = 1$. Here $\mu(\lambda,s|S) = \mu(\lambda|s,S)p(s|S)$. For a measurement M with outcome m, the ontological model associates a conditional probability $\xi(m|M,\lambda)$ such that $\sum_m \xi(m|M,\lambda) = 1$ for all $\lambda \in \Lambda$. Finally, the ontological model must reproduce the statistical predictions of the operational theory,

$$\mathrm{pr}(m,s|M,S) = \sum_{\lambda \in \Lambda} \xi(m|M,\lambda)\mu(\lambda,s|S). \qquad (1)$$

In the case of quantum theory, two measurement procedures differ *only* by context if and only if they yield the same statistics *for all quantum states*. In this case, they are represented by the *same* positive operator-valued measure (POVM).[3] Equivalently, two measurement *events* differ only by context if and only if they are assigned the same probability by all quantum states. In this case, they are represented by the same positive operator less than identity (or the same projector in the case of a sharp measurement).

By analogy, in an arbitrary operational theory, two measurement events, $[m|M]$ and $[m'|M']$, differ only by context if and only if for every preparation procedure, the probability of $[m|M]$ is the same as that of $[m'|M']$. The condition can be

---

the difference between Bell's 1964 no-go theorem establishing the impossibility of a locally causal model of quantum theory [8] and the CHSH inequalities [9], which are constraints on operational statistics that follow directly from the assumption of local causality, without presuming the validity of quantum theory.

[2]Note that our notational conventions here align with those of Ref. [11] rather than those of Ref. [10]. For example, the preparation procedure $[\top|S]$ would be represented as $P_S^{(\mathrm{ave})}$ in the notation of Ref. [10], where S is the choice of source setting.

[3]The type of measurement that is considered in most discussions of the Kochen-Specker theorem is a projective measurement (which we here refer to as a *sharp* measurement). Such measurements are a special class of POVMs wherein the positive operators are all projectors. Note that these are the only measurements in quantum theory that can be represented by a single Hermitian operator, namely, the one whose spectral projectors are the elements of the projector-valued measure.

formalized as[4]

$$\forall [s|S] : \text{pr}(m|M,s,S) = \text{pr}(m'|M',s,S). \quad (2)$$

When measurement events $[m|M]$ and $[m'|M']$ differ only by context, they are said to be *operationally equivalent*, denoted $[m|M] \simeq [m'|M']$.

In Ref. [13], a *measurement noncontextual* ontological model was defined to be one wherein operationally equivalent measurement events are represented by equivalent response functions:

$$[m|M] \simeq [m'|M'] \Rightarrow \xi(m|M,\lambda) = \xi(m'|M',\lambda), \forall \lambda \in \Lambda, \quad (3)$$

where one allows the measurements to respond *in*deterministically to λ: $\xi(m|M,\lambda) \in [0,1]$. On the other hand, many have proposed to generalize the notion of KS noncontextuality from quantum theory to arbitrary operational theories in a different manner, namely, by assuming Eq. (3) but with measurements responding deterministically, that is,

$$\xi(m|M,\lambda) \in \{0,1\}. \quad (4)$$

We term the latter proposal *KS noncontextuality*:

KS noncontextuality

= Measurement noncontextuality [Eq. (3)]

+ Outcome determinism [Eq. (4)] $\quad \forall \lambda \in \Lambda. \quad (5)$

In the following, we let $\mathfrak{M}$ denote a set of measurement procedures whose operational features include the compatibility relations and operational equivalences that their quantum counterparts satisfy in a proof of the KS theorem.

To operationalize the KS theorem, the notion of compatibility must also be generalized to an arbitrary operational theory. We follow the proposal of Ref. [14]: measurements M_1 and M_2 are deemed compatible if there is a third measurement with an outcome set that is the Cartesian product of the two outcome sets such that one simulates M_1 and M_2 by marginalization (i.e., for which the marginalized versions are operationally equivalent to M_1 and M_2).

III. FROM LOGICAL PROOFS OF THE KS THEOREM TO OPERATIONAL CRITERIA FOR UNIVERSAL NONCONTEXTUALITY

It follows from the above that the proof schema which generalizes a logical proof of the KS theorem from quantum theory to an arbitrary operational theory is of the following form:

Proposition 1 (no-go for KS noncontextuality from logical proof):

Measurement noncontextuality [Eq. (3)]

+ Outcome determinism [Eq. (4)] $\forall \lambda \in \Lambda, \forall M \in \mathfrak{M}$

+ Operational equivalences in the set $\mathfrak{M}$ (proof-dependent)

$\Rightarrow$ Contradiction.

In the quantum case, this contradiction can be inferred from KS uncolorability; Kochen and Specker's original proof of the KS theorem is an example [1].

Note that in the face of this contradiction, one can always salvage the spirit of noncontextuality (measurement noncontextuality) simply by abandoning outcome determinism. By contrast, this is not a way out of Bell's theorem because the notion of local causality does not presume outcome determinism.

For these reasons, it was argued in Ref. [13] that one should *drop* the assumption of outcome determinism that is part of KS noncontextuality and simply assume measurement noncontextuality. Such a move blocks the derivation of the contradiction in Proposition 1. It might appear, therefore, that there is in fact *no* conflict between quantum theory and the spirit of noncontextuality if one excises the notion of outcome determinism from the latter. However, it turns out that the property of outcome determinism can be *inferred* for certain measurements by applying a notion of noncontextuality to preparations [13], as we now explain.

Two source events, $[s|S]$ and $[s'|S']$, are operationally equivalent, denoted $[s|S] \simeq [s'|S']$, if for every measurement event $[m|M]$, the joint probability of obtaining $[s|S]$ and $[m|M]$ is the same as that of obtaining $[s'|S']$ and $[m|M]$,

$$\forall [m|M] : p(m,s|M,S) = p(m,s'|M,S'). \quad (6)$$

The assumption of preparation noncontextuality requires that operationally equivalent source events should be represented equivalently in the ontological model:[5]

$$[s|S] \simeq [s'|S'] \Rightarrow \mu(\lambda,s|S) = \mu(\lambda,s'|S'), \quad \forall \lambda \in \Lambda. \quad (7)$$

It was argued in Ref. [13] that whatever reasons can be given in support of measurement noncontextuality, these are also reasons to believe in preparation noncontextuality and therefore that the only reasonable assumption to make is noncontextuality for all experimental procedures, termed *universal noncontextuality*. It is the assumption we make here.

Reference [15] showed that for quantum theory, preparation noncontextuality implies that measurements should be assigned outcomes deterministically if and only if they are projective. Reference [10] generalized the logic to all operational theories by focusing not just on the set $\mathfrak{M}$ of measurements, but on a corresponding set $\mathfrak{S}$ of sources as well. Outcome determinism is justified when the following operational criteria are satisfied:

(i) For each equivalence class of measurements, $\mathcal{M}_i \in \mathfrak{M}$, there must exist an equivalence class of sources, $\mathcal{S}_i \in \mathfrak{S}$, such that the outcomes of $\mathcal{M}_i$ and $\mathcal{S}_i$, denoted m_i and $\mathcal{s}_i$, respectively, are perfectly correlated. This implies that the average correlation over the pairings $\{(\mathcal{M}_i, \mathcal{S}_i) : i \in \{1, \ldots, n\}\}$,

$$\text{Corr} \equiv \frac{1}{n} \sum_{i=1}^{n} \sum_{m_i, \mathcal{s}_i} \delta_{m_i, \mathcal{s}_i} \text{pr}(m_i, \mathcal{s}_i | \mathcal{M}_i, \mathcal{S}_i), \quad (8)$$

satisfies

$$\text{Corr} = 1. \quad (9)$$

---

[4]By a Bayesian inversion, this condition is equivalent to $\forall [s|S] :$ $\text{pr}(m,s|M,S) = \text{pr}(m',s|M',S)$, a form that makes more apparent the close analogy with the operational equivalence relation among source events which we define further on.

[5]Note that Eq. (6) is analogous to Eq. (2) and Eq. (7) is analogous to Eq. (3).

(ii) The sources must obey the following operational equivalence relations:

$$\forall i, i' : [\top | \mathcal{S}_i] \simeq [\top | \mathcal{S}_{i'}], \tag{10}$$

where, as stipulated earlier, $[\top | \mathcal{S}]$ denotes the event corresponding to implementing $\mathcal{S}$ and not conditioning on its outcome.

Note that if the source event $[s | S]$ is represented in the ontological model by $\mu(\lambda, s | S)$, then $[\top | S]$ is represented by

$$\mu(\lambda | S) \equiv \sum_s \mu(\lambda, s | S). \tag{11}$$

The inference established in Ref. [10] can then be expressed as follows:

Proposition 2 (justifying outcome determinism):

Preparation noncontextuality [Eq. (7)]

   + Operational equivalences in the set $\mathfrak{S}$ [Eq. (10)]

   + Perfect correlation between outcomes of $\mathcal{S}_i$ and $\mathcal{M}_i$ for all i [Eq. (9)]

     $\Rightarrow$ Outcome determinism $\forall \lambda \in \cup_{\mathcal{S} \in \mathfrak{S}} \text{supp}[\mu(\cdot | \mathcal{S})]$ [Eqs. (14), (15)] $\forall M \in \mathfrak{M}$.

