Contextual Advantage for State Discrimination

David Schmid[*] and Robert W. Spekkens

Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo N2L 2Y5, Ontario, Canada
and Institute for Quantum Computing and Department of Physics and Astronomy,
University of Waterloo, Waterloo N2L 3G1, Ontario, Canada

Finding quantitative aspects of quantum phenomena which cannot be explained by any classical model has foundational importance for understanding the boundary between classical and quantum theory. It also has practical significance for identifying information processing tasks for which those phenomena provide a quantum advantage. Using the framework of generalized noncontextuality as our notion of classicality, we find one such nonclassical feature within the phenomenology of quantum minimum-error state discrimination. Namely, we identify quantitative limits on the success probability for minimum-error state discrimination in any experiment described by a noncontextual ontological model. These constraints constitute noncontextuality inequalities that are violated by quantum theory, and this violation implies a quantum advantage for state discrimination relative to noncontextual models. Furthermore, our noncontextuality inequalities are robust to noise and are operationally formulated, so that any experimental violation of the inequalities is a witness of contextuality, independently of the validity of quantum theory. Along the way, we introduce new methods for analyzing noncontextuality scenarios and demonstrate a tight connection between our minimum-error state discrimination scenario and a Bell scenario.

Subject Areas: Quantum Physics, Quantum Information

I. INTRODUCTION

Understanding the boundary between the quantum and the classical is of fundamental importance for understanding quantum theory. One successful metric for nonclassicality, violation of Bell's notion of local causality [1], defines a clear departure from classicality in relativistic theories, but is relevant only for experiments with spacelike separated measurements. Another notion of classicality, which concerns context independence, was proposed by Kochen-Specker [2] and Bell [3], and has since been significantly refined and generalized [4]. It is the generalized notion of noncontextuality from Ref. [4] which we study in this paper, but we refer to it simply as "noncontextuality" hereafter. As a metric for nonclassicality, the failure of noncontextuality has a broader scope than the failure of local causality insofar as it does not require spacelike separation. It has also been shown to subsume many other preexisting notions of nonclassicality, such as the negativity of quasiprobability representations [5], the generation of anomalous weak values [6], and even the aforementioned violations of local causality [4].

The quantum-classical boundary is also of practical importance in identifying tasks which admit of a quantum advantage. For example, violations of Bell inequalities have been shown to be a resource for device-independent key distribution [7], certified randomness [8], and communication complexity [9]. The failure of noncontextuality has also been shown to be a resource, leading to advantages for cryptography [10–12] and computation [13–15].

We here analyze minimum-error state discrimination (MESD) from the point of view of noncontextuality. Quantum state discrimination is a task wherein one must guess which quantum state describes a given quantum system when the state of that system is drawn from a known set of possibilities with a known prior distribution, and the estimation is based on the outcome of a measurement of one's choosing. In the "minimum-error" variety of state discrimination, the objective is to minimize the probability that the estimate is in error. We here focus on the simplest case of a set containing just two states having equal *a priori* probability.

Although it is common to assert that the impossibility of perfectly discriminating nonorthogonal quantum states is an intrinsically nonclassical effect, this claim does not meet the minimal standard that one should require of *any* claim that some operational feature of quantum theory cannot be explained classically: namely, that it be justified by a rigorous no-go theorem. Such a theorem articulates a principle of classicality which has implications for operational statistics, and then proves that these implications are

[*]dschmid@perimeterinstitute.ca

inconsistent with some operational feature(s) of quantum theory. Because the principle of noncontextuality constrains operational statistics and also has very broad scope, it is a particularly useful notion of classicality. If one does take it as one's principle of classicality, then the impossibility of discriminating nonorthogonal pure quantum states *cannot* be considered a nonclassical effect because there are subtheories of quantum theory (containing a strict subset of the states, measurements, and transformations of the full theory) [16] wherein this phenomenon arises and which admit of a noncontextual model. (Within such models, the phenomenon can be attributed to the fact that the probability distributions associated to such quantum states are overlapping [17].) It follows that one must look at more nuanced aspects of the phenomenology of quantum state discrimination to identify features which are truly nonclassical according to this principle.

We identify one such strongly nonclassical aspect of minimum-error state discrimination: the particular dependence of the probability of successful discrimination on the overlap of the quantum states. For a given overlap, the quantum probability of discrimination is larger than can be accounted for by a noncontextual model. After presenting this result as a no-go theorem—that no noncontextual model can reproduce certain features of quantum MESD—we reformulate the problem in a manner which makes no reference to quantum theory, and which does not rely on any theoretical idealizations such as noise-free measurements or preparations. Our entirely operational formulation allows us to derive inequalities which can experimentally witness a contextual advantage for state discrimination, in the presence of noise and independently of the validity of quantum theory.

Our result identifies a key feature of quantum state discrimination which cannot be understood in any noncontextual model, and hence which is strongly nonclassical. Because quantum state discrimination is a primitive in many important quantum information processing protocols [19,20], this work constitutes a first step towards identifying contextuality as a resource for more tasks concerning communication, computation, and cryptography.

We also prove an isomorphism between our operational MESD scenario and a two-party Bell test in which one party performs one of a pair of binary-outcome measurements and the other performs one of three binary-outcome measurements. This is similar to the fact that the noncontextuality inequality delimiting the success rate for parity-oblivious multiplexing [10] is isomorphic to the Clauser-Horne-Shimony-Holte (CHSH) inequality in the Bell scenario [10].

Finally, we introduce two powerful new technical tools. First, we generalize existing methods for simulating exact operational equivalences [21]. Namely, while Ref. [21] shows how one may find a set of procedures which respects certain operational equivalences exactly, we further

demonstrate that one can find procedures which respect operational equivalences and *simultaneously* obey useful auxiliary constraints, such as the symmetries native to our ideal MESD scenario. This tool may have more general applications in the comparison of experimental data with theoretical expectations. More importantly, we find our noncontextuality inequalities using a novel algorithm (presented in Appendix B) for deriving the full set of necessary and sufficient noncontextuality inequalities for *any* finite prepare-and-measure scenario, with respect to any fixed operational equivalences (a full description of this algorithm can be found in Ref. [22]).

II. OPERATIONAL THEORIES AND ONTOLOGICAL MODELS

An operational theory is a specification of sets of primitive laboratory operations (e.g., preparations and measurements) and a prescription for finding the probabilities $p(k|M, P)$ for each outcome k given any measurement M performed on any preparation P. Two preparations P and P' are termed *operationally equivalent* if they cannot be differentiated by the statistics of any measurement; we denote this operational equivalence by

$$P \simeq P'. \tag{1}$$

In this article, quantum theory is understood as an operational theory. In the quantum formalism, the density operator specifies the statistics for all measurements, so that two preparation procedures are operationally equivalent if and only if they are represented by the same density operator.

An ontological model of an operational theory has the following form. To every system, there is associated an ontic state space Λ, where each ontic state $\lambda \in \Lambda$ specifies all the physical properties of the system. Each preparation P of a system is presumed to sample the system's ontic state λ at random from a probability distribution, denoted $\mu_P(\lambda)$ and termed the *epistemic state* associated to P, where

$$\forall \lambda \colon\ 0 \leq \mu_P(\lambda), \tag{2}$$

$$\int_\Lambda d\lambda \mu_P(\lambda) = 1. \tag{3}$$

Each measurement M on a system is presumed to have its outcome k sampled at random in a manner that depends on the ontic state λ. The term *effect* is used to refer to the pair consisting of a measurement M, together with one of its outcomes k, and is denoted by $k|M$. The probability of outcome k given measurement M, considered as a function of λ, is termed the *response function* associated to $k|M$, and denoted $\xi_{k|M}(\lambda)$, where

$$\forall \lambda, \forall k \colon\ 0 \leq \xi_{k|M}(\lambda), \tag{4}$$

$$\forall \lambda : \sum_k \xi_{k|M}(\lambda) = 1. \qquad (5)$$

Finally, an ontological model of an operational theory must reproduce the latter's empirical predictions; that is,

$$p(k|M, P) = \int_\Lambda d\lambda\, \xi_{k|M}(\lambda)\mu_P(\lambda). \qquad (6)$$

We are now in a position to describe the assumption of preparation noncontextuality defined in Ref. [4]. An ontological model is said to be *preparation noncontextual* if it assigns the same epistemic state to all operationally equivalent preparations [4]:

$$P \simeq P' \Rightarrow \mu_P(\lambda) = \mu_{P'}(\lambda). \qquad (7)$$

In operational quantum theory, the principle of preparation noncontextuality is respected whenever any two preparations that are associated to the same density operator are represented by the same epistemic state. For instance, different ensembles of states that average to the same mixed state (and for which one discards the information about which element of the ensemble was prepared) are operationally equivalent, and must be assigned the same epistemic state in a preparation-noncontextual model.

Although there is a corresponding notion of measurement noncontextuality (namely, that operationally equivalent outcomes of measurements are represented by the same response functions), we do not require its use in this article.

A few terminological conventions are useful. A measurement is said to be represented as *outcome-deterministic* in the ontological model if the associated response functions all take values in $\{0, 1\}$. The support of an epistemic state is defined as the set of $\lambda \in \Lambda$ which are assigned nonzero probability by it, $\mathrm{supp}[\mu_P(\lambda)] \equiv \{\lambda : \mu_P(\lambda) \neq 0\}$, while the support of a response function is defined as the set of $\lambda \in \Lambda$ for which the response function is nonzero, $\mathrm{supp}[\xi_{k|M}(\lambda)] \equiv \{\lambda : \xi_{k|M}(\lambda) \neq 0\}$.

III. QUANTUM MINIMUM-ERROR STATE DISCRIMINATION

We begin with the problem of discriminating two nonorthogonal pure quantum states $|\phi\rangle$ and $|\psi\rangle$. These two states span a two-dimensional space, so we can represent them as points in an equatorial plane of the Bloch ball, as in Fig. 1.

First, we consider the operational signature of their nonorthogonality. A measurement of the ϕ basis, $B_\phi \equiv \{|\phi\rangle\langle\phi|, |\bar\phi\rangle\langle\bar\phi|\}$, perfectly distinguishes between state $|\phi\rangle$ and its complement; we denote the associated outcomes by ϕ and $\bar\phi$, respectively. A measurement of the ψ basis, $B_\psi \equiv \{|\psi\rangle\langle\psi|, |\bar\psi\rangle\langle\bar\psi|\}$, does the same for the state $|\psi\rangle$ and its complement, with associated outcomes ψ and $\bar\psi$. If one implements the ψ basis measurement on the state ϕ, the probability of obtaining the ψ outcome is
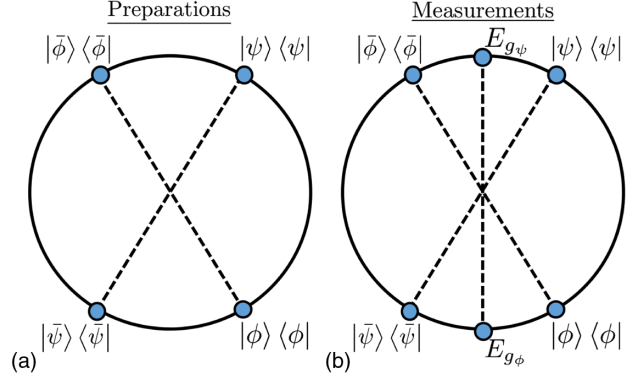

Preparations   Measurements

FIG. 1. The quantum states and measurements in our scenario, depicted as Bloch vectors in an equatorial plane of the Bloch ball.

$$c_q = \mathrm{Tr}[|\phi\rangle\langle\phi|\psi\rangle\langle\psi|] = |\langle\phi|\psi\rangle|^2. \qquad (8)$$

Because one could think of this quantity as the probability that ϕ passes the test for ψ and thus is confusable with ψ, we henceforth call it the *confusability*. Note that if one implements the ϕ basis measurement on the state ψ, the probability of obtaining the ϕ outcome is also c_q.

Second, we consider the operational probability of success at MESD for these states. If $|\phi\rangle$ and $|\psi\rangle$ have nonzero confusability (i.e., if they are not orthogonal), then no measurement can distinguish between the two without incurring a nonzero probability of error. We denote the discriminating measurement by $B_d \equiv \{E_{g_\phi}, E_{g_\psi}\}$, where the outcome for which one should guess ϕ (ψ) is denoted g_ϕ (g_ψ). Assuming equal prior probabilities of $|\phi\rangle$ and $|\psi\rangle$, the probability of guessing the state correctly with this measurement is

$$s_q \equiv \frac{1}{2}\mathrm{Tr}[E_{g_\phi}|\phi\rangle\langle\phi|] + \frac{1}{2}\mathrm{Tr}[E_{g_\psi}|\psi\rangle\langle\psi|]. \qquad (9)$$

We assume that the discriminating measurement has the natural symmetry property $\mathrm{Tr}[E_{g_\phi}|\phi\rangle\langle\phi|] = \mathrm{Tr}[E_{g_\psi}|\psi\rangle\langle\psi|]$, so that

$$s_q = \mathrm{Tr}[E_{g_\phi}|\phi\rangle\langle\phi|] = \mathrm{Tr}[E_{g_\psi}|\psi\rangle\langle\psi|]. \qquad (10)$$

The measurement scheme that yields the greatest probability of guessing correctly which of two nonorthogonal states was prepared is called the minimum-error state discrimination scheme. Since $|\phi\rangle$ and $|\psi\rangle$ are prepared with equal probability, the measurement $\{E_{g_\phi}, E_{g_\psi}\}$ achieving MESD is the one consisting of projectors onto the basis that straddles $|\phi\rangle$ and $|\psi\rangle$ in Hilbert space, which is depicted in the Bloch sphere in Fig. 1. This is called the Helstrom measurement [23]. It is well known that the probability of guessing the state correctly using the Helstrom measurement is

$$s_q = \frac{1}{2}\left(1 + \sqrt{1 - |\langle \phi | \psi \rangle|^2}\right) = \frac{1}{2}\left(1 + \sqrt{1 - c_q}\right). \quad (11)$$

We have now described all of the preparations and measurements that usually appear in a discussion of the problem of discriminating two nonorthogonal quantum states, and some basic facts about the relations that hold among the operational quantities characterizing the discrimination problem (i.e., facts about the phenomenology of quantum state discrimination). However, these facts are insufficient for deriving a no-go theorem for noncontextuality. The reason is that the preparations and measurements described thus far do not exhibit any operational equivalences via which the assumption of noncontextuality could imply nontrivial constraints on the ontological model.

However, there is a simple solution: we also consider the problem of discriminating the pair of quantum states that are complementary to $|\phi\rangle$ and $|\psi\rangle$, namely, $|\bar{\phi}\rangle$ and $|\bar{\psi}\rangle$, also depicted in Fig. 1. By symmetry, the confusability of $|\bar{\phi}\rangle$ and $|\bar{\psi}\rangle$ is also equal to c_q, and the success rate for distinguishing $|\bar{\phi}\rangle$ and $|\bar{\psi}\rangle$ when they have equal prior probability is also equal to s_q (where the optimal measurement is again $\{E_{g_\phi}, E_{g_\psi}\}$, but now the outcomes g_ϕ and g_ψ signal one to guess preparations $|\bar{\psi}\rangle$ and $|\bar{\phi}\rangle$, respectively). So the $|\bar{\phi}\rangle$ vs $|\bar{\psi}\rangle$ discrimination problem is a mirror image of the $|\phi\rangle$ vs $|\psi\rangle$ discrimination problem, and consequently does not require specifying any additional facts about the phenomenology of quantum state discrimination. However, the inclusion of $|\bar{\phi}\rangle$ and $|\bar{\psi}\rangle$ in our analysis provides us with a nontrivial operational equivalence relation among the preparations, namely,

$$\frac{1}{2}|\phi\rangle\langle\phi| + \frac{1}{2}|\bar{\phi}\rangle\langle\bar{\phi}| = \frac{1}{2}|\psi\rangle\langle\psi| + \frac{1}{2}|\bar{\psi}\rangle\langle\bar{\psi}| = \frac{\mathbb{1}}{2}. \quad (12)$$

We show that this equivalence relation together with the phenomenology of quantum state discrimination described above is sufficient to derive a no-go theorem for noncontextuality.

The probability of a given measurement outcome occurring on a given preparation, for every possible pairing thereof, is summarized in Table I. Here, the columns correspond to the distinct state preparations and the rows correspond to the distinct effects (where one need only include a single effect for each binary-outcome

TABLE I. Data table in the ideal quantum case.

	$|\phi\rangle$	$|\psi\rangle$	$|\bar{\phi}\rangle$	$|\bar{\psi}\rangle$
$|\phi\rangle\langle\phi|$	1	c_q	0	$1 - c_q$
$|\psi\rangle\langle\psi|$	c_q	1	$1 - c_q$	0
E_{g_ϕ}	s_q	$1 - s_q$	$1 - s_q$	s_q

measurement given that the probability for the other effect is fixed by normalization).

IV. NONCONTEXTUALITY NO-GO THEOREM FOR MESD IN QUANTUM THEORY

The fact that the ontological model must reproduce the probabilities in Table I via Eq. (6) implies constraints on the epistemic states associated to the four preparations and the response functions associated to the three effects. For instance, to reproduce the first column of the table, one requires that

$$\int_\Lambda d\lambda \, \xi_{\phi|B_\phi}(\lambda) \mu_\phi(\lambda) = 1, \quad (13)$$

$$\int_\Lambda d\lambda \, \xi_{\psi|B_\psi}(\lambda) \mu_\phi(\lambda) = c_q, \quad (14)$$

$$\int_\Lambda d\lambda \, \xi_{g_\phi|B_d}(\lambda) \mu_\phi(\lambda) = s_q. \quad (15)$$

Given that convex mixtures of preparations are represented in an ontological model by the corresponding mixture of epistemic states [see Eq. (7) of Ref. [5] and the surrounding discussion], it follows that $\frac{1}{2}|\phi\rangle\langle\phi| + \frac{1}{2}|\bar{\phi}\rangle\langle\bar{\phi}|$ is represented by $\frac{1}{2}\mu_\phi(\lambda) + \frac{1}{2}\mu_{\bar{\phi}}(\lambda)$, and $\frac{1}{2}|\psi\rangle\langle\psi| + \frac{1}{2}|\bar{\psi}\rangle\langle\bar{\psi}|$ is represented by $\frac{1}{2}\mu_\psi(\lambda) + \frac{1}{2}\mu_{\bar{\psi}}(\lambda)$. But because both of these mixtures of preparations are associated to the completely mixed state [Eq. (12)], they are operationally equivalent, and thus by the assumption of preparation noncontextuality, they are represented by the same epistemic state. It follows that

$$\frac{1}{2}\mu_\phi(\lambda) + \frac{1}{2}\mu_{\bar{\phi}}(\lambda) = \frac{1}{2}\mu_\psi(\lambda) + \frac{1}{2}\mu_{\bar{\psi}}(\lambda). \quad (16)$$

Any ontological model satisfying noncontextuality, and consequently Eq. (16), and reproducing the form of the data in Table I, and consequently Eqs. (13)–(15) and their kin, can be shown to satisfy the following trade-off between s_q and c_q:

$$s_q \leq 1 - \frac{c_q}{2}. \quad (17)$$

An intuitive proof is provided in Section IV A, where we also discuss how this result is related to the results of Refs. [24–26]. (In Appendix A, we provide a proof using more general methods, which generalizes more easily to the noisy case discussed later, in Sec. VI.)

This trade-off relation contradicts the one known to be optimal in quantum theory, Eq. (11). The optimal quantum trade-off generally allows *higher* success rates for a given confusability than the noncontextual trade-off. Therefore, we conclude that the phenomenology of minimum-error
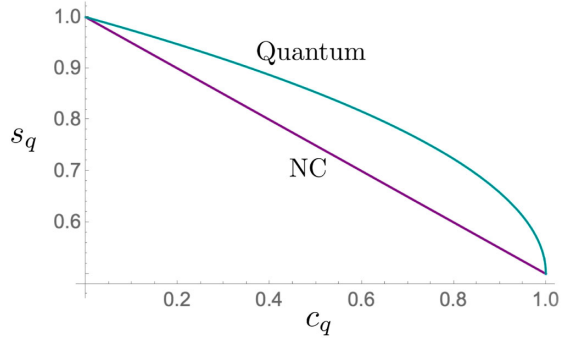
FIG. 2. Optimal trade-off for a noncontextual (NC) model (purple line) and for quantum theory (light blue curve).

state discrimination in the noiseless quantum case is inconsistent with the principle of noncontextuality.

In Fig. 2, we plot the maximum success rate for MESD as a function of the confusability for both quantum theory [Eq. (11)] and for a noncontextual model [the trade-off that saturates the inequality of Eq. (17)].

A. Intuitive proof of the noncontextual trade-off

We now introduce some basic facts from classical probability theory, which we then leverage to prove Eq. (17).

Suppose that a classical variable λ has been sampled from one of two overlapping probability distributions, $p(\lambda|a)$ and $p(\lambda|b)$. Absent additional information, it is straightforward to see that in trying to guess which of the two distributions a given λ was drawn from, one cannot do better than guessing "distribution a" for the values of λ for which $p(a|\lambda) > p(b|\lambda)$, and guessing "distribution b" when the opposite is true. [Of course, it is irrelevant what one guesses for the values of λ for which $p(a|\lambda) = p(b|\lambda)$.] In the special case we are considering, with equal prior probability $p(a) = p(b) = \frac{1}{2}$ for the two options, if we perform a Bayesian inversion, we find $p(\lambda|a) > p(\lambda|b)$ if and only if $p(a|\lambda) > p(b|\lambda)$, and hence one should guess distribution a for the values of λ for which $p(\lambda|a) > p(\lambda|b)$, and guess distribution b when the opposite is true.

The probability that the guess $g \in \{a, b\}$ was correct given a particular value of λ is simply $p(g|\lambda)$. Since we always guess the distribution a or b that has the higher likelihood of being correct, the probability that we are right in each run is simply $\max\{p(a|\lambda), p(b|\lambda)\}$. On average, then, the success probability r is

$$r = \int_\Lambda d\lambda\, p(\lambda) \max\{p(a|\lambda), p(b|\lambda)\} \tag{18}$$

$$= \int_\Lambda d\lambda\, p(\lambda)[1 - \min\{p(a|\lambda), p(b|\lambda)\}] \tag{19}$$

$$= 1 - \int_\Lambda d\lambda \min\{p(a|\lambda)p(\lambda), p(b|\lambda)p(\lambda)\} \tag{20}$$

$$= 1 - \int_\Lambda d\lambda \min\{p(\lambda|a)p(a), p(\lambda|b)p(b)\} \tag{21}$$

$$= 1 - \frac{1}{2}\int_\Lambda d\lambda \min\{p(\lambda|a), p(\lambda|b)\}, \tag{22}$$

where the equality in Eq. (19) uses the fact that $p(a|\lambda) + p(b|\lambda) = 1$ for all λ. The quantity $\int_\Lambda d\lambda \min\{p(\lambda|a), p(\lambda|b)\}$ is termed the *classical overlap* of the probability distributions $p(\lambda|a)$ and $p(\lambda|b)$.

In a MESD scenario, the task is to guess, in each particular run of the experiment, whether a system was prepared by state preparation $|\phi\rangle$ or by state preparation $|\psi\rangle$. If the experiment is described by an ontological model, then this task corresponds to guessing, from a single sample of the ontic state λ of the system, whether it was sampled from the distribution $\mu_\phi(\lambda)$ or from $\mu_\psi(\lambda)$. Given that we do not assume any operational equivalence relations among the measurements in the experiment, the assumption of measurement noncontextuality does not place any constraints on the ontological representation of the measurements. Therefore, in particular, the Helstrom measurement is *at best* represented in the ontological model by the set of response functions that yield the maximum probability of guessing which distribution the ontic state λ was sampled from. From our discussion concerning two overlapping classical probability distributions, it is clear that this corresponds to a measurement that returns the g_ϕ outcome whenever $\mu_\phi(\lambda) > \mu_\psi(\lambda)$ and the g_ψ outcome whenever $\mu_\phi(\lambda) < \mu_\psi(\lambda)$, and that the probability of guessing correctly based on the outcome of the Helstrom measurement is upper bounded as follows [27]:

$$s_q \le 1 - \frac{1}{2}\int_\Lambda d\lambda \min\{\mu_\phi(\lambda), \mu_\psi(\lambda)\}. \tag{23}$$

We now show that in a noncontextual model,

$$c_q = \int_\Lambda d\lambda \min\{\mu_\phi(\lambda), \mu_\psi(\lambda)\}, \tag{24}$$

so that substituting Eq. (24) into Eq. (23), we infer that $s_q \le 1 - (c_q/2)$, the noncontextual bound on the trade-off between s_q and c_q described in Eq. (17).