The proof is as follows. From Eqs. (7) and (10) we conclude that

$$\forall i, i' : \mu(\lambda | \mathcal{S}_i) = \mu(\lambda | \mathcal{S}_{i'})$$
$$\equiv \nu(\lambda). \tag{12}$$

By Bayesian inversion, $\mu(s | \lambda, \mathcal{S}) = \mu(\lambda, s | \mathcal{S}) / \nu(\lambda)$. Substituting this into Eq. (1), we have

$$\text{pr}(m_i, s_i | \mathcal{M}_i, \mathcal{S}_i) = \sum_{\lambda \in \Lambda} \xi(m_i | \mathcal{M}_i, \lambda) \mu(s_i | \lambda, \mathcal{S}_i) \nu(\lambda). \tag{13}$$

Given this expression, the only way to explain the perfect correlation of Eq. (9), then, is if the measurements respond deterministically for all ontic states in the support of ν, that is,

$$\forall \lambda \in \text{supp}(\nu) : \xi(m_i | \mathcal{M}_i, \lambda) \in \{0, 1\}, \tag{14}$$

where $\text{supp}(\nu) \equiv \{\lambda \in \Lambda : \nu(\lambda) > 0\}$. But given Eq. (12),

$$\text{supp}(\nu) = \text{supp}[\mu(\cdot | \mathcal{S})] \ \forall \mathcal{S} \in \mathfrak{S}, \tag{15}$$

which concludes the proof.

To obtain a contradiction in the no-go theorem of Proposition 1, it is sufficient to assume outcome determinism just for the ontic states in the union of the ontic supports of the distributions representing the sources in $\mathfrak{S}$, even though this might be a subset of the full set of ontic states. This is because any ontic state outside this subset does not affect the operational statistics of any experiment involving measurements on sources in $\mathfrak{S}$. Hence, the premise of outcome determinism $\forall \lambda \in \Lambda$ in the no-go of Proposition 1 can be replaced by the same premise $\forall \lambda \in \cup_{\mathcal{S} \in \mathfrak{S}} \text{supp}[\mu(\cdot | \mathcal{S})]$ and thus by the antecedent of Proposition 2.

The no-go that one obtains by combining Proposition 2 with Proposition 1 is a no-go for universal noncontextuality based on a logical proof of the KS theorem.

Proposition 3 (no-go for universal noncontextuality from logical proof):

Universal noncontextuality [Eq. (3), Eq. (7)]

   + Operational equivalences in the set $\mathfrak{S}$ [Eq. (10)]

   + Operational equivalences in the set $\mathfrak{M}$ (proof-dependent)

   + Perfect correlation between outcomes of $\mathcal{S}_i$ and $\mathcal{M}_i$ for all i [Eq. (9)]

     $\Rightarrow$ Contradiction.

As noted in Ref. [10], it implies that any operational theory that *does* admit of a universally noncontextual model while exhibiting the appropriate operational features of $\mathfrak{M}$ and $\mathfrak{S}$ must exhibit *imperfect* correlations for the pairings $\{(\mathcal{S}_i, \mathcal{M}_i)\}$; that is, it must satisfy Corr < 1.

The precise amount by which Corr is bounded away from 1 is determined as follows. Substituting Eq. (13) into the definition of Corr [Eq. (8)], we obtain

$$\text{Corr} = \sum_\lambda \text{Corr}(\lambda) \nu(\lambda), \tag{16}$$

where

$$\text{Corr}(\lambda) \equiv \frac{1}{n} \sum_{i=1}^n \sum_{m_i, s_i} \delta_{m_i, s_i} \xi(m_i | \mathcal{M}_i, \lambda) \mu(s_i | \mathcal{S}_i, \lambda). \tag{17}$$

For a given choice of λ corresponding to a noncontextual assignment to the $\{\mathcal{M}_i\}_i$, Corr(λ) is maximized by taking $\mu(s_i | \mathcal{S}_i, \lambda) = 1$ for $s_i = m_i^{\max}$, where $m_i^{\max}$ is any value of m_i such that $\max_{m_i} \xi(m_i | \mathcal{M}_i, \lambda) = \xi(m_i^{\max} | \mathcal{M}_i, \lambda)$. We then have Corr$(\lambda) = \frac{1}{n} \sum_{i=1}^n \max_{m_i} \xi(m_i | \mathcal{M}_i, \lambda)$. Because every noncontextual assignment is indeterministic for some $\mathcal{M}_i$, Corr(λ) is bounded away from 1 for all λ. Letting Corr$_{\text{ind}}$ denote the maximum value of $\frac{1}{n} \sum_{i=1}^n \max_{m_i} \xi(m_i | \mathcal{M}_i, \lambda)$ in a variation over λ that correspond to indeterministic noncontextual assignments, we have

$$\text{Corr} \leqslant \text{Corr}_{\text{ind}}. \tag{18}$$

The qualifier that λ correspond to indeterministic noncontextual assignments in the variation over λ that defines Corr$_{\text{ind}}$ may seem unnecessary at this stage since *all* λ in a logical proof of the KS theorem correspond to such assignments. However, we will soon consider statistical proofs of the KS theorem, where there exist λ that correspond to deterministic noncontextual assignments and where the qualifier that Corr$_{\text{ind}}$ is computed by varying over λ that correspond to indeterministic noncontextual assignments becomes necessary.

The compatibility and operational equivalence relations on the $\{\mathcal{M}_i\}_i$, combined with the assumption of measurement noncontextuality, define linear constraints on the n-tuple of response functions $\{\xi(m_i | \mathcal{M}_i, \lambda)\}_i$. These linear constraints describe the facets of a polytope, termed the noncontextual measurement-assignment polytope. One can obtain the vertices of this polytope from its facets using convex hull algorithms. One determines Corr$_{\text{ind}}$ by determining the maximum value of $\frac{1}{n} \sum_{i=1}^n \max_{m_i} \xi(m_i | \mathcal{M}_i, \lambda)$ in a variation over the (indeterministic) vertices. The noncontextuality inequality one obtains for a given logical proof of the KS theorem, therefore, is simply the inequality one obtains by substituting the determined value of Corr$_{\text{ind}}$ into Eq. (18). References [10,11] provide examples of how to derive such inequalities for specific logical proofs of the KS theorem.

IV. FROM STATISTICAL PROOFS OF THE KS THEOREM TO OPERATIONAL CRITERIA FOR UNIVERSAL NONCONTEXTUALITY

We can now turn to the question of how to obtain operational criteria for the failure of universal noncontextuality from statistical, rather than logical, proofs of the KS theorem. In the operationalized version of such proofs (i.e., one that makes no reference to the quantum formalism), the contradiction is achieved by noting that there is a special source event, that is, the event of obtaining a special outcome $\jmath_* = 0$ of a special source $\mathcal{S}_*$, denoted $[\jmath_* = 0|\mathcal{S}_*]$, such that the measurement statistics one obtains for the special preparation associated with this event are inconsistent with KS noncontextuality.

We index the compatible subsets of $\mathfrak{M}$ by α and denote the equivalence class of measurements that jointly simulates the elements of this subset by $\mathcal{M}^{(\alpha)}$, with the vector of outcomes denoted by $\vec{m}^{(\alpha)}$.

Let R be the value of a particular linear function F of the operational statistics for the compatible subsets of measurements when these are implemented on the special preparation, i.e.,

$$R = F(\{\mathrm{pr}(\vec{m}^{(\alpha)}|\mathcal{M}^{(\alpha)}, \jmath_* = 0, \mathcal{S}_*)\}_\alpha). \tag{19}$$

We define R_{det} to be the largest value of R consistent with *deterministic* noncontextual measurement assignments, and R_{ind} to be the largest value of R consistent with *indeterministic* noncontextual measurement assignments. For every statistical proof of the KS theorem, it is possible to construct a function F such that

$$R_{\mathrm{ind}} \geqslant R > R_{\mathrm{det}}. \tag{20}$$

[An example is given in Eq. (23).]

We define

$$p_* \equiv \mathrm{pr}(\jmath_* = 0|\mathcal{S}_*) \tag{21}$$

to be the probability of the source $\mathcal{S}_*$ yielding the outcome $\jmath_* = 0$. The assumption that the special preparation sometimes occurs can be formalized as

$$p_* > 0. \tag{22}$$

Recalling (5), it follows that the proof schema for a no-go theorem for KS noncontextuality based on a statistical proof of the KS theorem is as follows:

Proposition 4 (no-go for KS noncontextuality from statistical proof):

Measurement noncontextuality [Eq. (3)]

+ Outcome determinism [Eq. (4)] $\forall \lambda \in \Lambda, \forall M \in \mathfrak{M}$

+ Operational equivalences in the set $\mathfrak{M}$ (proof-dependent)

+ Features of correlations among compatible subsets of $\mathfrak{M}$ for the special preparation [Eq. (20)]

+ Nonzero probability of the special preparation [Eq. (22)]

$\Rightarrow$ Contradiction.