First, in any preparation-noncontextual model the response function $\xi_i(\lambda)$ for a projector onto pure state $|i\rangle$ satisfies

$$\xi_i(\lambda) = \begin{cases} 1 & \text{if } \lambda \in \text{supp}[\mu_i(\lambda)] \\ 0 & \text{otherwise.} \end{cases} \tag{25}$$

This outcome determinism for sharp measurements was first proven in Ref. [4]. It can be seen by considering the projector as part of some projective measurement M with effects $\{E_i = |i\rangle\langle i|\}$, and the corresponding basis of pure

states $\{\rho_i = |i\rangle\langle i|\}$, so that $\mathrm{Tr}[E_i\rho_j] = \delta_{i,j}$. Denoting the epistemic state of ρ_j as $\mu_j(\lambda)$ and the response function for E_i as $\xi_{i|M}(\lambda)$, this implies that $\int \mu_j(\lambda)\xi_{i|M}(\lambda)d\lambda = \delta_{i,j}$. Because $\mu_j(\lambda)$ is a normalized probability distribution, this implies that, for *any* ontological model,

$$\xi_{i|M}(\lambda) = \begin{cases} 1 & \text{if } \lambda \in \mathrm{supp}[\mu_i(\lambda)] \\ 0 & \text{if } \lambda \in \mathrm{supp}[\mu_{j\neq i}(\lambda)]. \end{cases} \quad (26)$$

Equation (26) is not equivalent to Eq. (25), since there may exist ontic states that are not in the support of *any* of the $\mu_i(\lambda)$, and Eq. (26) does not constrain such ontic states in any way. In a preparation-noncontextual model, however, we can furthermore show that there are no ontic states outside of the union of the supports of the set of basis states, $\cup_i\mathrm{supp}[\mu_i(\lambda)]$, as follows. Every density operator ρ appears in *some* decomposition of the maximally mixed state $(1/d)\mathbb{1}$. By preparation noncontextuality, every such decomposition has the *same* distribution $\mu_{(1/d)\mathbb{1}}(\lambda)$ over ontic states. Thus, every ontic state in the support of the corresponding $\mu_\rho(\lambda)$ *also* appears in the support of $\mu_{(1/d)\mathbb{1}}(\lambda)$, so the full state space Λ is equivalent to $\mathrm{supp}[\mu_{(1/d)\mathbb{1}}(\lambda)]$. Furthermore, for the basis of states $\{\rho_i\}$ above, $(1/d)\sum_i\rho_i = (1/d)\mathbb{1}$, so preparation noncontextuality implies that $\sum_i(1/d)\mu_i(\lambda) = \mu_{(1/d)\mathbb{1}}(\lambda)$, and therefore $\cup_i\mathrm{supp}[\mu_i(\lambda)] = \mathrm{supp}[\mu_{(1/d)\mathbb{1}}(\lambda)] = \Lambda$. Thus, every ontic state λ must be in the support of exactly one of the ρ_i, and Eq. (26) can be strengthened to Eq. (25).

Recalling the expression for the confusability of quantum states $|\phi\rangle$ and $|\psi\rangle$ in an ontological model, $c_q = \int_\Lambda d\lambda\xi_{\phi|B_\phi}(\lambda)\mu_\psi(\lambda)$, Eq. (25) implies that for a preparation-noncontextual model

$$c_q = \int_{\mathrm{supp}[\mu_\phi(\lambda)]} d\lambda\mu_\psi(\lambda). \quad (27)$$

By virtue of the symmetry of the problem, the analogous expression with the roles of ϕ and ψ reversed also holds. The fact that the expression for the ideal confusability $c_q = |\langle\phi|\psi\rangle|^2$ of ϕ and ψ in a preparation-noncontextual model is given by Eq. (27) was noted by Leifer and Maroney [25].

The second implication of preparation noncontextuality which we require to prove Eq. (24) is that for each of the four quantum states $\Psi \in \{\phi, \psi, \bar{\phi}, \bar{\psi}\}$, $\mu_\Psi(\lambda) = 2\mu_{\frac{1}{2}}(\lambda)$ for all $\lambda \in \mathrm{supp}[\mu_\Psi(\lambda)]$, where $\mu_{\frac{1}{2}}(\lambda)$ is the distribution associated with the maximally mixed state $\frac{1}{2}$. This was also first proven in Ref. [4], and follows immediately from preparation noncontextuality, $\frac{1}{2}\mu_\phi(\lambda)+\frac{1}{2}\mu_{\bar{\phi}}(\lambda) = \frac{1}{2}\mu_\phi(\lambda)+\frac{1}{2}\mu_{\bar{\phi}}(\lambda) = \mu_{\frac{1}{2}}(\lambda)$, and the fact that an ontic state can be in the support of at most one state from a set of orthogonal states; that is, $\mu_\phi(\lambda)\mu_{\bar{\phi}}(\lambda) = 0$ and $\mu_\psi(\lambda)\mu_{\bar{\psi}}(\lambda) = 0$.

Hence, for all $\lambda \in \mathrm{supp}[\mu_\phi(\lambda)] \cap \mathrm{supp}[\mu_\psi(\lambda)]$, we have $\mu_\phi(\lambda) = \mu_\psi(\lambda) = 2\mu_{\frac{1}{2}}(\lambda)$. It follows that $\min\{\mu_\phi(\lambda), \mu_\psi(\lambda)\} = \mu_\phi(\lambda) = \mu_\psi(\lambda)$ for all $\lambda \in \mathrm{supp}[\mu_\phi(\lambda)] \cap \mathrm{supp}[\mu_\psi(\lambda)]$, and is equal to 0 everywhere else, and consequently

$$\int_{\mathrm{supp}[\mu_\phi(\lambda)]} d\lambda\mu_\psi(\lambda)$$

$$= \int_{\mathrm{supp}[\mu_\phi(\lambda)]\cap\mathrm{supp}[\mu_\psi(\lambda)]} d\lambda\mu_\psi(\lambda)$$

$$= \int_{\mathrm{supp}[\mu_\phi(\lambda)]\cap\mathrm{supp}[\mu_\psi(\lambda)]} d\lambda \min\{\mu_\phi(\lambda), \mu_\psi(\lambda)\}$$

$$= \int_\Lambda d\lambda \min\{\mu_\phi(\lambda), \mu_\psi(\lambda)\}. \quad (28)$$

Finally, Eqs. (27) and (28) together imply Eq. (24), which is what we sought to prove.

B. Graphical summary of the proof

The intuitive proof is best summarized graphically, by contrasting a preparation-contextual ontological model, Fig. 3, with a preparation noncontextual ontological model, Fig. 4. For visual simplicity, we have chosen a continuous, one-dimensional, bounded ontic state space. We arrange
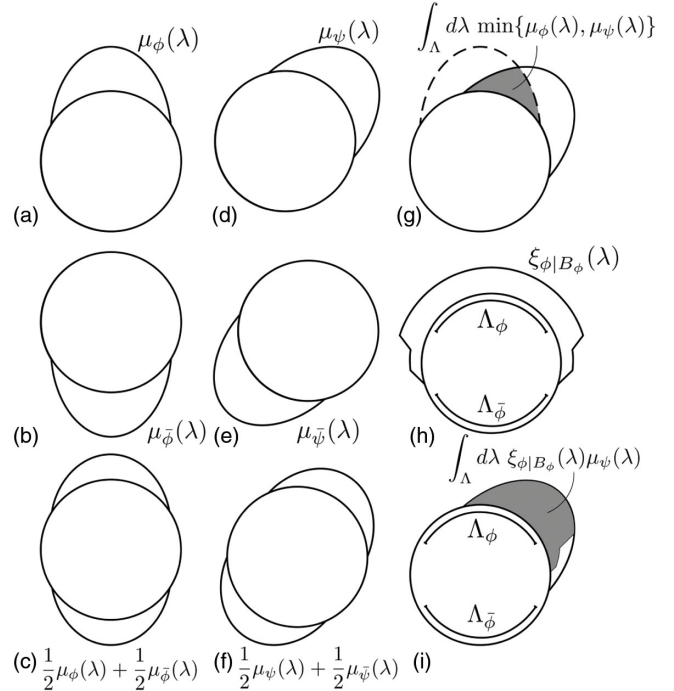

FIG. 3. In a contextual model of a MESD scenario: (a)–(f) epistemic states, (g) classical overlap between $\mu_\phi(\lambda)$ and $\mu_\psi(\lambda)$, (h) response function $\xi_{\phi|B_\phi}(\lambda)$, with indication of $\Lambda_\phi \equiv \mathrm{supp}[\mu_\phi(\lambda)]$ and $\Lambda_{\bar{\phi}} \equiv \mathrm{supp}[\mu_{\bar{\phi}}(\lambda)]$, (i) confusability defined by $\xi_{\phi|B_\phi}(\lambda)$, also with indication of Λ_ϕ and $\Lambda_{\bar{\phi}}$.
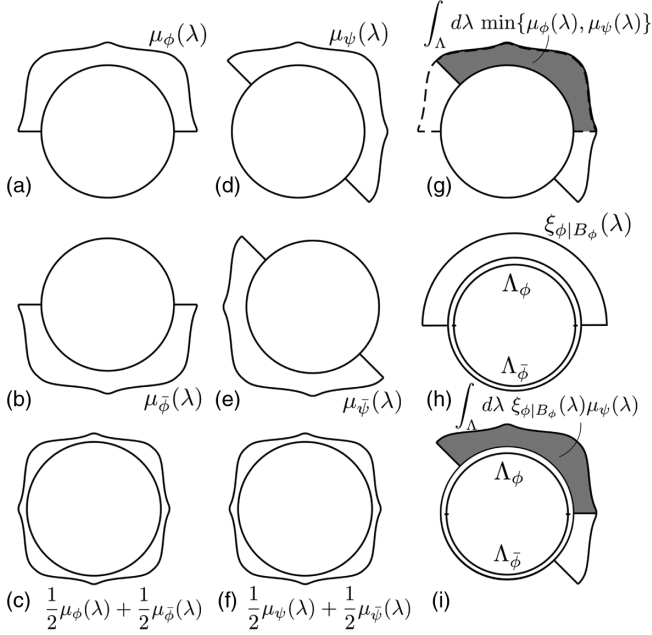
FIG. 4. In a noncontextual model of an MESD scenario: (a)–(f) epistemic states, (g) classical overlap between $\mu_\phi(\lambda)$ and $\mu_\psi(\lambda)$, (h) response function $\xi_{\phi|B_\phi}(\lambda)$, with indication of $\Lambda_\phi \equiv \mathrm{supp}[\mu_\phi(\lambda)]$ and $\Lambda_{\bar{\phi}} \equiv \mathrm{supp}[\mu_{\bar{\phi}}(\lambda)]$, (i) confusability defined by $\xi_{\phi|B_\phi}(\lambda)$, also with indication of Λ_ϕ and $\Lambda_{\bar{\phi}}$.

the state space into a circle, so that each point on the circle is a unique ontic state, and epistemic states are represented as probability distributions on the surface of the circle (where the probability density corresponds to the radial height). In each figure, we show the epistemic states for the four preparations and for the two mixed preparations, the classical overlap for two epistemic states, a representative response function, and the confusability generated by that response function. We then show that in the contextual model, the classical overlap and confusability can differ, while in the noncontextual model, they must be identical.

In the ontological model of a MESD scenario shown in Fig. 3, the distributions $\frac{1}{2}\mu_\phi(\lambda) + \frac{1}{2}\mu_{\bar{\phi}}(\lambda)$ and $\frac{1}{2}\mu_\psi(\lambda) + \frac{1}{2}\mu_{\bar{\psi}}(\lambda)$ are not identical; hence, this model is preparation-contextual. The classical overlap $\int_\Lambda d\lambda \min\{\mu_\phi(\lambda), \mu_\psi(\lambda)\}$ is equal to the area of the shaded region in Fig. 3(g). The response function $\xi_{\phi|B_\phi}(\lambda)$ must have value 0 on the support of $\mu_{\bar{\phi}}(\lambda)$ and value 1 on the support of $\mu_\phi(\lambda)$, as pictured in Fig. 3(h); however, in the region outside both of these supports, its value is arbitrary, as indicated schematically. Given the response function pictured, the confusability $c_q = \int_\Lambda d\lambda \xi_{\phi|B_\phi}(\lambda)\mu_\psi(\lambda)$ equals the area of the shaded region in Fig. 3(i). One can clearly see that the classical overlap and the confusability need not be the same in a preparation-contextual model.

In the ontological model of a MESD scenario shown in Fig. 4, the distributions $\frac{1}{2}\mu_\phi(\lambda) + \frac{1}{2}\mu_{\bar{\phi}}(\lambda)$ and $\frac{1}{2}\mu_\psi(\lambda) + \frac{1}{2}\mu_{\bar{\psi}}(\lambda)$ *are* identical; hence, this model is preparation-

noncontextual. Furthermore, these two distributions are equal to the unique distribution $\mu_{\frac{1}{2}}(\lambda)$ (whose support must span the entire ontic state space), and the epistemic states $\mu_\phi(\lambda)$, $\mu_{\bar{\phi}}(\lambda)$, $\mu_\psi(\lambda)$, and $\mu_{\bar{\psi}}(\lambda)$ must both be equal on their support to $2\mu_{\frac{1}{2}}(\lambda)$. Thus, in a preparation-noncontextual model, the classical overlap is given simply by the integral of $2\mu_{\frac{1}{2}}(\lambda)$ in the region of common support, as shown by the shaded region in Fig. 4(g). Furthermore, preparation noncontextuality implies that the response function $\xi_{\phi|B_\phi}(\lambda)$ is 1 on the support of $\mu_\phi(\lambda)$ and 0 on *all* other ontic states, as shown in Fig. 4(h). Given this form for the response function, the confusability $c_q = \int_\Lambda d\lambda \xi_{\phi|B_\phi}(\lambda)\mu_\psi(\lambda)$ is given by the area of the shaded region in Fig. 4(i). Clearly, the classical overlap and the confusability are identical in a preparation-noncontextual model.

C. Relation to previous work

Leifer and Maroney [25] consider the assumption that Eq. (27) should hold for *every* possible pair of quantum states ϕ and ψ as a constraint on ontological models that is worthy of investigation in its own right. They term ontological models that satisfy this assumption *maximally ψ-epistemic*. As we note in Sec. IV A (and as demonstrated in their Letter), this assumption follows from preparation noncontextuality (and hence from universal noncontextuality). However, Leifer and Maroney investigate the consequences of making the assumption of maximal ψ-epistemicity without also assuming other consequences of universal noncontextuality, in particular, without assuming other consequences of preparation noncontextuality.

They establish their no-go theorem for maximal ψ-epistemicity (and hence for universal noncontextuality) by demonstrating that maximal ψ-epistemicity implies the Kochen-Specker notion of noncontextuality (which is measurement noncontextuality together with the assumption of outcome determinism for sharp measurements), and then relying on the fact that quantum theory does not admit of a Kochen-Specker noncontextual model (the Kochen-Specker theorem).

Both our article and theirs explore senses in which a pair of quantum states may be said to be "indistinguishable," and to what extent some operational counterpart of this indistinguishability can be explained in an ontological model satisfying certain properties. But there are key differences. As we have noted, the property of ontological models that we focus on is different: we consider the assumption of universal noncontextuality rather than just maximal ψ-epistemicity [29]. The more important difference between our work and that of Leifer and Maroney, however, is in how we operationalize the notion of indistinguishability.

To explain the difference, it is useful to highlight two distinct facts about a pair of nonorthogonal pure quantum states [i.e., a pair $|\psi\rangle$ and $|\phi\rangle$ for which $|\langle\psi|\phi\rangle|^2 > 0$]:

(i) they are not perfectly *discriminable*, which is to say that there is no quantum measurement that achieves zero error in the discrimination task, formalized as $s_q > 0$, and (ii) they are *confusable*, which is to say that the ideal quantum measurement that tests for being in the state $|\phi\rangle$ has a nonzero probability of being passed by the state $|\psi\rangle$, and similarly for $|\phi\rangle$ and $|\psi\rangle$ interchanged, formalized as $c_q > 0$.

The determination of the maximum probability of discrimination for a given confusability, that is, the optimal trade-off relation that holds between s_q and c_q, is one of the central results in the field of quantum state estimation. Our work seeks to determine constraints on this trade-off relation from assumptions about the ontological model.

Leifer and Maroney, by contrast, do not consider this trade-off relation, nor the expression for the discriminability of quantum states. Rather, they address (and answer in the negative) the question of whether the degree of confusability of nonorthogonal pure quantum states can be given a particular expression in the ontological model, namely, that of Eq. (27), which asserts that the test associated to the state $|\phi\rangle$ is a test for whether the ontic state λ is inside the ontic support of the distribution representing $|\phi\rangle$ [30]. While the expression for the confusability of two quantum states is a feature of their indistinguishability, it is not one that has previously been of interest in the field of quantum state estimation.

Thus, whereas Leifer and Maroney show the impossibility of a particular ontological expression for the confusability from a known no-go result for Kochen-Specker noncontextuality (the Kochen-Specker theorem), we begin with the native phenomenology of minimum-error state discrimination (the quantum trade-off between s_q and c_q), and we derive a novel no-go result for universal noncontextuality from it.

The form of the trade-off relation between discriminability and confusability has relevance for quantum information processing tasks that make use of minimum-error state discrimination. For instance, it is used in Ref. [32] to derive the trade-off relation between concealment and bindingness in quantum bit commitment protocols [33,34], and such protocols can be used as subroutines in protocols for other tasks, such as strong coin flipping [32,35]. It has also been used in the analysis of quantum protocols for the task of oblivious transfer [36]. Our results may be useful, therefore, in determining whether or not the failure of universal noncontextuality is a resource for such tasks.

Note that because MESD for two pure quantum states is a phenomenon occurring in a two-dimensional Hilbert space (the subspace spanned by the two states) while the Kochen-Specker theorem can be proven only in Hilbert spaces of dimension three or greater, there is no possibility of leveraging facts about Kochen-Specker-uncolorable sets to infer anything about which aspects of MESD resist explanation within a universally noncontextual model [37].

A final crucial advantage of our approach over that of Ref. [25] is that it can be used to derive noncontextuality inequalities that are noise robust and hence experimentally testable, as we show in the next section. Noise robustness is critical if one hopes to leverage contextuality as a resource in real (hence, noisy) implementations of information processing protocols.

V. DEALING WITH NOISE

It is important to recognize that the inequality of Eq. (17) is not experimentally testable. To clarify this point, we first draw a distinction between noncontextuality no-go results and noncontextuality inequalities. A noncontextuality no-go result is a proof that no noncontextual model can reproduce certain predictions of quantum theory; as such, a no-go result can contain idealizations (such as perfect correlations) which are justified by quantum theory but which never hold in real experiments. In some cases (as above), a no-go result may derive an inequality on the way to deriving a logical contradiction, but such an inequality may not qualify as a proper noncontextuality inequality. In our usage, a noncontextuality inequality makes no reference to the quantum formalism and must not invoke idealized assumptions in its derivation. We give such an inequality for MESD in Sec. VI.

The distinction between no-go results and robust inequalities has historical precedent. In his 1964 paper [1], in deriving an inequality that could be shown to be violated by quantum correlations, Bell assumed an experiment wherein certain pairs of measurements had perfectly correlated outcomes. Such perfect correlations hold for ideal quantum states and measurements, but are never observed in nature. Hence, Bell's 1964 result is a no-go result, with consequences for the interpretation of quantum theory, but the inequality he derives en route to this contradiction does not provide a means of experimentally testing the principle of local causality. In 1969, Clauser, Horne, Shimony, and Holte [38] derived an inequality without assuming these idealizations. Because their inequality makes no reference to perfect correlations or to any other feature of quantum theory, its violation rules out all locally causal ontological models, independently of the validity of quantum theory. Only inequalities of this type are termed "Bell inequalities" in modern usage (so that the inequality in Bell's 1964 paper is not a "Bell inequality").

Similarly, Eq. (17) is not a proper noncontextuality inequality because it relies upon the idealization of perfect correlations between which of the states $|\phi\rangle$ or $|\bar{\phi}\rangle$ was prepared and which of the outcomes will occur in the measurement of the B_ϕ basis (and similarly for ψ and $\bar{\psi}$). To get a noncontextuality inequality, we must allow these correlations to be imperfect. Thus, in Table I, the entries that take the values 0 and 1 must instead be presumed to take the values ϵ and $1 - \epsilon$, respectively, such that ϵ becomes a parameter in our noncontextuality inequality

which quantifies the degree of imperfection of the correlations. We then show that quantum mechanics still allows higher success rates for a given confusability than any noncontextual model, even when $\epsilon \neq 0$.

Before proving this, we first rephrase the scenario as a totally operational prepare-and-measure experiment, with no reference to the quantum formalism (despite the suggestive notation below). This is a necessary first step for deriving any proper noncontextuality inequality.

A. Operationalizing MESD

We imagine an experiment involving four preparations $\{P_\phi, P_\psi, P_{\bar{\phi}}, P_{\bar{\psi}}\}$ and three binary-outcome measurements, $\{M_\phi, M_\psi, M_d\}$, with outcome sets denoted $\{\phi, \bar{\phi}\}$, $\{\psi, \bar{\psi}\}$, and $\{g_\phi, g_\psi\}$, respectively. An arbitrary data table for such an experiment would contain 12 independent parameters, specifying the probability of the first outcome of each measurement when acting on each preparation (the probability of obtaining the second outcome being fixed by normalization).

However, we wish to study the scenario in which preparations P_ϕ, P_ψ, $P_{\bar{\phi}}$, and $P_{\bar{\psi}}$ satisfy the following relation: the procedure $P_{\frac{1}{2}\phi+\frac{1}{2}\bar{\phi}}$ defined by sampling from preparations P_ϕ and $P_{\bar{\phi}}$ uniformly at random (and then forgetting which preparation occurred) is indistinguishable from the similarly defined procedure $P_{\frac{1}{2}\psi+\frac{1}{2}\bar{\psi}}$. We denote this operational equivalence by

$$P_{\frac{1}{2}\phi+\frac{1}{2}\bar{\phi}} \simeq P_{\frac{1}{2}\psi+\frac{1}{2}\bar{\psi}}. \tag{29}$$

This implies that only three of the parameters in each row are independent, so only nine independent parameters remain.

Previously the operational equivalence of Eq. (29) was guaranteed by quantum theory [Eq. (12)], but now we wish to justify it experimentally. In order to do so, one must show that the statistics for $P_{\frac{1}{2}\phi+\frac{1}{2}\bar{\phi}}$ and for $P_{\frac{1}{2}\psi+\frac{1}{2}\bar{\psi}}$ are identical for all measurements. Because the statistics of a tomographically complete set of measurements allows one to predict the statistics for *all* measurements, it suffices to verify this identity for such a tomographically complete set. Accumulating evidence that a given set of measurements is indeed tomographically complete represents the most difficult challenge for an experimental test of noncontextuality (see Refs. [21,39] for a more detailed discussion).

Note that in a realistic experiment, the four preparations that are realized, called the primary preparations, will not satisfy Eq. (29) perfectly. However, this problem can be solved by postprocessing these into "secondary preparations" which are chosen to enforce this equivalence [21,40], as we discuss in Sec. VII.

For this nine-parameter problem, the algorithm we describe in Appendix B gives the full set of necessary and sufficient noncontextuality inequalities, which we list in Appendix D. For now, however, we consider a special case with just three parameters, which captures the essence of minimum-error state discrimination. Namely, we assume symmetries that parallel those in the ideal quantum case:

$$s \equiv p(g_\phi|M_d, P_\phi) = 1 - p(g_\phi|M_d, P_\psi)$$
$$= p(g_{\bar{\phi}}|M_d, P_{\bar{\phi}}) = 1 - p(g_{\bar{\phi}}|M_d, P_{\bar{\psi}}), \tag{30}$$

$$c \equiv p(\phi|M_\phi, P_\psi) = p(\psi|M_\psi, P_\phi)$$
$$= p(\bar{\phi}|M_{\bar{\phi}}, P_{\bar{\psi}}) = p(\bar{\psi}|M_{\bar{\psi}}, P_{\bar{\phi}}), \tag{31}$$

and

$$1 - \epsilon \equiv p(\psi|M_\psi, P_\psi) = p(\phi|M_\phi, P_\phi)$$
$$= p(\bar{\psi}|M_{\bar{\psi}}, P_{\bar{\psi}}) = p(\bar{\phi}|M_{\bar{\phi}}, P_{\bar{\phi}}). \tag{32}$$

We denote the three free parameters that remain after imposing the symmetries by s, c, and $1 - \epsilon$, paralleling their ideal quantum counterparts, s_q, c_q, and 1, respectively. Just like the operational equivalence, these symmetries will never hold exactly for the primary procedures, but we can enforce them while choosing secondary procedures, as we discuss in Sec. VII.

The notation P_ϕ, P_ψ, $P_{\bar{\phi}}$, $P_{\bar{\psi}}$, M_ϕ, M_ψ, and M_d will henceforth be used to denote the secondary procedures, for which the operational equivalence and symmetries are exact.

The resulting data table, Table II, is similar to the ideal scenario of Table I, but contains the noise parameter ϵ $(1 - \epsilon)$ in place of the probability 0 (1).

Note that for each row, the average of the entries in the P_ϕ and $P_{\bar{\phi}}$ columns is $\frac{1}{2}$ (and similarly for P_ψ and $P_{\bar{\psi}}$). Here, this follows from the assumed symmetries, not from the operational equivalence (which specifies that the average of the entries for P_ϕ and $P_{\bar{\phi}}$ is the same as the average of the entries for P_ψ and $P_{\bar{\psi}}$, but not necessarily $\frac{1}{2}$); in Table I, the same averaging property is implied by the operational equivalence of each of the two mixtures to the maximally mixed quantum state in Eq. (12) (and redundantly implied by these symmetries).

Finally, we assume that the measurements and outcomes are labeled in the natural way; e.g., the outcome of M_ϕ that is more likely to occur given the preparation P_ϕ is ϕ rather than $\bar{\phi}$, and so on. Then, the data satisfy the constraint that

$$\epsilon \leq c \leq 1 - \epsilon. \tag{33}$$

TABLE II. Data table for our operational scenario.