A simple example of such a no-go theorem is based on the n-cycle scenario for odd n [5,14,16,17]. Here, there are n equivalence classes of binary-outcome measurements, $\mathfrak{M} \equiv \{\mathcal{M}_i\}_{i=1}^n$, where adjacent pairs are compatible, so that there are n compatible subsets, which we can index by
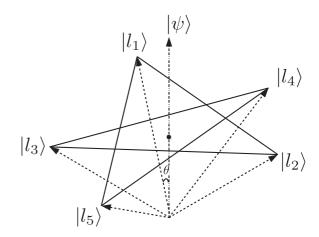

FIG. 1. Quantum construction for a statistical proof of the KS theorem based on the 5-cycle [5].

$\{(1,2),(2,3),\ldots,(n-1,n),(n,1)\}$. Let $\mathcal{M}^{(i,i\oplus1)}$ denote the equivalence class of measurements which jointly simulates $\mathcal{M}_i$ and $\mathcal{M}_{i\oplus1}$ (here, $\oplus$ denotes sum modulo n). Let $M^{(i,i\oplus1)}$ denote a procedure in the equivalence class $\mathcal{M}^{(i,i\oplus1)}$, and let $M_{i(i\oplus1)}$ be the procedure one obtains by implementing $M^{(i,i\oplus1)}$ and marginalizing over the outcome of $\mathcal{M}_{i\oplus1}$. Note that by this definition, $M_{i(i\oplus1)}$ is in the equivalence class $\mathcal{M}_i$. Define $M_{i(i\ominus1)}$ similarly. The difference between $M_{i(i\oplus1)}$ and $M_{i(i\ominus1)}$ is merely a difference of context, corresponding to the neighbor with which $\mathcal{M}_i$ is jointly implemented. The relevant operational equivalence relations (implicit in the definition of the equivalence classes) are therefore $\forall i : M_{i(i\oplus1)} \simeq M_{i(i\ominus1)}$. An assignment of outcomes (deterministic or indeterministic) to these measurements in the ontological model is noncontextual if it is independent of this choice.

Clearly, if n is odd, then not all adjacent pairs of measurements can have anticorrelated outcomes if the assignment is deterministic. At most, this can occur for $n-1$ out of the n pairs. Consequently, if one defines R to be the probability of seeing anticorrelated outcomes when jointly implementing an adjacent pair of measurements ($\mathcal{M}_i$ and $\mathcal{M}_{i\oplus1}$) on the special preparation, averaged over all such pairs, i.e.,

$$R \equiv \frac{1}{n} \sum_{i=1}^n \mathrm{pr}(m_i \neq m_{i\oplus1}|\mathcal{M}^{(i,i\oplus1)}, \jmath_* = 0, \mathcal{S}_*), \tag{23}$$

then the largest value achievable by deterministic noncontextual assignments is $R_{\mathrm{det}} = \frac{n-1}{n}$, and KS noncontextuality implies $R \leqslant \frac{n-1}{n}$. It follows that in any operational theory that predicts $p_* > 0$ and $R > \frac{n-1}{n}$, we obtain a contradiction with KS noncontextuality.

As is well known, this is the case for quantum theory. The first instance of such a proof, due to Klyachko, Can, Binicioglu, and Shumovsky (KCBS) [5], was for the 5-cycle. It showed that the compatibility relations required to hold among the $\{\mathcal{M}_i\}_{i=1}^5$ can be achieved with sharp quantum measurements on a qutrit if $\mathcal{M}_i$ corresponds to the projector-valued measure $\{|l_i\rangle\langle l_i|, \mathbb{1} - |l_i\rangle\langle l_i|\}$, where $|l_i\rangle = (\sin\theta\cos\phi_i, \sin\theta\sin\phi_i, \cos\theta)$, $\phi_i = \frac{4\pi i}{5}$, and $\cos\theta = \frac{1}{\sqrt[4]{5}}$. These are depicted as 3-dimensional vectors in Fig. 1. The special preparation event $[\jmath_* = 0|\mathcal{S}_*]$ corresponds to the quantum state $|\psi\rangle = (0,0,1)$, also depicted

in Fig. 1. Consequently, by letting $\mathcal{S}_*$ be a quantum source that prepares $|\psi\rangle$ with nonzero probability, we ensure that $p_* > 0$.

The average anticorrelation for $|\psi\rangle$ is found to be $R = \frac{2}{\sqrt{5}} \approx 0.89442$, contradicting the prediction of KS noncontextuality that $R \leqslant \frac{4}{5}$. All of this generalizes to arbitrary odd $n \geqslant 5$ [14,16,17]: the $|l_i\rangle$ have the same form with $i \in \{1, 2, \ldots, n\}$, $\phi_i = \frac{n-1}{n}\pi i$, and $\cos^2\theta = \cos(\pi/n)/[1 + \cos(\pi/n)]$, with $|\psi\rangle$ as before, so that $p_* > 0$ and $R = \frac{2\cos(\frac{\pi}{n})}{1+\cos(\frac{\pi}{n})} > \frac{n-1}{n}$.

To convert a statistical proof of the KS theorem into a proof of the failure of universal noncontextuality, we proceed analogously to the conversion procedure for logical proofs. We note that in the no-go of Proposition 4, one can continue to derive a contradiction if one assumes outcome determinism just for the ontic states in the union of the supports of the distributions for the sources in the set $\mathfrak{S}$ (redefined to include the special source $\mathcal{S}_*$), as long as one assumes that the operational features of the set $\mathfrak{S}$ now include the fact that $[\top|\mathcal{S}_*]$ is operationally equivalent with all other marginalized sources,

$$\forall i, i' : [\top|\mathcal{S}_i] \simeq [\top|\mathcal{S}'_{i'}] \simeq [\top|\mathcal{S}_*]. \tag{24}$$

Therefore, among the premises of Proposition 4, we can replace the assumption of outcome determinism with the antecedent of the inference of Proposition 2 [with $\mathfrak{S}$ now including $\mathcal{S}_*$ and Eq. (10) replaced by Eq. (24)], to obtain a no-go theorem for universal noncontextuality:

Proposition 5 (no-go for universal noncontextuality from statistical proof):

Universal noncontextuality [Eq. (3), Eq. (7)]
- Operational equivalences in the set $\mathfrak{S}$ [Eq. (24)]
- Operational equivalences in the set $\mathfrak{M}$ (proof-dependent)
- Perfect correlation between outcomes of $\mathcal{S}_i$ and $\mathcal{M}_i$ for all i [Eq. (9)]
- Features of correlations among compatible subsets of $\mathfrak{M}$ for the special preparation [Eq. (20)]
- Nonzero probability of the special preparation [Eq. (22)]
  $\Rightarrow$ Contradiction.

It is useful to see how this proof schema yields a no-go for universal noncontextuality in quantum theory when one particularizes to the n-cycle scenario. Let $\{\mathcal{M}_i\}_i$ be the set of projective binary-outcome measurements on a qutrit specified earlier. Let $\mathcal{S}_i$ be the quantum source that prepares $|l_i\rangle\langle l_i|$ with probability $\frac{1}{3}$ and $\frac{\mathbb{1}-|l_i\rangle\langle l_i|}{2}$ with probability $\frac{2}{3}$. Similarly, let $\mathcal{S}_*$ be the quantum source that prepares $|\psi\rangle\langle\psi|$ with probability $\frac{1}{3}$ (corresponding to outcome $\jmath_* = 0$) and $\frac{\mathbb{1}-|\psi\rangle\langle\psi|}{2}$ with probability $\frac{2}{3}$. Clearly, we have the operational equivalences of (24) by virtue of the fact that all of these ensembles average to $\frac{1}{3}\mathbb{1}$. Furthermore, we have Corr $= 1$. Thus all of the antecedents of the inference of Proposition 5 are satisfied, and we have a proof of the failure of universal noncontextuality in quantum theory.

With the proof schema of Proposition 5 in hand, we can finally turn to the question of how to derive a noncontextuality inequality from statistical proofs of the KS theorem.

It suffices to note that Proposition 5 implies that any operational theory that *does* admit of a universally noncontextual model while exhibiting the specified operational features of the sets $\mathfrak{M}$ and $\mathfrak{S}$ cannot satisfy Eqs. (9), (20), and (22); that is, it cannot satisfy Corr $= 1$, $R > R_{\det}$, and $p_* > 0$. To derive a noncontextuality inequality, therefore, one must simply determine the precise trade-off relation satisfied by Corr, R, and p_* in a universally noncontextual model.

Applying the assumption of preparation noncontextuality to Eq. (24), we can infer that

$$\nu(\lambda) = \mu(\lambda|\mathcal{S}_*) = \sum_{\jmath_*} \mu(\lambda|\jmath_*, \mathcal{S}_*)\mathrm{pr}(\jmath_*|\mathcal{S}_*). \tag{25}$$

Substituting this into Eq. (16), we obtain

$$\mathrm{Corr} = \sum_{\jmath_*} \mathrm{pr}(\jmath_*|\mathcal{S}_*)\mathrm{Corr}(\jmath_*), \tag{26}$$

where

$$\mathrm{Corr}(\jmath_*) \equiv \sum_{\lambda} \mathrm{Corr}(\lambda)\mu(\lambda|\jmath_*, \mathcal{S}_*). \tag{27}$$

$\mathrm{Corr}(\jmath_*)$ quantifies the average degree of correlation for the pairings $\{(\mathcal{S}_i, \mathcal{M}_i)\}$ predicted by ontic state λ, averaged over the ontic states in the support of $\mu(\cdot|\jmath_*, \mathcal{S}_*)$.