	P_ϕ	P_ψ	$P_{\bar{\phi}}$	$P_{\bar{\psi}}$
$\phi\|M_\phi$	$1 - \epsilon$	c	ϵ	$1 - c$
$\psi\|M_\psi$	c	$1 - \epsilon$	$1 - c$	ϵ
$g_\phi\|M_d$	s	$1 - s$	$1 - s$	s

VI. NONCONTEXTUALITY INEQUALITIES FOR MESD

The operational equivalence relation of Eq. (29) together with the assumption of preparation noncontextuality implies via Eq. (7) that

$$\frac{1}{2}\mu_{P_\phi}(\lambda) + \frac{1}{2}\mu_{P_{\hat{\phi}}}(\lambda) = \frac{1}{2}\mu_{P_\psi}(\lambda) + \frac{1}{2}\mu_{P_{\hat{\psi}}}(\lambda), \quad (34)$$

where we again use the fact that convex mixtures of preparations are represented in an ontological model by the corresponding mixture of epistemic states. The fact that the ontological model must reproduce Table II implies constraints analogous to Eqs. (13)–(15) and their kin.

As we prove in Appendix B, the trade-off between s, c, and ϵ in any noncontextual model of our operational scenario must satisfy

$$s \le 1 - \frac{c - \epsilon}{2}. \quad (35)$$

In Appendix C, we show that quantum theory allows a trade-off of

$$s = \frac{1}{2}\left[1 + \sqrt{1 - \epsilon + 2\sqrt{\epsilon(1-\epsilon)c(c-1)} + c(2\epsilon - 1)}\right]. \quad (36)$$

Thus, quantum theory predicts a higher state discrimination success rate for any given c and ϵ than a noncontextual model allows. One easily verifies that Eq. (35) reduces to Eq. (17) in the limit of $\epsilon \to 0$, and that Eq. (36) reduces to Eq. (11) in the same limit. It is an open question whether Eq. (36) is the optimal trade-off that quantum theory allows. We conjecture that it is optimal for pairs of states in a two-dimensional Hilbert space.

The noncontextual and quantum trade-offs are shown in Fig. 5. The purple surface represents the triples (s, c, ϵ) saturating the inequality of Eq. (35), while the light blue surface represents the triples (s, c, ϵ) corresponding to the quantum success rate of Eq. (36).

If an experiment generates data having the form of Table II and satisfying Eq. (29), and it is found to lie above the purple shaded surface, then one has experimental evidence for the failure of noncontextuality. This evidence is independent of the validity of quantum theory, and signals a contextual advantage for state discrimination, even when one's preparations and measurements are imperfect.

A. Understanding the quantum and noncontextual bounds

For both quantum and noncontextual models, we adopt the natural labeling convention described above Eq. (33), so that all operational data necessarily satisfy
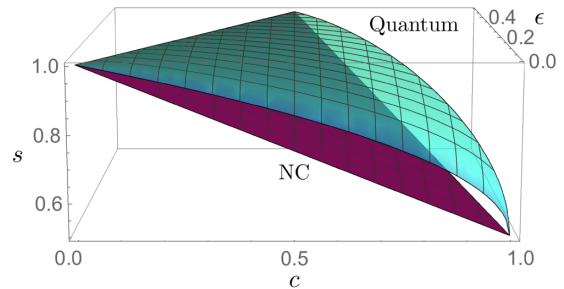

FIG. 5. Maximum success rate achievable in a noncontextual model (purple surface), and quantumly acheivable success rate (light blue surface).

$\epsilon \le c \le 1 - \epsilon$. In the $c - \epsilon$ plane of Fig. 5, these constraints describe a triangular wedge which points into the page.

In the plane with $\epsilon = 0$, Sec. IV A provides an intuitive explanation for the trade-off relation.

In the plane with $\epsilon = c$, we can see that for both quantum and noncontextual models, the preparations can be perfectly distinguishable, $s = 1$. This follows from the fact that the value of ϵ quantifies the noise in M_ϕ and M_ψ, and when c is no larger than ϵ, we can attribute *all* of the confusability to this noise. Explicitly, one can construct a quantum model where preparation P_ϕ is represented by $|0\rangle\langle 0|$ and P_ψ is represented by $|1\rangle\langle 1|$ and where effect $E_{\phi|M_\phi}$ is represented by $(1-\epsilon)|0\rangle\langle 0| + \epsilon|1\rangle\langle 1|$ and $E_{\psi|M_\psi}$ is represented by $\epsilon|0\rangle\langle 0| + (1-\epsilon)|1\rangle\langle 1|$, which implies that $c = \epsilon$, while $s = 1$ for the Helstrom measurement $\{|0\rangle\langle 0|, |1\rangle\langle 1|\}$. Furthermore, since these states and effects are all diagonal in the same basis, we can take the eigenvalues of these to define the conditional probabilities of a noncontextual model which achieves $c = \epsilon$ and $s = 1$.

Whenever $c > \epsilon$, however, the noise in M_ϕ and M_ψ cannot explain all of the confusability, and therefore some of this confusability must be explained by the lack of perfect distinguishability of the preparations; that is, in a quantum model, the preparations must be represented by nonorthogonal states, while in a noncontextual model, they must be represented by overlapping probability distributions. Thus, the maximum value of s falls away from 1 as we move away from the $\epsilon = c$ plane. In a noncontextual model, it falls off linearly, interpolating between its value for $\epsilon = c$ and its value for $\epsilon = 0$. The quantum bound falls off more slowly.

B. Robustness to depolarizing noise

We can get a sense for the robustness of our noncontextuality inequalities by considering a specific noise model in quantum theory. Imagine that one's attempts to implement the ideal quantum preparations and measurements are thwarted by a depolarizing channel which has the same noise parameter v for all states and effects:
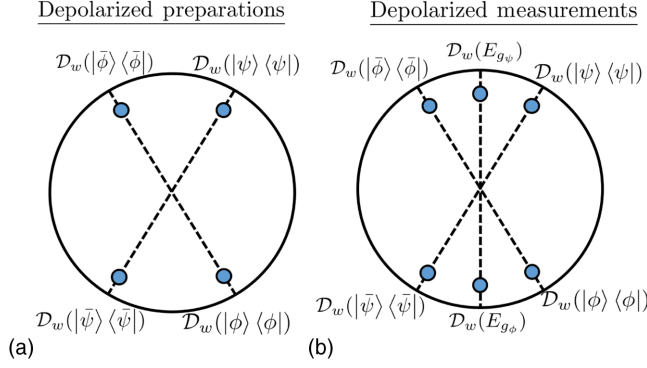
FIG. 6. The images of the ideal quantum states and effects under a depolarization map for some fixed value of v.

$$\mathcal{D}_v(\rho) = (1 - v)\rho + v\frac{\mathbb{1}}{2}, \tag{37}$$

$$\mathcal{D}_v(E_k) = (1 - v)E_k + v\frac{\mathbb{1}}{2}. \tag{38}$$

The resulting states and effects are shown in Fig. 6 for some fixed v. One can graphically see that this uniform depolarization map generates a new set of states and measurements which satisfy the symmetries and operational equivalence we require. However, if the noise is too large, our noncontextuality inequality will not be violated, as we now show.

This noisy model generates a data table of the form of Table II with

$$s = \frac{1}{2} + (1 - v)^2 \left\{ \left[\frac{1}{2}\left(1 + \sqrt{1 - c_q}\right) \right] - \frac{1}{2} \right\}, \tag{39}$$

$$c = \frac{1}{2} + (1 - v)^2 \left(c_q - \frac{1}{2}\right), \tag{40}$$

$$\epsilon = \frac{1}{2}[1 - (1 - v)^2]. \tag{41}$$

As always, $c_q = |\langle\phi|\psi\rangle|^2$.

The maximum level of noise v that still violates our noncontextuality inequality, Eq. (35), is easily calculated as a function of the Bloch sphere angle θ between the two states [defined by $\cos^2(\theta/2) = |\langle\phi|\psi\rangle|^2$], by substituting Eqs. (39)–(41) into Eq. (35):

$$v = 1 - \frac{1}{c_q + \sqrt{1 - c_q}} = 1 - \frac{1}{\cos^2\left(\frac{\theta}{2}\right) + \sin\left(\frac{\theta}{2}\right)}. \tag{42}$$

Equation (42) is plotted in Fig. 7. For $\theta = 0$ or $\theta = \pi$, the noncontextual bound equals the ideal quantum bound, and hence no experiment can violate our noncontextuality inequality at these extremal angles. For all other θ, an experiment with depolarizing noise such that $v \leq 1 - \{1/[\cos^2(\theta/2) + \sin(\theta/2)]\}$ can violate the inequality.
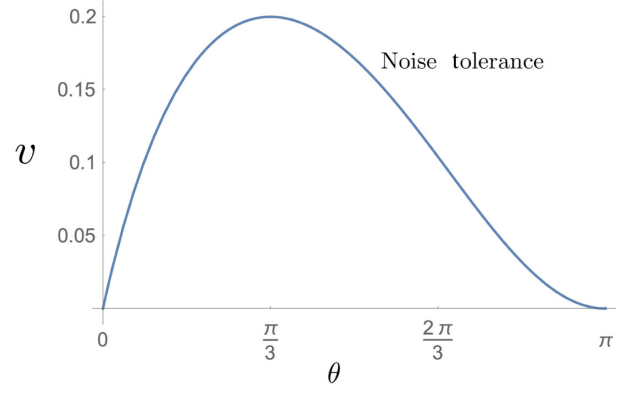

FIG. 7. The maximum value of the parameter v for the depolarizing noise model that allows a violation of our noncontextuality inequality, as a function of the Bloch sphere angle θ between the two states.

The maximum tolerance to noise ($v = 0.2$) occurs when $\theta = (\pi/3)$.

VII. ENFORCING SYMMETRIES AND OPERATIONAL EQUIVALENCES

In Sec. V A, we predicate our noncontextuality inequalities on the exact operational equivalence of Eq. (29) and exact operational symmetries of Eqs. (30)–(32), yet we claim that these idealizations *can* in fact be realized in realistic, noisy experiments. Of course, no experimental data will *directly* satisfy either of these requirements; rather, one performs a postprocessing of the data, as originally outlined in Ref. [21].

For pedagogical clarity, we discuss this data processing under the assumption that the operational theory is quantum theory. Note, however, that our comments can easily be generalized to the framework of generalized probabilistic theories (defined in Refs. [41,42]), as demonstrated in Refs. [21,39]. Indeed, the analysis *must* be performed in this framework if one hopes to directly test the hypothesis of noncontextuality against one's experimental data (i.e., without assuming the validity of quantum theory).

For any set $\mathfrak{P}$ of noisy preparations that has been performed experimentally, one can simulate perfectly the statistics of all other preparations in the convex hull of $\mathfrak{P}$, viewed as points in the quantum state space (here, simply a plane of the Bloch sphere). Similarly, for any set $\mathfrak{E}$ of noisy measurement effects, one can perfectly simulate the statistics of all other effects in the convex hull of $\mathfrak{E}$, viewed as points in the space of valid quantum effects. In Ref. [21], this fact was leveraged to simulate exact operational equivalences for a set of secondary preparations from data on a set of primary preparations that failed to satisfy the operational equivalences exactly. Here, we leverage this trick to simulate preparations and measurements which simultaneously satisfy our operational equivalence *as well as* the symmetries. We now argue that this can always be
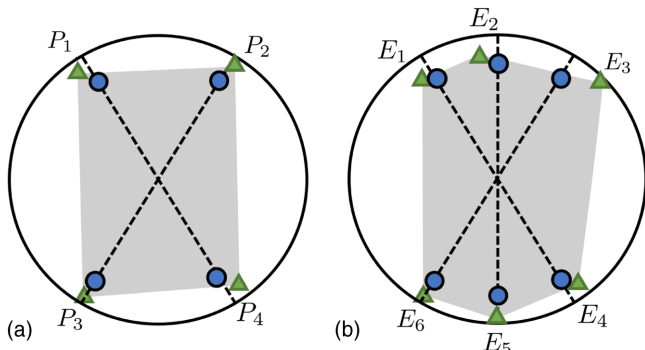
FIG. 8. (a) If one can perform the four primary preparations P_1–P_4 (shown as green triangles), then one can simulate any preparation in their convex hull (shown as a light gray shaded region). In particular, one can simulate secondary procedures that are depolarized versions of the ideal preparations (shown as blue circles like those in Fig. 6). (b) Similarly for the measurements.

done, although if the primary preparations or measurements are too noisy, the resulting simulated data will not violate our inequalities.

As we show explicitly in Sec. VI B, even a partially depolarized set of states and effects can violate our inequality. Hence, one need only realize experimental sets $\mathfrak{P}$ and $\mathfrak{E}$ which contain in their convex hull the images of our ideal states and effects under the depolarization map $\mathcal{D}_v$ with $v \leq 1 - \{1/[\cos^2(\theta/2) + \sin(\theta/2)]\}$. Then, one can postprocess the data obtained from $\mathfrak{P}$ and $\mathfrak{E}$ to obtain a physically meaningful set of data which satisfies the operational equivalence and symmetries that we assume in the main text, and our inequality will still be violated. Geometrically, this simply means that the primary preparations must have a convex hull which contains the image of the ideal states under a depolarizing map with $v \leq 1 - \{1/\cos^2(\theta/2) + \sin(\theta/2)\}$, as pictured in Fig. 8 (and similarly for the measurements, also pictured).

In fact, there are other noisy sets of preparations and measurements besides the depolarized versions of the corresponding ideals which satisfy the operational equivalence and symmetries needed for the noncontextuality inequality to apply. A simple example is states and measurements that are depolarized versions of the ideals that are also rotated in the plane by the same angle. By doing such a rotation, one may be able to simulate a set of states and effects with less depolarization, which then leads to larger violations. In general, there are many sets of states and effects that satisfy our operational equivalence and symmetries. Given a set of primary procedures that one has performed and characterized, finding the states and measurements satisfying our constraints which maximize the violation of our inequality is a straightforward linear program [21].

Leveraging the convex structure of operational theories in order to define secondary laboratory procedures which respect certain theoretical idealizations is a powerful tool

which we expect to have broad applicability. To date, this method has been proposed to identify operational procedures which respect exact operational equivalences. What we have just shown is that the method also allows one to enforce natural symmetries which greatly simplify the problem at hand [as evidenced by comparing Eq. (35) to the set of inequalities in Appendix D]. Of course, this tool does not allow one to define laboratory procedures which satisfy *any* desired idealizations; for example, one could never generate a pure state or a sharp measurement effect by convexly mixing the noisy procedures actually performed in the lab. We expect future work to continue expanding the range of practical applicability of the technique of secondary procedures.

VIII. ISOMORPHISM BETWEEN MESD AND A BELL SCENARIO

Any noncontextuality scenario that makes no assumptions of measurement noncontextuality, and for which there is a *single* mixed preparation whose various ensemble decompositions generate *all* of the operational equivalences of interest, is isomorphic to a related Bell scenario [43]. Both of these conditions hold for our MESD scenario, since we do not consider any operational equivalences among the measurements, and the operational equivalences among the preparations are generated by decompositions of a single mixed preparation (e.g., the maximally mixed state in the ideal case). The operational Bell scenario which relates to our MESD scenario is one with two parties, whom we denote by S and M (for reasons that will become apparent), where S has two binary measurements, denoted S_1 and S_2, and M has three binary measurements, denoted M_1, M_2, and M_3. The outcomes (which we denote s_i for S_i and m_j for M_j) take values in the set $\{-1, +1\}$.

For such a scenario, the set of constraints which define the local set of correlations is given by positivity inequalities, $p(s_i m_j | S_i M_j) \geq 0$, the normalization condition $\sum_{s_i m_j} p(s_i m_j | S_i M_j) \geq 0$, and the CHSH inequalities [38] (applied to any of the three possible pairings of two measurement settings on S with two measurement settings on M) [44]. As we show, the bound on our MESD success rate follows under our assumed symmetries from the CHSH inequality

$$\langle s_1 m_1 \rangle + \langle s_1 m_3 \rangle + \langle s_2 m_1 \rangle - \langle s_2 m_3 \rangle \leq 2, \quad (43)$$

where

$$\langle s_i m_j \rangle = \sum_{s_i m_j} s_i m_j p(s_i m_j | S_i M_j)$$

$$= 2p(s_i = m_j | S_i M_j) - 1. \quad (44)$$

The connection between this Bell scenario and our MESD scenario is most easily seen in the ideal quantum

realization. Imagine that the two parties share a maximally entangled state $|\Phi^+\rangle_{SM} = (1/\sqrt{2})(|00\rangle_{SM} + |11\rangle_{SM})$ (with $|0\rangle$ and $|1\rangle$ defined so that $|\phi\rangle$ and $|\psi\rangle$ have real coefficients when written in this basis), and imagine that their measurements correspond to the quantum measurements from the main text, as follows:

$$
\begin{aligned}
S_1 &= \{|\phi\rangle_S\langle\phi|, |\bar{\phi}\rangle_S\langle\bar{\phi}|\}, \\
S_2 &= \{|\psi\rangle_S\langle\psi|, |\bar{\psi}\rangle_S\langle\bar{\psi}|\}, \\
M_1 &= \{|\phi\rangle_M\langle\phi|, |\bar{\phi}\rangle_M\langle\bar{\phi}|\}, \\
M_2 &= \{|\psi\rangle_M\langle\psi|, |\bar{\psi}\rangle_M\langle\bar{\psi}|\}, \\
M_3 &= \{E_M^{g_\phi}, E_M^{g_\psi}\}.
\end{aligned}
\tag{45}
$$

We take the $+1$ outcome for each measurement to correspond to the first quantum effect for that measurement. This ideal quantum realization of this Bell scenario is conceptually transformed into our ideal quantum realization of the MESD scenario by viewing a measurement by party S to be a remote preparation (via quantum steering) for party M. For example, outcome $+1$ for S_1 remotely prepares the state $|\phi\rangle_M$ (which is why we have chosen the notation S, for "source"). Similarly, outcome -1 for measurement S_2 prepares the state $|\bar{\psi}\rangle_M$, and so on.

Thus, one can verify that in the ideal quantum realization, s_q and c_q become [in our new notation, and assuming the symmetries in Eqs. (30)–(32)]

$$
\begin{aligned}
s_q &= p(s_1 = m_3|S_1M_3) = 1 - p(s_2 = m_3|S_2M_3), \\
c_q &= p(s_1 = m_2|S_1M_2) = p(s_2 = m_1|S_2M_1),
\end{aligned}
\tag{46}
$$

while the fact that paired preparations and measurements are perfectly correlated in the ideal quantum realization corresponds to

$$
0 = 1 - p(s_1 = m_1|S_1M_1) = 1 - p(s_2 = m_2|S_2M_2).
\tag{47}
$$

Furthermore, the no-signaling condition in the Bell scenario implies the operational equivalence of our MESD scenario. If party S performs measurement S_1, the updated state on M will be either $|\phi\rangle$ or $|\bar{\phi}\rangle$ with equal likelihood, and if party S performs measurement S_2, the updated state on M will be either $|\psi\rangle$ or $|\bar{\psi}\rangle$ with equal likelihood. In quantum theory, the no-signaling condition implies that the average density operator prepared on M is the same for either choice of measurement by S, which is precisely the operational equivalence of Eq. (12).

Using Eq. (44), we can write Eqs. (46) and (47) in terms of expectation values:

$$
\begin{aligned}
s_q &= \frac{1}{2}(1 + \langle s_1 m_3\rangle) = \frac{1}{2}(1 - \langle s_2 m_3\rangle), \\
c_q &= \frac{1}{2}(1 + \langle s_1 m_2\rangle) = \frac{1}{2}(1 + \langle s_2 m_1\rangle), \\
0 &= \frac{1}{2}(1 - \langle s_1 m_1\rangle) = \frac{1}{2}(1 - \langle s_2 m_2\rangle).
\end{aligned}
\tag{48}
$$

Rewriting Eq. (43) in terms of s_q and c_q instead of expectation values, one obtains

$$
s_q \le 1 - \frac{c_q}{2},
\tag{49}
$$

recovering Eq. (17), our bound for the success rate in state discrimination.

Because both the Bell scenario and our MESD scenario are operationally defined, one can also make the translation without assuming the ideal quantum realizations. In a realistic operational scenario, ϵ will be nonzero, and one obtains

$$
\begin{aligned}
s &= p(s_1 = m_3|S_1M_3) = 1 - p(s_2 = m_3|S_2M_3), \\
c &= p(s_1 = m_2|S_1M_2) = p(s_2 = m_1|S_2M_1), \\
\epsilon &= 1 - p(s_1 = m_1|S_1M_1) = 1 - p(s_2 = m_2|S_2M_2).
\end{aligned}
\tag{50}
$$

Rewriting Eq. (43) in terms of s, c, and ϵ instead of expectation values, one obtains

$$
s \le 1 - \frac{c - \epsilon}{2},
\tag{51}
$$

recovering Eq. (35), our bound for the success rate in state discrimination.

Because of the redundancies induced by our assumed symmetries, Eq. (35) follows also from the CHSH inequality,

$$
\langle s_1 m_2\rangle + \langle s_1 m_3\rangle + \langle s_2 m_2\rangle - \langle s_2 m_3\rangle \le 2,
\tag{52}
$$

by the same logic. More generally, if we do not assume any symmetries, then there are no redundant inequalities. If we furthermore do not assume the natural labeling constraint [Eq. (33)], then the full polytope of local correlations for this Bell scenario [44] [and described just above Eq. (43)] is isomorphic to the full polytope of noncontextual correlations for our MESD scenario.

IX. FUTURE DIRECTIONS

We have derived noncontextuality inequalities that minimum-error state discrimination in quantum theory that fails to admit of a noncontextual model. We derive noncontextuality inequalities that delimit the trade-off between

success rate, error rate, and confusability in state discrimination, independently of the validity of quantum theory.

Our results show that contextuality is a resource for state discrimination, even in realistic, noisy experiments. This suggests many directions for future research. One important question is how our results translate into advantages for quantum information processing tasks which have MESD as a subroutine. Because many such tasks (e.g., key distribution) consider consecutive measurements on the system, this research program would require further analysis regarding the consequences of noncontextuality for experiments involving sequential measurements [6,45,46].

It would also be interesting to generalize these results to other types of state discrimination, such as unambiguous state discrimination. Indeed, one can easily derive a relevant no-go theorem. The challenge is to define an operational notion of "unambiguous" given that no measurement yields truly unambiguous knowledge in the presence of noise. Once this challenge is met, it should be straightforward to apply the general algorithm that we have introduced in this article in order to derive the noncontextuality inequalities for this scenario. Understanding the relation between noncontextuality and other kinds of state discrimination should translate into new kinds of quantum advantages for information processing tasks.

APPENDIX A: PROOF OF NONCONTEXTUALITY NO-GO THEOREM FOR MESD

Herein we provide an alternative proof of our no-go theorem, Eq. (17), that is, of the fact that the inequality

$$s_q \leq 1 - \frac{c_q}{2} \tag{A1}$$

must be satisfied for any s_q and c_q arising in a noncontextual model that reproduces the data in Table I and respects the operational equivalence of Eq. (12). While the

proof provided in the main text uses an intuitive argument that is native to the task of state discrimination, the argument in this section abstracts away from the specific problem at hand, and as such extends naturally to the more general method required for proving Eq. (35) (as we discuss in Appendix B).

First, we allow the ontological model to have an ontic state space of arbitrary form, and we allow the response functions to be outcome-indeterministic. Second, we show that for any such model, there exists a simpler ontological model which is equally general, but which has only eight ontic states and has response functions that are purely outcome-deterministic. Third, we show that two of these ontic states are superfluous if B_d is optimal for state discrimination. Fourth, we show that the forms of the epistemic states are greatly constrained by their perfectly predictable responses on the corresponding measurements. Fifth, we parametrize the set of possible epistemic states as probability distributions over the remaining six ontic states in accordance with these constraints. Sixth, we calculate the values of s_q and c_q in terms of these response functions and epistemic states. Finally, we impose preparation noncontextuality and eliminate the unobserved variables to obtain the optimal trade-off between s_q and c_q.

As one ranges over the ontic states in our ontological model, the vector $\left(\xi_{\phi|B_\phi}(\lambda), \xi_{\psi|B_\psi}(\lambda), \xi_{g_\phi|B_d}(\lambda)\right)$ of valid probability assignments to our three binary basis measurements defines a unit cube. The most obvious ontological model would have one λ for each possible probability assignment (including the indeterministic ones), defining an ontic state space isomorphic to the unit cube. The epistemic states in such an outcome-indeterministic model would be arbitrary normalized probability densities over this set of ontic states (that is, all the interior points of the cube).

However, we can always simplify matters without loss of generality by decomposing each nonextremal probability value assignment into extremal assignments. (These extremal points are outcome-deterministic if and only if there are no nontrivial constraints from measurement noncontextuality, but this is indeed the case here.)

Let us define a variable κ which runs over the eight extremal points in the cube of ontic states. Then, there exists a $p(\kappa|\lambda)$ such that $\xi_{k|M}(\lambda) = \sum_\kappa \xi_{k|M}(\kappa)p(\kappa|\lambda)$. We can thus write any observable probability $p(k|M, P)$ as

$$p(k|M, P) = \int_\Lambda \xi_{k|M}(\lambda)\mu_P(\lambda)d\lambda$$
$$= \sum_\kappa \xi_{k|M}(\kappa)\mu_P(\kappa), \tag{A2}$$

where $\mu_P(\kappa) \equiv \int_\Lambda d\lambda p(\kappa|\lambda)\mu_P(\lambda)$. This construction lets us write observed probabilities in terms of extremal value

assignments by effectively moving uncertainty into the new state distributions $\mu_P(\kappa)$.