The argument proceeds by showing that the quantity $\mathrm{Corr}(\jmath_* = 0)$ has a nontrivial upper bound.

Define

$$R(\lambda) = F(\{\xi(\vec{m}^{(\alpha)}|\mathfrak{M}^{(\alpha)}, \lambda)\}_\alpha), \tag{28}$$

where F is the linear function specified in Eq. (19), so that the expression for R in the ontological model is

$$R = \sum_{\lambda} R(\lambda)\mu(\lambda|\jmath_*, \mathcal{S}_*). \tag{29}$$

Recalling that $R_{\det}$ denotes the maximum value that can be achieved by deterministic noncontextual assignments to the measurements, if $R > R_{\det}$, then some of the ontic states in the support of $\mu(\cdot|\jmath_* = 0, \mathcal{S}_*)$ must be inconsistent with a convex mixture of deterministic noncontextual assignments. In this case, $\mathrm{Corr}(\jmath_* = 0)$ must be bounded away from 1,

$$\mathrm{Corr}(\jmath_* = 0) < 1. \tag{30}$$

By contrast, given that the no-go result does not make any appeal to the statistics of measurements on the preparations associated with $[\jmath_* \neq 0|\mathcal{S}_*]$, the ontic states in the support of $\mu(\cdot|\jmath_* \neq 0, \mathcal{S}_*)$ could potentially assign outcomes to the measurements deterministically, which in turn implies that $\mathrm{Corr}(\jmath_* \neq 0)$ can only be upper bounded by its logical maximum,

$$\mathrm{Corr}(\jmath_* \neq 0) \leqslant 1. \tag{31}$$

In all, therefore, we have

$$\mathrm{Corr} \leqslant p_*\mathrm{Corr}(\jmath_* = 0) + (1 - p_*). \tag{32}$$

Given the dependence of $\mathrm{Corr}(\jmath_* = 0)$ on R, this equation specifies a trade-off relation between Corr, R, and p_*. Such a trade-off relation constitutes a noncontextuality inequality derived from a statistical proof of the KS theorem.

The precise amount by which Corr is bounded away from 1 for a given value of R in Eq. (30) depends on three quantities: (i)

the maximum value of $R(\lambda)$ for any deterministic noncontextual assignment, denoted here by R_{det}, (ii) the maximum value of $R(\lambda)$ for any indeterministic noncontextual assignment, denoted here by R_{ind}, and (iii) the maximum value of $\text{Corr}(\lambda)$ for any indeterministic noncontextual assignment, which (as in the case of logical proofs of the KS theorem) we denote by Corr_{ind}. The values of R_{det}, R_{ind}, and Corr_{ind} depend on the particular statistical proof of the KS theorem one is considering. We will show that

$$\text{Corr}(\jmath_* = 0) \leqslant \frac{R_{\text{ind}} - R}{R_{\text{ind}} - R_{\text{det}}}(1 - \text{Corr}_{\text{ind}}) + \text{Corr}_{\text{ind}}. \quad (33)$$

Substituting this into Eq. (32), we obtain the main result of this article.

Theorem 1. In a prepare-and-measure experiment that admits of a universally noncontextual ontological model, the following trade-off relation between Corr, R, and p_* [defined in Eqs. (8), (19), and (21), respectively] holds:

$$\text{Corr} \leqslant 1 - p_*(1 - \text{Corr}_{\text{ind}})\left(\frac{R - R_{\text{det}}}{R_{\text{ind}} - R_{\text{det}}}\right). \quad (34)$$

This is our noise-robust noncontextuality inequality.

Recall that by assumption, the precise form of R depends on which statistical proof of the KS theorem one is considering, and that the definition of R is such that $R_{\text{ind}} > R_{\text{det}}$.

Note that this inequality implies that if $\text{Corr} = 1$ and $p_* > 0$, then $R \leqslant R_{\text{det}}$, so that our noise-robust noncontextuality inequality for a given statistical proof of the KS theorem reduces to the KS noncontextuality inequality that is conventionally associated with that proof [5,17–22]. Experimentally, however, one *never* achieves perfect correlation; that is, one always finds $\text{Corr} < 1$, so that our criterion for noncontextuality *never* reduces to a conventional KS-noncontextuality inequality in a real experiment. It follows that a violation of a conventional KS-noncontextuality inequality ($R \leqslant R_{\text{det}}$) in a real experiment is insufficient to demonstrate the failure of noncontextuality.

A question that arises at this point is whether it might still be appropriate to test the inequality $R \leqslant R_{\text{det}}$ on the grounds that it tests the assumption of KS noncontextuality rather than the assumption of universal noncontextuality. Recalling from Eq. (5) that the assumption of KS noncontextuality incorporates an assumption of outcome determinism, to adopt such a view would be to simply *assume* outcome determinism rather than seeking to justify it from preparation noncontextuality and perfect correlations between sources and measurements. However, it was shown in Ref. [15] (see also Ref. [23]) that assuming KS noncontextuality (hence outcome determinism) for unsharp quantum measurements leads to absurd conclusions, such as the failure of KS noncontextuality for experiments that are completely classical (in the sense that all states and measurements are diagonal in the same basis). Therefore, the assumption of KS noncontextuality is only applicable to sharp, i.e., noiseless, quantum measurements, which are never achieved experimentally. For operational theories other than quantum theory, the same argument holds: every measurement that can be achieved in a real experiment fails to satisfy the ideal

of noiselessness, and it is only for such noiseless measurements that the assumption of KS noncontextuality is justified.[6]

The noncontextuality inequality of Eq. (34), on the other hand, accommodates noisy experimental data for which $\text{Corr} < 1$. Thus, even if sources and measurements deviate from the ideal of sharpness in an experiment, one can still see a violation of our inequality. It is in this sense that it is robust to noise.

One can also deduce the precise limit to noise tolerance for such an inequality. For a fixed p_*, in order to obtain a nontrivial upper bound on R (i.e., a bound smaller than R_{ind}), one must have $\text{Corr} > 1 - p_*(1 - \text{Corr}_{\text{ind}})$. If the noise is such that Corr is reduced to a value below this bound, then it becomes impossible to witness contextuality via this inequality. Further, if in addition to fixing p_*, one achieves a certain value of R, say $R = R_{\text{expt}}$, then contextuality is witnessed if and only if $\text{Corr} > 1 - p_*(1 - \text{Corr}_{\text{ind}})(\frac{R_{\text{expt}} - R_{\text{det}}}{R_{\text{ind}} - R_{\text{det}}})$ [which is just a rewriting of the violation of Eq. (34)]. In Appendix C, we provide further details about what an experiment must achieve in order to test an inequality of the form of Eq. (34).

The proof of Eq. (33) proceeds as follows. Without any loss of generality, we identify the set of ontic states Λ with the set of vertices of the polytope of noncontextual assignments to the elements of $\mathfrak{M}$. We divide the vertices into two sets, corresponding to deterministic and indeterministic assignments, denoted Λ_{det} and Λ_{ind}, respectively, so that $\Lambda = \Lambda_{\text{det}} \cup \Lambda_{\text{ind}}$. For all $\lambda \in \Lambda_{\text{det}}$, $R(\lambda)$ satisfies the nontrivial upper bound $R(\lambda) \leqslant R_{\text{det}}$ (where $R_{\text{det}} < R_{\text{ind}}$ because R is, by construction, a function that cannot achieve the logically maximal value of R_{ind} for deterministic noncontextual assignments) while $\text{Corr}(\lambda)$ can always achieve its logical maximum of 1, so that the bound is trivial, $\text{Corr}(\lambda) \leqslant 1$. By contrast, for all $\lambda \in \Lambda_{\text{ind}}$, $\text{Corr}(\lambda)$ satisfies the nontrivial upper bound $\text{Corr}(\lambda) \leqslant \text{Corr}_{\text{ind}}$ (where $\text{Corr}_{\text{ind}} < 1$ because indeterministic noncontextual assignments necessarily imply a failure to achieve perfect source-measurement correlations), but because there exist $\lambda \in \Lambda_{\text{ind}}$ such that $R(\lambda)$ achieves its maximum of R_{ind}, we have only the trivial bound $R(\lambda) \leqslant R_{\text{ind}}$. We now make use of these facts to determine the upper bound on $\text{Corr}(\jmath_* = 0)$ for a given value of R. Defining

$$\mu_{\text{det}} \equiv \sum_{\lambda \in \Lambda_{\text{det}}} \mu(\lambda | \jmath_* = 0, \mathcal{S}_*)$$

and

$$\mu_{\text{ind}} \equiv \sum_{\lambda \in \Lambda_{\text{ind}}} \mu(\lambda | \jmath_* = 0, \mathcal{S}_*),$$

so that $\mu_{\text{det}} + \mu_{\text{ind}} = 1$, and recalling Eqs. (27) and (29), we have

$$\text{Corr}(\jmath_* = 0) \leqslant \mu_{\text{det}} + \text{Corr}_{\text{ind}}\mu_{\text{ind}} \quad (35)$$

---

[6]Note that the inappropriateness of applying KS noncontextuality to real experiments has been argued in detail elsewhere over the years [10,11,13,15,23–25] and our comments here are meant merely to highlight the precise sense in which this plays out for statistical proofs of the KS theorem.

and

$$R \leqslant R_{\text{det}}\mu_{\text{det}} + R_{\text{ind}}\mu_{\text{ind}}. \tag{36}$$

Eliminating μ_{det} and μ_{ind} from these constraints, we obtain Eq. (33).

The noncontextuality inequality of Eq. (34) can be saturated by a noncontextual ontological model in certain circumstances.[7] Define $\Lambda_{\text{det}}^{\max} \equiv \{\lambda \in \Lambda_{\text{det}} : R(\lambda) = R_{\text{det}}, \text{Corr}(\lambda) = 1\}$ and $\Lambda_{\text{ind}}^{\max} \equiv \{\lambda \in \Lambda_{\text{ind}} : R(\lambda) = R_{\text{ind}}, \text{Corr}(\lambda) = \text{Corr}_{\text{ind}}\}$. Clearly, $\Lambda_{\text{det}}^{\max} \subseteq \Lambda_{\text{det}}$ and $\Lambda_{\text{ind}}^{\max} \subseteq \Lambda_{\text{ind}}$. For noncontextual measurement-assignment polytopes based on statistical proofs of the KS theorem, $\Lambda_{\text{det}}^{\max}$ is always a nonempty set. If $\Lambda_{\text{ind}}^{\max}$ is nonempty, then the noncontextuality inequality of Eq. (34) can be saturated. The reason is that in this case one can choose the support of $\mu(\lambda|\jmath_* = 0, \mathcal{S}_*)$ on Λ_{det} to be restricted to $\Lambda_{\text{det}}^{\max}$ and the support of $\mu(\lambda|\jmath_* = 0, \mathcal{S}_*)$ on Λ_{ind} to be restricted to $\Lambda_{\text{ind}}^{\max}$, in which case the inequalities in Eqs. (35) and (36) become equalities. The condition that $\Lambda_{\text{ind}}^{\max}$ be nonempty is satisfied for the case of odd n-cycle scenario, and therefore the noncontextuality inequalities that we derive in this case will be tight. To illustrate our technique on a concrete example, we now turn to the odd n-cycle scenario.

For the case of the n-cycle scenario, the function F defining the quantity R is specified in Eq. (23). As noted in our previous discussion of the n-cycle scenario, for deterministic vertices of the noncontextual measurement-assignment polytope, there is a nontrivial upper bound on $R(\lambda)$, namely, $R(\lambda) \leqslant R_{\text{det}}$, where

$$R_{\text{det}} = \frac{n-1}{n}. \tag{37}$$

Similarly, for every indeterministic vertex of the noncontextual measurement-assignment polytope, there is a nontrivial upper bound on $\text{Corr}(\lambda)$, namely, $\text{Corr}(\lambda) \leqslant \text{Corr}_{\text{ind}}$, where

$$\text{Corr}_{\text{ind}} = \tfrac{1}{2}. \tag{38}$$

Also, we have

$$R_{\text{ind}} = 1. \tag{39}$$

(For further discussion of the vertices of the polytope in the case of the n-cycle scenario, see Appendix B.) Substituting Eqs. (37) and (38) into Eq. (34), we find that a noise-robust noncontextuality inequality for the n-cycle scenario is

$$\text{Corr} \leqslant 1 - p_* \frac{n}{2}\left(R - \frac{n-1}{n}\right). \tag{40}$$

The quantum realization of the n-cycle scenario [5,14] that was discussed earlier clearly violates this inequality. It suffices to note that by Eq. (40), if $\text{Corr} = 1$ and $p_* > 0$, then $R \leqslant \frac{n-1}{n}$, while this quantum realization achieves $\text{Corr} = 1$, $p_* = \frac{1}{3}$, and $R = \frac{2\cos(\frac{\pi}{n})}{1+\cos(\frac{\pi}{n})} > \frac{n-1}{n}$. The violation persists in the presence of noise: if $p_* = \frac{1}{3}$, there is a range of values of Corr and R below those achieved in the ideal quantum realization—that is, where $\text{Corr} < 1$ and $R < \frac{2\cos(\frac{\pi}{n})}{1+\cos(\frac{\pi}{n})}$—such that the inequality is still violated.

V. DISCUSSION

As noted earlier, Xu *et al.* [12] have previously obtained noise-robust noncontextuality inequalities starting from the KCBS [5] and Yu-Oh [6] statistical proofs of the KS theorem. In Appendix E, we compare our approach to theirs for the case of the odd n-cycle scenario. One difference is that all of the simulating measurements in their approach are 3-outcome measurements, rather than the 4-outcome case we have considered here (see also Appendix D). We also show that their inequality for the odd n-cycle case is a special case of our noncontextuality inequality for the odd n-cycle scenario, Eq. (40), when p_* is presumed to take the value of 1/3. Hence, unlike Ref. [12], our inequality does not presume that all the ensembles of preparations in the experiment correspond to uniformly random probability distributions (i.e., probability 1/3 for each preparation). Rather, it merely presumes that the relevant operational equivalences hold. If the value of p_* realized in an experiment is different from the 1/3 value of the ideal quantum realization, then our inequality specializes to one that is different from that of Ref. [12]. These differences are not very significant, however. We consider the main advantage of our approach over that of Ref. [12] to be that it makes clear precisely which aspects of the polytope of noncontextual measurement assignments (for any given statistical proof of the KS theorem) need to be identified in order to determine the form of the noncontextuality inequality.

The value of upper bounds on $R(\lambda)$ for deterministic and indeterministic vertices of the noncontextual measurement-assignment polytope, denoted here by R_{det} and R_{ind}, respectively, are well studied for many statistical proofs of the KS theorem [16,17]. The values of the upper bounds on $\text{Corr}(\lambda)$ for indeterministic vertices of this polytope, denoted here by Corr_{ind}, have not been studied previously, but are just as easy to determine.[8] From these, one can determine a noncontextuality inequality for any statistical proof of the KS theorem via Eq. (34). A study of how our technique allows one to convert the graph-theoretic framework of [17] to a hypergraph-theoretic framework for noise-robust noncontextuality inequalities is carried out in [23].

APPENDIX A: ELABORATION OF THE IDEAS UNDERLYING THE TECHNIQUE

The features of the correlations within compatible subsets of $\mathfrak{M}$ (for the special preparation $[\jmath_* = 0|\mathcal{S}_*]$) which underlie the

---

[7]See Sec. VI B of Ref. [23] for a detailed discussion of this noncontextual ontological model.

[8]Indeed, this quantity makes appearance as a hypergraph invariant in Ref. [23], in addition to the usual invariants in the graph-theoretic approach of Ref. [17].

no-go theorem of Proposition 5 constitute a witness that these correlations *cannot* arise from a distribution $\mu(\lambda|s_* = 0, \mathcal{S}_*)$ supported *only* on ontic states corresponding to measurement-noncontextual and outcome-deterministic assignments. These features are represented by a quantity R that is upper bounded by a constant number if the correlations *do* arise from a distribution $\mu(\lambda|s_* = 0, \mathcal{S}_*)$ supported only on such ontic states.

This is precisely analogous to how, in a Bell scenario, the violation of a Bell inequality witnesses the fact that the distribution over ontic states has support on indeterministic vertices of the no-signalling polytope, e.g., PR boxes [26] in the case of the CHSH scenario. The ontic states corresponding to deterministic vertices of the no-signaling polytope define the Bell polytope. A Bell violation thus rules out outcome-deterministic locally-causal ontological models.

By Fine's theorem [27], a Bell violation also rules out locally-causal ontological models that are outcome-*in*deterministic. This is because the notion of local causality implies factorizability of the joint response function, so that these joint response functions are convex mixtures of outcome-deterministic assignments to the local measurements. In other words, Fine's theorem shows that there is no loss of generality in assuming outcome determinism in tests of locality.[9] However, for tests of noncontextuality involving nontrivial applications of the assumption of measurement non-contextuality (as is the case for every test of noncontextuality arising from a statistical or logical proof of the KS theorem), there is no analog of Fine's theorem. The reasons for this are described in Ref. [15].

Now, for Corr to be bounded away from 1, a nonempty subset of the ontic states in the union of the supports of sources in $\mathfrak{S}$ must correspond to the indeterministic vertices of the polytope of measurement noncontextual assignments of probabilities to measurement outcomes. In the case of noncontextuality inequalities inspired by logical proofs of the KS theorem [10], *all* the ontic states in the union of the supports of sources in $\mathfrak{S}$ correspond to indeterministic vertices of this polytope, simply because the polytope admits *no* deterministic vertices on account of the KS uncolorability that such proofs hinge upon. On the other hand, for noncontextuality inequalities inspired by statistical proofs of the KS theorem, we need a witness for the fact that some nonempty subset of the union of the ontic supports of sources in $\mathfrak{S}$ corresponds to indeterministic vertices of the measurement noncontextuality polytope. Only then can we expect Corr to be bounded away from 1. This witness corresponds to the quantity R exceeding its KS noncontextual bound R_{det}.