We sometimes simplify the notation by letting the distributions and response functions be vectors of probabilities indexed by the ontic states κ; e.g.,

$$p(k|M, P) = \sum_\kappa \xi_{k|M}(\kappa)\mu_P(\kappa) = \vec{\xi}_{k|M} \cdot \vec{\mu}_P. \quad (A3)$$

We thus convert an outcome-indeterministic model over a continuum of ontic states (the unit cube) to an outcome-deterministic model over just eight ontic states (its vertices), without any loss of generality. The vertices κ_1–κ_8 correspond to the deterministic triples,

$$(\xi_{\phi|B_\phi}(\kappa), \xi_{\psi|B_\psi}(\kappa), \xi_{g_\phi|B_d}(\kappa))$$
$$\in \{(0,0,0), (0,0,1), (0,1,0), \dots, (1,1,1)\}, \quad (A4)$$

so the three response functions are

$$\vec{\xi}_{\phi|B_\phi} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}, \qquad \vec{\xi}_{\psi|B_\psi} = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 1 \\ 1 \end{bmatrix}, \qquad \vec{\xi}_{g_\phi|B_d} = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 1 \\ 0 \\ 1 \\ 0 \\ 1 \end{bmatrix}. \quad (A5)$$

In fact, if we assume B_d is optimal, the fourth and fifth of these value assignments will never occur. Consider, for example, the triple $(1, 0, 0)$ (which occurs for κ_5). Since $\xi_{\phi|B_\phi}(\kappa_5) = 1$, the state cannot have been $\bar{\phi}$. Since $\xi_{\psi|B_\psi}(\kappa_5) = 0$, the state cannot have been ψ. Thus, we know the state must have been ϕ or $\bar{\psi}$; in either case, the winning strategy is for B_d to return the outcome g_ϕ. Therefore, the winning strategy has $\xi_{g_\phi|B_d}(\kappa_5) = 1$, and thus the triple $(1, 0, 0)$ never occurs in the winning strategy. Similar logic applies to the triple $(0, 1, 1)$, and hence we need not consider these two assignments [47]. The remaining value assignments are

$$(\xi_{\phi|B_\phi}(\kappa), \xi_{\psi|B_\psi}(\kappa), \xi_{g_\phi|B_d}(\kappa))$$
$$\in \{(0,0,0), (0,0,1), (0,1,0), (1,0,1),$$
$$(1,1,0), (1,1,1)\}. \quad (A6)$$

Thus, six ontic states are sufficient for describing our experiment: one for each remaining deterministic assignment. It follows that the vectors representing each of the three response functions are

$$\vec{\xi}_{\phi|B_\phi} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 1 \end{bmatrix}, \qquad \vec{\xi}_{\psi|B_\psi} = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 1 \\ 1 \end{bmatrix}, \qquad \vec{\xi}_{g_\phi|B_d} = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 1 \\ 0 \\ 1 \end{bmatrix}. \quad (A7)$$

We can constrain the most general form of the epistemic states using the perfect predictability of measurements B_ϕ and B_ψ on their corresponding states. Namely, recalling Eq. (25) and the form of the response functions, $\xi_{\phi|B_\phi}(\kappa)$, $\xi_{\psi|B_\psi}(\kappa)$, $\xi_{\bar{\psi}|B_\psi}(\kappa) \equiv 1 - \xi_{\psi|B_\psi}(\kappa)$, and $\xi_{\bar{\phi}|B_\phi}(\kappa) \equiv 1 - \xi_{\phi|B_\phi}(\kappa)$, we can see that our epistemic states must have the form

$$\vec{\mu}_\phi = \begin{bmatrix} 0 \\ 0 \\ 0 \\ a_4 \\ a_5 \\ a_6 \end{bmatrix}, \quad \vec{\mu}_{\bar{\phi}} = \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad \vec{\mu}_\psi = \begin{bmatrix} 0 \\ 0 \\ b_3 \\ 0 \\ b_5 \\ b_6 \end{bmatrix}, \quad \vec{\mu}_{\bar{\psi}} = \begin{bmatrix} b_1 \\ b_2 \\ 0 \\ b_4 \\ 0 \\ 0 \end{bmatrix}, \quad (A8)$$

where normalization requires that $a_4 + a_5 + a_6 = 1$, and so on.

The definitions of c_q and s_q in Eqs. (8) and (10) translated into our ontological model become

$$c_q = \vec{\mu}_\psi \cdot \vec{\xi}_{\phi|B_\phi} = \vec{\mu}_\phi \cdot \vec{\xi}_{\psi|B_\psi} = 1 - \vec{\mu}_{\bar{\psi}} \cdot \vec{\xi}_{\phi|B_\phi}$$
$$= 1 - \vec{\mu}_{\bar{\phi}} \cdot \vec{\xi}_{\psi|B_\psi}, \quad (A9)$$

$$s_q = \vec{\mu}_\phi \cdot \vec{\xi}_{g_\phi|B_d} = 1 - \vec{\mu}_\psi \cdot \vec{\xi}_{g_\phi|B_d} = 1 - \vec{\mu}_{\bar{\phi}} \cdot \vec{\xi}_{g_\phi|B_d}$$
$$= \vec{\mu}_{\bar{\psi}} \cdot \vec{\xi}_{g_\phi|B_d}. \quad (A10)$$

Taking these dot products using the vectors in Eqs. (A7) and (A8) gives

$$c_q = b_5 + b_6 = a_5 + a_6 = 1 - b_4 = 1 - a_3 \quad (A11)$$

and

$$s_q = a_4 + a_6 = 1 - b_6 = 1 - a_2 = b_2 + b_4. \quad (A12)$$

Because the epistemic states must be normalized, it follows that $b_5 + b_6 = 1 - b_3$, $a_5 + a_6 = 1 - a_4$, $a_4 + a_6 = 1 - a_5$, and $b_2 + b_4 = 1 - b_1$. Substituting these four expressions, we obtain

$$c_q = 1 - b_3 = 1 - a_4 = 1 - b_4 = 1 - a_3 \qquad \text{(A13)}$$

and

$$s_q = 1 - a_5 = 1 - b_6 = 1 - a_2 = 1 - b_1, \qquad \text{(A14)}$$

and hence $b_3 = a_4 = b_4 = a_3$ and $a_5 = b_6 = a_2 = b_1$.

Let us take $s_q = 1 - a_2$ and $c_q = 1 - a_3$. If there were no more constraints, then a_2 and a_3 could range from 0 to 1 independently, and s_q and c_q could take any values from 0 to 1. By imposing preparation noncontextuality, however, we have

$$\vec{\mu}_{\frac{1}{2}} = \frac{1}{2} \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \\ a_6 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} b_1 \\ b_2 \\ b_3 \\ b_4 \\ b_5 \\ b_6 \end{bmatrix}. \qquad \text{(A15)}$$

This implies $b_i = a_i$ for all i. Since $a_1 + a_2 + a_3 = 1$ from normalization, $a_1 = b_1$ from preparation noncontextuality, and $b_1 = a_2$ as derived above, we also have $2a_2 + a_3 = 1$ and hence $c_q = 2a_2$. Finally, writing s_q in terms of c_q yields

$$s_q = 1 - \frac{c_q}{2}. \qquad \text{(A16)}$$

APPENDIX B: PROOF OF NONCONTEXTUALITY INEQUALITY FOR MESD

Herein we prove our noncontextuality inequality, Eq. (35); that is, we prove that

$$s \leq 1 - \frac{c - \epsilon}{2} \qquad \text{(B1)}$$

must be satisfied for any s, c, and ϵ arising in a non-contextual model that reproduces data in Table II and respects Eq. (29).

First, we use the arguments of Appendix A to write down an ontological model with eight ontic states and purely outcome-deterministic response functions. Second, we parametrize the set of possible epistemic states for this second model in accordance with preparation noncontextuality. Third, we calculate expressions for s, c, and ϵ in terms of these response functions and epistemic states. These manipulations reduce the problem to a small set of linear equalities and inequalities over unobserved and observed variables. Finally, we eliminate the unobserved variables to obtain inequalities concerning only the observed variables s, c, and ϵ.

Exactly as before, we can convert a general, outcome-indeterministic model over a continuum of ontic states (the unit cube) to an outcome-deterministic model over just eight ontic states (its vertices), without any loss of generality. (As before, this is simply a mathematical construction, and in no way commits us to a fundamental principle of outcome determinism.) The vertices of the unit cube, κ_1–κ_8, again correspond to the deterministic triples

$$(\xi_{\phi|M_\phi}(\kappa), \xi_{\psi|M_\psi}(\kappa), \xi_{g_\phi|M_d}(\kappa))$$
$$\in \{(0,0,0), (0,0,1), (0,1,0), \ldots, (1,1,1)\}, \qquad \text{(B2)}$$

and the three response functions are again

$$\vec{\xi}_{\phi|M_\phi} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}, \quad \vec{\xi}_{\psi|M_\psi} = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 1 \\ 1 \end{bmatrix}, \quad \vec{\xi}_{g_\phi|M_d} = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 1 \\ 0 \\ 1 \\ 0 \\ 1 \end{bmatrix}. \qquad \text{(B3)}$$

In a more general situation in which measurement non-contextuality is also leveraged, there will be linear constraints on this set of response functions, and the extremal response functions will no longer all be outcome-deterministic. In this case, one can still explicitly enumerate the finite set of extremal response functions by taking the intersection of the linear constraints with the above cube of value assignments. These extremal points modify the specific form of Eq. (B3), and our methods would proceed largely unchanged.

Each preparation generates a probability distribution over κ, so we can write the epistemic states as

$$\vec{\mu}_{P_\phi} = \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \\ a_6 \\ a_7 \\ a_8 \end{bmatrix}, \vec{\mu}_{P_\psi} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \\ b_4 \\ b_5 \\ b_6 \\ b_7 \\ b_8 \end{bmatrix}, \vec{\mu}_{P_{\bar{\phi}}} = \begin{bmatrix} c_1 \\ c_2 \\ c_3 \\ c_4 \\ c_5 \\ c_6 \\ c_7 \\ c_8 \end{bmatrix}, \vec{\mu}_{P_{\bar{\psi}}} = \begin{bmatrix} d_1 \\ d_2 \\ d_3 \\ d_4 \\ d_5 \\ d_6 \\ d_7 \\ d_8 \end{bmatrix}, \qquad \text{(B4)}$$

where the parameters in each vector are positive and sum to 1.

Dot products between a vector in Eq. (B3) and a vector in Eq. (B4) can produce any set of observable statistics, and thus constitute a general ontological model for our measurements and preparations. The values of (s, c, ϵ) that we

can observe in a noncontextual model with our assumed symmetries, however, are restricted by the above constraints, all of which we repeat here for convenience.

Equations (2) and (3) imply that for all four preparations,

$$\forall \kappa: \ 0 \leq [\vec{\mu}_P]_\kappa \leq 1 \qquad (B5)$$

and

$$\sum_k [\vec{\mu}_P]_k = 1. \qquad (B6)$$

Equation (29) gives

$$\vec{\mu}_{P_\phi} + \vec{\mu}_{P_{\bar{\phi}}} = \vec{\mu}_{P_\psi} + \vec{\mu}_{P_{\bar{\psi}}}. \qquad (B7)$$

Equations (30)–(32) are, respectively,

$$\vec{\mu}_{P_\phi} \cdot \vec{\xi}_{g_\phi | M_d} = 1 - \vec{\mu}_{P_\psi} \cdot \vec{\xi}_{g_\phi | M_d} = s. \qquad (B8)$$

$$\vec{\mu}_{P_\psi} \cdot \vec{\xi}_{\phi | M_\phi} = \vec{\mu}_{P_\phi} \cdot \vec{\xi}_{\psi | M_\psi} = c, \qquad (B9)$$

$$\vec{\mu}_{P_\psi} \cdot \vec{\xi}_{\psi | M_\psi} = \vec{\mu}_{P_\phi} \cdot \vec{\xi}_{\phi | M_\phi} = 1 - \epsilon. \qquad (B10)$$

Equation (33) gives

$$\epsilon \leq c \leq 1 - \epsilon. \qquad (B11)$$

Equations (B5)–(B11) define a set of constraints over the variables s, c, ϵ, a_i, b_i, c_i, and d_i (where $i \in \{1, 2, \ldots, 8\}$). Although the parameters a_i, b_i, c_i, d_i in our epistemic states are not observable, constraints upon them [Eqs. (B5) and (B6)] have consequences for the set of possible triples (s, c, ϵ). Finding the set of inequalities over only (s, c, ϵ) that is implied by the full set of linear equalities and inequalities of Eqs. (B5)–(B11) is algebraically tedious by hand, but straightforward using the well-known Fourier-Motzkin elimination algorithm, which returns our result

$$s \leq 1 - \frac{c - \epsilon}{2}. \qquad (B12)$$

It is worth noting that the technique for deriving noncontextuality inequalities we introduce here, insofar as it reduces to a convex hull problem, is an instance of the problem of quantifier elimination. Recent work in quantum foundations has seen increasing use of quantifier elimination algorithms, in noncontextuality [48,49] as well as other scenarios. Fourier-Motzkin elimination, which is appropriate for problems wherein the dependence on the variables to be eliminated is linear, has been used to derive Bell inequalities [50], and also recently, to derive Bell-like inequalities for novel causal scenarios [51–53]. In Ref. [53], where the

problem is reduced to what is known as the classical marginals problem—that of determining whether a given set of distributions on various subsets of a set of variables can arise as the marginals of a single joint distribution over all of the variables—this problem can be solved by performing quantifier elimination on the probabilities in the joint distributions using convex hull algorithms. Nonlinear quantifier elimination using cylindrical algebraic decomposition has also found application in deriving Bell-like inequalities in simple scenarios [52,54]. We anticipate that these more general techniques for quantifier elimination will ultimately also find applications to the derivation of noncontextuality inequalities.