APPENDIX B: NONCONTEXTUAL MEASUREMENT-ASSIGNMENT POLYTOPE FOR THE n-CYCLE SCENARIO FOR ODD n

An ontological model must specify a conditional probability distribution for every compatible subset of measurements. In the n-cycle scenario, there are n such subsets, corresponding to
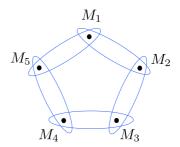
FIG. 2. Compatibility relations for the 5-cycle case: the vertices represent the equivalence classes of binary-outcome measurements and the edges denote compatibility (i.e., joint measurability) of the vertices they contain. Specifically, the ith vertex denotes the equivalence class $\mathcal{M}_i$, and the ith edge denotes the equivalence class $\mathcal{M}^{(i,i\oplus 1)}$.

all adjacent pairs of measurements in the cycle. This is depicted in Fig. 2. Recalling that $\mathcal{M}^{(i,i\oplus 1)}$ denotes the equivalence class of measurement procedures that jointly simulate $\mathcal{M}_i$ and $\mathcal{M}_{i\oplus 1}$, an ontological model must specify an n-tuple of response functions of the form

$$\xi(m_i, m_{i\oplus 1}|\mathcal{M}^{(i,i\oplus 1)}, \lambda). \tag{B1}$$

Marginalization of an outcome of a procedure is modeled in the ontological model by marginalization of the response function, so that the ontological representation of $\mathcal{M}_i$ satisfies

$$\xi(m_i|\mathcal{M}_i, \lambda) = \sum_{m_{i\oplus 1}} \xi(m_i, m_{i\oplus 1}|\mathcal{M}^{(i,i\oplus 1)}, \lambda). \tag{B2}$$

The assumption of noncontextuality implies that the response function representing a measurement depends only on its equivalence class, so that we infer the constraints

$$\forall i : \sum_{m_{i\ominus 1}} \xi(m_i, m_{i\ominus 1}|\mathcal{M}^{(i,i\ominus 1)}, \lambda) = \sum_{m_{i\oplus 1}} \xi(m_i, m_{i\oplus 1}|\mathcal{M}^{(i,i\oplus 1)}, \lambda).$$

$$\tag{B3}$$

Together with the conditions of being a probability distribution, $0 \leqslant \xi(m_i, m_{i\oplus 1}|\mathcal{M}^{(i,i\oplus 1)}, \lambda) \leqslant 1$, and of normalization, $\sum_{m_i, m_{i\oplus 1}} \xi(m_i, m_{i\oplus 1}|\mathcal{M}^{(i,i\oplus 1)}, \lambda) = 1$, Eq. (B3) defines the constraints on the response functions. The set of solutions to these constraints in turn determines the set of solutions for the n-tuple of response functions for the binary-outcome measurements, $\{\xi(m_i|\mathcal{M}_i, \lambda)\}_i$, through Eq. (B2).

For every λ, such an n-tuple of response functions defines a possible n-tuple of (deterministic or indeterministic) assignments to all of the measurements. Following Ref. [11], we term the latter set the noncontextual measurement-assignment polytope. (It is equivalent to what is termed the "no-disturbance" polytope elsewhere; for the n-cycle scenario, it was characterized in Ref. [16].)

There are two types of vertex for this polytope, corresponding to noncontextual measurement assignments that are deterministic and indeterministic respectively. If we identify the set of ontic states Λ with the set of vertices of the polytope, as in the main text, then the two types define a partition of Λ into subsets Λ_{det} and Λ_{ind}.

The deterministic vertices are simply those that assign an outcome to each of the $\mathcal{M}_i$ independently. There are

$$\xi(m_1|M_1,\lambda) = \delta_{m_1,+1}$$

$$\xi(m_5|M_5,\lambda) = \delta_{m_5,+1} \qquad \xi(m_2|M_2,\lambda) = \delta_{m_2,-1}$$

$$\xi(m_4|M_4,\lambda) = \delta_{m_4,-1} \qquad \xi(m_3|M_3,\lambda) = \delta_{m_3,+1}$$

$$\xi(m_i, m_{i\oplus}|M^{(i,i\oplus1)},\lambda) = \xi(m_i|M_i,\lambda)\xi(m_{i\oplus1}|M_{i\oplus1},\lambda)$$
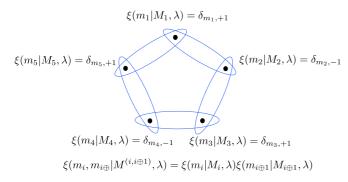
FIG. 3. An example of a deterministic measurement-noncontextual assignment that maximizes the amount of anticorrelation achievable by such assignments, $R(\lambda) = \frac{4}{5}$.

consequently 2^n of these. All of these can achieve the logical maximum value of 1 for Corr(λ), and of these, there are $2n$ that achieve $R(\lambda) = \frac{n-1}{n}$. For instance, this is achieved for the vertices $\{\kappa_j\}_{j=1}^n$ of the form

$$\forall i \in [n] : \xi(m_i, m_{i\oplus1}|\kappa_j) = \xi(m_i|\kappa_j)\xi(m_{i\oplus1}|\kappa_j), \quad \text{(B4)}$$

where

$$\xi(m_i|\kappa_j) = \delta_{m_i,+1} \text{ for } i \in \{j, j \oplus 2, \ldots, j \oplus (n-1)\}$$
$$= \delta_{m_i,-1} \text{ for } i \in \{j \oplus 1, j \oplus 3, \ldots, j \oplus (n-2)\}. \quad \text{(B5)}$$

κ_1 is depicted in Fig. 3.

To obtain another set of n vertices that achieve $R(\lambda) = \frac{n-1}{n}$, it suffices to flip the sign of all the assignments.

These $2n$ vertices constitute the subset of Λ_{det} that can saturate the inequalities Corr(λ) $\leqslant 1$ and $R(\lambda) \leqslant \frac{n-1}{n}$ described in the main text.

The indeterministic vertices are those that exhibit either perfect positive correlation or perfect negative correlation for each of the adjacent pairs of measurements with the number of pairs that exhibit perfect negative correlation being odd. There are 2^{n-1} such assignments. All of these achieve Corr(λ) $= \frac{1}{2}$. One of them also achieves $R(\lambda) = 1$, namely, the one, denoted κ_*, corresponding to perfect negative correlation for all pairs,

$$\forall i \in [n] : \xi(m_i, m_{i\oplus}|\kappa_*) = \frac{1}{2}\delta_{m_i,+1}\delta_{m_{i\oplus},-1} + \frac{1}{2}\delta_{m_i,-1}\delta_{m_{i\oplus},+1}, \quad \text{(B6)}$$

such that the assignment to each $\mathcal{M}_i$ is uniformly random,

$$\forall i \in [n] : \xi(m_i|\kappa_*) = \frac{1}{2}\delta_{m_i,-1} + \frac{1}{2}\delta_{m_i,+1}. \quad \text{(B7)}$$

The vertex κ_* is depicted in Fig. 4. It is the only element of the set Λ_{ind} that saturates the inequalities Corr(λ) $\leqslant \frac{1}{2}$ and $R(\lambda) \leqslant 1$ described in the main text.

For each statistical proof of the KS theorem, one can determine the polytope of noncontextual measurement assignments. The details of this polytope will determine the nontrivial upper bound on $R(\lambda)$ for the deterministic vertices [where R is defined in a manner that is specific to the statistical proof one is considering; see Eq. (19) for the general form and Eq. (23) for an example from the n-cycle scenario] and the nontrivial upper bound on Corr(λ) for the indeterministic vertices. These determinations are all that one requires to derive

$$\xi(m_1|M_1,\lambda)$$
$$= \frac{1}{2}\delta_{m_1,+1} + \frac{1}{2}\delta_{m_1,-1}$$

$$\xi(m_5|M_5,\lambda)$$
$$= \frac{1}{2}\delta_{m_5,+1} + \frac{1}{2}\delta_{m_5,-1}$$

$$\xi(m_2|M_2,\lambda)$$
$$= \frac{1}{2}\delta_{m_2,+1} + \frac{1}{2}\delta_{m_2,-1}$$

$$\xi(m_4|M_4,\lambda)$$
$$= \frac{1}{2}\delta_{m_4,+1} + \frac{1}{2}\delta_{m_4,-1}$$

$$\xi(m_3|M_3,\lambda)$$
$$= \frac{1}{2}\delta_{m_3,+1} + \frac{1}{2}\delta_{m_3,-1}$$

$$\xi(m_i, m_{i\oplus1}|M^{(i,i\oplus1)},\lambda)$$
$$= \frac{1}{2}\delta_{m_i,+1}\delta_{m_{i\oplus1},-1} + \frac{1}{2}\delta_{m_i,-1}\delta_{m_{i\oplus1},+1}$$
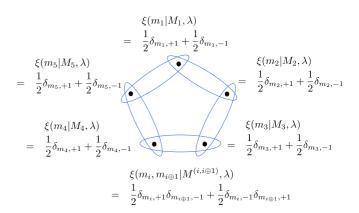
FIG. 4. The indeterministic measurement-noncontextual assignment that achieves perfect anticorrelation, $R(\lambda) = 1$.

a noncontextuality inequality for any given statistical proof of the KS theorem.

APPENDIX C: HOW TO TEST SUCH INEQUALITIES EXPERIMENTALLY

Recall that α is a variable that runs over the compatible subsets of measurements in $\mathfrak{M}$ and that $\mathcal{M}^{(\alpha)}$ denotes the measurement that jointly simulates the compatible subset associated with α, that is, $\{\mathcal{M}_i\}_{i \in \alpha}$. Strictly speaking, $\mathcal{M}^{(\alpha)}$ denotes an equivalence class of measurement procedures. Let $M^{(\alpha)}$ denote a specific procedure in the class $\mathcal{M}^{(\alpha)}$, and let $M_{i(\alpha)}$ denote the procedure in the equivalence class $\mathcal{M}_i$ that is obtained by implementing the joint measurement procedure $M^{(\alpha)}$ and post-processing its outcome (specifically, by marginalizing over the outcomes of all measurements other than $\mathcal{M}_i$ in the compatible subset). Supposing that $\alpha = a_i$ and $\alpha = a_i'$ both correspond to compatible subsets of measurements that include $\mathcal{M}_i$, then $M_{i(\alpha=a_i)}$ and $M_{i(\alpha=a_i')}$ are distinct procedures in the operational equivalence class $\mathcal{M}_i$.

Any experiment that involves the set of measurements $\mathfrak{M}$ and seeks to test noncontextuality must aim to implement a specific measurement procedure $M^{(\alpha)}$ for each α such that every operational equivalence relation of the form $M_{i(\alpha=a_i)} \simeq M_{i(\alpha=a_i')}$ holds.

For instance, in the case of the n-cycle scenario, where $\mathcal{M}^{(i,i\oplus1)}$ denotes the equivalence class of measurement procedures that jointly simulate $\mathcal{M}_i$ and $\mathcal{M}_{i\oplus1}$, where $M^{(i,i\oplus1)}$ denotes a specific procedure in the class $\mathcal{M}^{(i,i\oplus1)}$, and where $M_{i(i\oplus1)}$ denotes the procedure one obtains by implementing $M^{(i,i\oplus1)}$ and marginalizing over the outcome $m_{i\oplus1}$, any experiment that seeks to test noncontextuality must aim to implement a specific measurement procedure $M^{(i,i\oplus1)}$ for each $i \in [n]$ such that the operational equivalence relation $M_{i(i\oplus1)} \simeq M_{i(i\ominus1)}$ holds for all $i \in [n]$.

Furthermore, any experiment that involves the set of sources $\mathfrak{S}$ and seeks to test noncontextuality must aim to implement a specific binary-outcome source procedure S_i for each i, as well as a special source procedure S_* such that the operational equivalence relations $[\top|S_i] \simeq [\top|S_{i'}] \simeq [\top|S_*]$ for all i, i' hold [see Eq. (24)].

Whichever measurement procedures $\{M^{(\alpha)}\}_\alpha$ and source procedures $\{S_i\}_i$ and S_* one targets, however, experimental

imperfections ensure that the relevant operational equivalence relations are not achieved precisely. But given that noncontextuality is an inference from operational equivalences to equivalences in the ontological model, such imprecision blocks the derivation of any consequences for the ontological model of the experiment. This was termed the problem of no strict operational equivalences in Ref. [25]. It was shown there how to solve it using the technique of secondary procedures (see also Sec. V of Ref. [11]). The idea is to identify, within the convex hull of the sources and measurements that were experimentally implemented (termed the primary procedures), sources and measurements that satisfy the operational equivalence relations *exactly* (termed the secondary procedures), and then to test the noncontextuality inequalities on the secondary procedures.

Note that operational equivalence of two measurements (sources) requires equivalence of statistics for *all* sources (measurements), or equivalently, equivalence for a tomographically complete set of sources (measurements). Experiments seeking to test operational equivalence relations, therefore, must accumulate evidence in favor of a given set of procedures being tomographically complete. See, e.g., the evidence described in Ref. [25]. Note that there is a loophole, which we term the *tomography loophole*, for experiments testing universal noncontextuality: no matter how much evidence one accumulates for the tomographic completeness of some set of procedures, it is possible that future experiments will uncover new procedures whose statistics are *not* predicted by the statistics of those in the set. It follows that any hypothesis of tomographic completeness of some set is necessarily tentative. One should endeavor to falsify it experimentally, and as long as it resists falsification, one has good evidence for the hypothesis. But one can never verify it.

APPENDIX D: AN ALTERNATIVE WAY OF OPERATIONALIZING THE KCBS PROOF OF THE KS THEOREM

In this article, we have operationalized the KCBS proof of the KS theorem as an n-cycle scenario with odd n, that is, as n binary-outcome measurements arranged in a cycle such that adjacent pairs are jointly measurable. The outcome set of each joint measurement can be taken to be the Cartesian product of the outcome sets of the two measurements being simulated, so that each joint measurement has four outcomes. Recall that the binary-outcome measurements were denoted $\mathcal{M}_i$ with outcome m_i, and the joint measurements were denoted $\mathcal{M}^{(i,i\oplus1)}$ with outcome $(m_i, m_{i\oplus1})$. (For more details, see Appendix B.)

However, one can also imagine operationalizing the proof as an odd number n of three-outcome measurements. Denoting the ith measurement by $\mathcal{M}_i^{\mathrm{tri}}$, and taking the outcome set to be $m_i^{\mathrm{tri}} \in \{0,1,2\}$, the operational equivalence relations have the form $\forall i : [0|M_i^{\mathrm{tri}}] \simeq [2|M_{i\oplus1}^{\mathrm{tri}}]$. In words, the last outcome of one measurement in the cycle is operationally equivalent to the first outcome of the next measurement in the cycle.

To translate between the two approaches, it suffices to recognize that the three-outcome measurement $\mathcal{M}_i^{\mathrm{tri}}$ can be identified with the four-outcome measurement $\mathcal{M}^{(i,i\oplus1)}$ as long as one of the outcomes of the latter has probability zero for all preparations. Specifically, we take $(m_i = +1, m_{i\oplus1} = +1)$ to be the outcome that is always assigned zero probability,

and make the translations $m_i^{\mathrm{tri}} = 0 \leftrightarrow (m_i = +1, m_{i\oplus1} = -1)$, $m_i^{\mathrm{tri}} = 1 \leftrightarrow (m_i = -1, m_{i\oplus1} = -1)$, and $m_i^{\mathrm{tri}} = 2 \leftrightarrow (m_i = -1, m_{i\oplus1} = +1)$.

If an experiment implements n four-outcome measurements where all of the outcomes have nonzero probability, then one must use the n-cycle approach, whereas if the measurements that are implemented have only three outcomes that ever occur, then either approach can be used.

Some illustrative quantum examples help to clarify these ideas.

Recall that in KCBS's quantum construction [5], the binary-outcome measurement $\mathcal{M}_i$ is represented by a projection-valued measure (PVM) $\{\Pi_+^{(i)}, \Pi_-^{(i)}\}$ where $\Pi_+^{(i)} \equiv |l_i\rangle\langle l_i|$ and $\Pi_-^{(i)} = I - \Pi_+^{(i)}$. By construction, any two neighboring PVMs, $\{\Pi_+^{(i)}, \Pi_-^{(i)}\}$ and $\{\Pi_+^{(i\oplus1)}, \Pi_-^{(i\oplus1)}\}$, consist of projectors that commute and satisfy $\Pi_+^{(i)}\Pi_+^{(i\oplus1)} = 0$. It follows that the joint measurement $\mathcal{M}^{(i,i\oplus1)}$ is represented (uniquely) by the four-outcome PVM consisting of the products of the projectors from each of the neighboring PVMs, $\{\Pi_+^{(i)}\Pi_+^{(i\oplus1)}, \Pi_+^{(i)}\Pi_-^{(i\oplus1)}, \Pi_-^{(i)}\Pi_+^{(i\oplus1)}, \Pi_-^{(i)}\Pi_-^{(i\oplus1)}\}$. But this simplifies to $\{0, \Pi_+^{(i)}, \Pi_+^{(i\oplus1)}, I - \Pi_+^{(i)} - \Pi_+^{(i\oplus1)}\}$, so that it is clear that the first outcome always has probability zero of occurring and consequently the joint measurement is translatable into one of the three-outcome variety. This, therefore, is an example of the type described above, wherein the outcome $(m_i = +1, m_{i\oplus1} = +1)$ of the four-outcome joint measurement $\mathcal{M}^{(i,i\oplus1)}$ never occurs.

Now consider quantum realizations of the n-cycle scenario wherein the binary-outcome measurement $\mathcal{M}_i$ is *not* represented projectively, but rather by a nonprojective POVM, which we denote by $\{E_+^{(i)}, E_-^{(i)}\}$, where $E_-^{(i)} = I - E_+^{(i)}$. The constraint that $\mathcal{M}_i$ and $\mathcal{M}_{i\oplus1}$ be jointly simulatable implies that there must exist a four-outcome POVM, $\{G_{++}^{(i)}, G_{+-}^{(i)}, G_{-+}^{(i)}, G_{--}^{(i)}\}$, such that $G_{++}^{(i)} + G_{+-}^{(i)} = E_+^{(i)}$ and $G_{++}^{(i)} + G_{-+}^{(i)} = E_+^{(i\oplus1)}$.

For certain compatible pairs of nonprojective POVMs, namely, those for which $E_+^{(i)} + E_+^{(i\oplus1)} \leqslant I$, there exists a joint measurement POVM of the form $\{G_{++}^{(i)}, G_{+-}^{(i)}, G_{-+}^{(i)}, G_{--}^{(i)}\}$ where $G_{++}^{(i)} = 0$, $G_{+-}^{(i)} = E_+^{(i)}$, $G_{-+}^{(i)} = E_+^{(i\oplus1)}$, and $G_{--}^{(i)} = I - E_+^{(i)} - E_+^{(i\oplus1)}$. This constitutes another example of a four-outcome joint measurement where the first outcome never occurs, so that it is translatable into one of the three-outcome variety.

On the other hand, for *generic* compatible pairs of nonprojective POVMs, i.e., those for which it is not the case that $E_+^{(i)} + E_+^{(i\oplus1)} \leqslant I$, the joint measurement must be represented by a genuinely four-outcome POVM. It follows that if one considers a quantum realization of the n-cycle wherein the compatible pairs are of this sort, then the four-outcome joint measurement is not translatable into one of the three-outcome variety.

APPENDIX E: COMPARISON WITH THE APPROACH OF XU *ET AL.*

We here compare our approach to obtaining inequalities for universal noncontextuality from statistical proofs of the KS theorem to the one described in Xu *et al.* [12].

In fact, Ref. [12] describes *two* approaches to doing so. The first is described in Secs. IV A and IV B of that paper and the second is described in Sec. V. Neither approach, however, is presented in a manner that fully excises reference to the ideal quantum realization.

The first technique, for instance, makes explicit reference to predictions of quantum theory when determining the upper bound on their quantity $\mathcal{A}$ (in the case of the KCBS proof); they presume that their quantity $\mathcal{I}$ can achieve the maximum quantum value of $\sqrt{5}$ [their Eq. (17)].

In the second technique (also in the case of the n-cycle proof), the choice of coefficients in the inequality is particular to the case when the source outcomes are uniformly random, as in the ideal quantum realization.[10]

We are confident that generalizations of the two techniques can address these concerns, because the intuition behind them (articulated in the last paragraph of Sec. 2.2 of Ref. [12]) is in line with that of our approach. However, Xu *et al.* did not disentangle this intuition from features of the ideal quantum realization as cleanly as we do here.

We consider the main advantages of the approach described in this article, relative to those of Ref. [12], to be twofold: (i) we have derived the inequalities in a principled manner, motivating the logic with a no-go theorem for universal noncontextuality, and excising all features of the ideal quantum realization that are not needed to derive nontrivial inequalities (such as the particular choice of probabilities for source outcomes), and (ii)

_____

[10]Note that some of the commentary provided in Ref. [12] on their derivation of the inequality may create the impression that it is important that two preparation procedures (P and $\bar{P}$ in their notation) are perfectly distinguishable, so that their ontic supports are disjoint. See, e.g., the comment above Eq. (20) in their article. This idealization is not, however, required to derive the inequality.

we have described explicitly which parameters in the noncontextuality inequality depend on the choice of statistical proof and how to compute these parameters by characterizing the noncontextual measurement-assignment polytope associated with that proof.

For the case of the KS theorem based on the odd n-cycle scenario, the second technique described in Xu *et al.* leads to an inequality that is a special case of our noncontextuality inequality for this scenario when some parameters are fixed. In the rest of this section, we make the connection explicit.

First, we note that the manner in which the KCBS statistical proof of the KS theorem is operationalized in Ref. [12] differs from the manner in which we do so here in precisely the sort of way outlined in the previous appendix. Strictly speaking, therefore, their inequality is only applicable for experiments that aim to implement a set of n three-outcome measurements, $\{\mathcal{M}_i^{\mathrm{tri}}\}_i$, with operational equivalence relations of the form $\forall i :$ $[0|M_i^{\mathrm{tri}}] \simeq [2|M_{i\oplus 1}^{\mathrm{tri}}]$.

For ease of comparison with our results, however, we conceptualize these three-outcome measurements as four-outcome measurements wherein one of the outcomes never occurs, and we make the particular identification between outcomes outlined in the previous appendix. We can then rewrite their inequalities using the notational conventions of this article.

Doing so, the inequality of Ref. [12] becomes

$$\mathrm{Corr} \leqslant 1 - \frac{n}{6}\left(R - \frac{n-1}{n}\right), \tag{E1}$$

which is a special case of our inequality for the n-cycle scenario [Eq. (40)], where p_* is presumed to take the value that it takes in the ideal quantum realization of the no-go result for universal noncontextuality, namely $p_* = \frac{1}{3}$.

[1] S. Kochen and E. Specker, The problem of hidden variables in quantum mechanics, Indiana Univ. Math. J. **17**, 59 (1968).

[2] A. Peres, Two simple proofs of the Kochen-Specker theorem, J. Phys. A **24**, L175 (1991).

[3] N. D. Mermin, Hidden variables and the two theorems of John Bell, Rev. Mod. Phys. **65**, 803 (1993).

[4] A. Cabello, J. Estebaranz, and G. Garcia-Alcaine, Bell-Kochen-Specker theorem: A proof with 18 vectors, Phys. Lett. A **212**, 183 (1996).

[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] S. Yu and C. H. Oh, State-Independent Proof of Kochen-Specker Theorem with 13 Rays, Phys. Rev. Lett. **108**, 030402 (2012).

[7] A. Cabello, M. Kleinmann, and J. R. Portillo, Quantum state-independent contextuality requires 13 rays, J. Phys. A: Math. Theor. **49**, 38LT01 (2016).

[8] J. S. Bell, On the Einstein-Podolsky-Rosen paradox, Phys. **1**, 195 (1964); *Speakable and Unspeakable in Quantum Mechanics* (Cambridge University Press, New York, 1987).

[9] 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).

[10] R. Kunjwal and R. W. Spekkens, From the Kochen-Specker Theorem to Noncontextuality Inequalities Without Assuming Determinism, Phys. Rev. Lett. **115**, 110403 (2015).

[11] 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).

[12] Z-P. Xu, D. Saha, H-Y. Su, M. Pawlowski, and J-L. Chen, Reformulating noncontextuality inequalities in an operational approach, Phys. Rev. A **94**, 062103 (2016).

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

[14] 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).

[15] R. W. Spekkens, The status of determinism in proofs of the impossibility of a noncontextual model of quantum theory, Found. Phys. **44**, 1125 (2014).

[16] 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).

[17] A. Cabello, S. Severini, and A. Winter, (Non-)contextuality of physical theories as an axiom, arXiv:1010.2163; Graph-Theoretic Approach to Quantum Correlations, Phys. Rev. Lett. **112**, 040401 (2014).

[18] A. Cabello, S. Filipp, H. Rauch, and Y. Hasegawa, Proposed Experiment for Testing Quantum Contextuality with Neutrons, Phys. Rev. Lett. **100**, 130404 (2008).

[19] H. Bartosik, J. Klepp, C. Schmitzer, S. Sponar, A. Cabello, H. Rauch, and Y. Hasegawa, Experimental Test of Quantum Contextuality in Neutron Interferometry, Phys. Rev. Lett. **103**, 040403 (2009).

[20] R. Lapkiewicz, P. Li, C. Schaeff, N. K. Langford, S. Ramelow, M. Wieśniak, and A. Zeilinger, Experimental non-classicality of an indivisible quantum system, Nature (London) **474**, 490 (2011).

[21] C. Zu, Y.-X. Wang, D.-L. Deng, X.-Y. Chang, K. Liu, P.-Y. Hou, H.-X. Yang, and L.-M. Duan, State-Independent Experimental Test of Quantum Contextuality in an Indivisible System, Phys. Rev. Lett. **109**, 150401 (2012).

[22] F. M. Leupold, M. Malinowski, C. Zhang, V. Negnevitsky, J. Alonso, A. Cabello, and J. P. Home, Sustained state-independent quantum contextual correlations from a single ion, arXiv:1706.07370.

[23] R. Kunjwal, Beyond the Cabello-Severini-Winter framework: Making sense of contextuality without sharpness of measurements, arXiv:1709.01098.

[24] R. Kunjwal, Fine's theorem, noncontextuality, and correlations in Specker's scenario, Phys. Rev. A **91**, 022108 (2015).

[25] 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).

[26] S. Popescu and D. Rohrlich, Quantum nonlocality as an axiom, Found. Phys. **24**, 379 (1994).

[27] A. Fine, Hidden Variables, Joint Probability, and the Bell Inequalities, Phys. Rev. Lett. **48**, 291 (1982).