APPENDIX C: NOISY QUANTUM REALIZATION WHICH VIOLATES OUR NONCONTEXTUALITY INEQUALITY

We now sketch a quantum realization of the MESD scenario for any given values of c and ϵ satisfying the assumed symmetries and operational equivalences and violating our noncontextuality inequality for all values of c and ϵ. (The ideal quantum realization of the MESD scenario, given earlier, was defined only for $\epsilon = 0$.)

There is no general technique for finding the set of all data tables achievable in quantum theory for some prepare-and-measure scenario. For some cases (e.g., Bell tests), this set can be approximated efficiently via the Navascues-Pironio-Acin hierarchy [55]. For situations with multiple preparations or additional constraints, no such method exists yet.

However, we can apply our understanding from Sec. VI A to construct a quantum model which recovers Eq. (36), which we conjecture is optimal for qubits. Namely, because we want to find the maximum value of s consistent with a given c and ϵ, we should attribute as much of the confusability as possible to noise in the M_ϕ and M_ψ measurements, and only attribute the remainder of the confusability to nonorthogonality of the states. As such, in this section we allow the effects E_ϕ, E_ψ, $E_{\bar{\phi}}$, and $E_{\bar{\psi}}$ to be noisy. That is, the measurements are allowed to be any positive-operator valued measures, unlike in Appendix A (where E_ϕ denotes a projector onto $|\phi\rangle$, and so on).

Imagine P_ϕ prepares state $|0\rangle$ on the Bloch sphere and P_ψ prepares a pure state $|\theta\rangle$ rotated by $\theta \in [0, \pi]$ with respect to $|0\rangle$ in the $X - Z$ plane. We specify the value of θ later. Within this plane, the effect E_ϕ must lie on the green line shown in Fig. 9, since only these effects imply $\langle 0 | E_\phi | 0 \rangle = 1 - \epsilon$.

The choice of E_ϕ that yields the maximum confusability is the one on the green line, closest to $|\theta\rangle$ (but not closer to $|\theta\rangle$ than to $|0\rangle$, since that would imply that $c \geq 1 - \epsilon$). The remainder of the confusability must then be attributed to the nonzero inner product between the two pure states, so θ is fixed by $\langle \theta | E_\phi | \theta \rangle = c$. Now that the two states
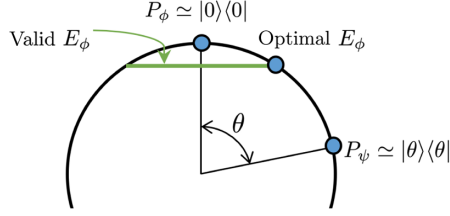
FIG. 9. Sketch of the quantum model which yields Eq. (36).

are specified, calculating the optimal (Helstrom) probability is a simple quantum calculation whose result gives Eq. (36); that is,

$$s = \frac{1}{2}\left[1 + \sqrt{1 - \epsilon + 2\sqrt{\epsilon(1-\epsilon)c(c-1)} + c(2\epsilon - 1)}\right]. \tag{C1}$$

The remaining states and effects are completely fixed by the assumed symmetries and operational equivalence. For a general pair of c and ϵ, this quantum model outperforms the optimal noncontextual model, as seen in Fig. 5.

APPENDIX D: FULL SET OF NONCONTEXTUALITY INEQUALITIES FOR MESD WITHOUT SYMMETRIES

As promised in Sec. V A, we now derive the full set of noncontextuality inequalities for our operational MESD scenario when the symmetries of Eqs. (30)–(32) are not assumed. In Table III, we show a general data table for three binary measurements and four preparations which respect our operational equivalence. There are nine free parameters, since the probabilities in the last column are fixed by those in the first three.

The procedure from Appendix B yields the following set of inequalities over the nine free parameters, which are necessary and sufficient for the data to have been generated by a noncontextual model respecting operational equivalence Eq. (29):

TABLE III. Data table for our operational scenario with no symmetries assumed. There are nine free parameters.

	P_ϕ	P_ψ	$P_{\bar\phi}$	$P_{\bar\psi}$
$\phi\|M_\phi$	$1 - \epsilon_\phi$	c_ψ	$\epsilon_{\bar\phi}$	$1 - \epsilon_\phi + \epsilon_{\bar\phi} - c_\psi$
$\psi\|M_\psi$	c_ϕ	$1 - \epsilon_\psi$	$1 - c_{\bar\phi}$	$c_\phi - c_{\bar\phi} + \epsilon_\psi$
$g_\phi\|M_d$	s_ϕ	$1 - s_\psi$	$1 - s_{\bar\phi}$	$s_\phi - s_{\bar\phi} - s_\psi$

$$0 \le s_\phi \le 1,$$
$$0 \le s_{\bar\phi} \le 1,$$
$$0 \le s_\psi \le 1,$$
$$0 \le \epsilon_\phi \le c_\phi \le 1 - \epsilon_\phi,$$
$$0 \le \epsilon_{\bar\phi} \le c_{\bar\phi} \le 1 - \epsilon_{\bar\phi},$$
$$0 \le \epsilon_\psi \le c_\psi \le 1 - \epsilon_\psi,$$
$$0 \le s_\phi - s_{\bar\phi} + s_\psi \le 1,$$
$$0 \le c_\phi - c_{\bar\phi} + \epsilon_\psi,$$
$$0 \le c_\psi + s_{\bar\phi} - s_\psi + \epsilon_\phi,$$
$$0 \le c_\psi - s_{\bar\phi} + s_\psi + \epsilon_\phi,$$
$$0 \le -c_\psi + s_\phi + s_\psi + \epsilon_{\bar\phi},$$
$$0 \le c_\phi + s_{\bar\phi} - s_\psi + \epsilon_\psi,$$
$$0 \le -c_{\bar\phi} + s_\phi + s_\psi + \epsilon_\psi,$$
$$0 \le c_\phi - s_{\bar\phi} + s_\psi + \epsilon_\psi,$$
$$0 \le -c_{\bar\phi} + c_\psi + \epsilon_\phi + \epsilon_\psi,$$
$$0 \le c_\phi - c_\psi + \epsilon_{\bar\phi} + \epsilon_\psi,$$
$$0 \le 2 - c_\psi - s_\phi - s_\psi + \epsilon_{\bar\phi},$$
$$0 \le 2 - c_{\bar\phi} - s_\phi - s_\psi + \epsilon_\psi,$$
$$0 \le -c_\phi + c_{\bar\phi} + c_\psi + \epsilon_\phi - \epsilon_{\bar\phi} - \epsilon_\psi,$$
$$0 \le 1 - c_\phi + c_{\bar\phi} - c_\psi - \epsilon_\phi + \epsilon_{\bar\phi} - \epsilon_\psi. \tag{D1}$$

Of course, these inequalities reproduce Eq. (35) if the symmetries are now imposed.

In deriving these inequalities, we assume the logical labeling of Eq. (33). If one drops the labeling condition, then the resulting inequalities are identical to the facets of the Bell polytope discussed in Sec. VIII (but have no practical relevance to minimum-error state discrimination).

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

[2] S. Kochen and E. Specker, *The Problem of Hidden Variables in Quantum Mechanics*, Indiana Univ. Math. J. **17**, 59 (1968).

[3] J. S. Bell, *On the Problem of Hidden Variables in Quantum Mechanics*, Rev. Mod. Phys. **38**, 447 (1966).

[4] R. W. Spekkens, *Contextuality for Preparations, Transformations, and Unsharp Measurements*, Phys. Rev. A **71**, 052108 (2005).

[5] R. W. Spekkens, *Negativity and Contextuality are Equivalent Notions of Nonclassicality*, Phys. Rev. Lett. **101**, 020401 (2008).

[6] M. F. Pusey, *Anomalous Weak Values Are Proofs of Contextuality*, Phys. Rev. Lett. **113**, 200401 (2014).

[7] J. Barrett, L. Hardy, and A. Kent, *No Signaling and Quantum Key Distribution*, Phys. Rev. Lett. **95**, 010503 (2005).

[8] A. Acin and L. Masanes, *Certified Randomness in Quantum Physics*, Nature (London) **540**, 213 (2016).

[9] H. Buhrman, R. Cleve, S. Massar, and R. de Wolf, *Nonlocality and Communication Complexity*, Rev. Mod. Phys. **82**, 665 (2010).

[10] 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).

[11] A. Chailloux, I. Kerenidis, S. Kundu, and J. Sikora, *Optimal Bounds for Parity-Oblivious Random Access Codes*, New J. Phys. **18**, 045003 (2016).

[12] 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.

[13] M. Howard, J. Wallman, V. Veitch, and J. Emerson, *Contextuality Supplies the 'Magic' for Quantum Computation*, Nature (London) **510**, 351 (2014).

[14] R. Raussendorf, *Contextuality in Measurement-Based Quantum Computation*, Phys. Rev. A **88**, 022322 (2013).

[15] M. J. Hoban, E. T. Campbell, K. Loukopoulos, and D. E. Browne, *Non-Adaptive Measurement-Based Quantum Computation and Multi-Party Bell Inequalities*, New J. Phys. **13**, 023014 (2011).

[16] R. W. Spekkens, *Quasi-Quantization: Classical Statistical Theories with an Epistemic Restriction* (Springer, Dordrecht, 2016), pp. 83–135.

[17] Such models are ψ-epistemic, in the terminology of Ref. [18].

[18] N. Harrigan and R. W. Spekkens, *Einstein, Incompleteness, and the Epistemic View of Quantum States*, Found. Phys. **40**, 125 (2010).

[19] C. H. Bennett, *Quantum Cryptography Using Any Two Nonorthogonal States*, Phys. Rev. Lett. **68**, 3121 (1992).

[20] T. S. Cubitt, D. Leung, W. Matthews, and A. Winter, *Improving Zero-Error Classical Communication with Entanglement*, Phys. Rev. Lett. **104**, 230503 (2010).

[21] 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).

[22] D. Schmid, R. W. Spekkens, and E. Wolfe, *All the Noncontextuality Inequalities for Arbitrary Prepare-and-Measure Experiments with Respect to Any Fixed Sets of Operational Equivalences*, arXiv:1710.08434.

[23] C. W. Helstrom, *Quantum Detection and Estimation Theory*, J. Stat. Phys. **1**, 231 (1969).

[24] O. J. E. Maroney, *How Statistical Are Quantum States?*, arXiv:1207.6906.

[25] M. S. Leifer and O. J. E. Maroney, *Maximally Epistemic Interpretations of the Quantum State and Contextuality*, Phys. Rev. Lett. **110**, 120401 (2013).

[26] J. Barrett, E. G. Cavalcanti, R. Lal, and O. J. E. Maroney, *No ψ-Epistemic Model Can Fully Explain the Indistinguishability of Quantum States*, Phys. Rev. Lett. **112**, 250403 (2014).

[27] A mathematically equivalent version of this upper bound was previously proven under different assumptions in Refs. [26,28]. The former article considered the assumption that this inequality is saturated as a constraint on ontological models, which they termed "maximal ψ-epistemicity." Note that this constraint is different from the constraint considered in Ref. [25] even though it has the same name.

[28] M. Ringbauer, B. Duffus, C. Branciard, E. G. Cavalcanti, A. G. White, and A. Fedrizzi, *Measurements on the Reality of the Wavefunction*, Nat. Phys. **11**, 249 (2015).

[29] Whereas we believe that the assumption of universal noncontextuality is well motivated (namely, by Leibniz's principle of the identity of indiscernibles), it is unclear to us whether any motivation can be given for maximal ψ-epistemicity that is not simultaneously a motivation for universal noncontextuality. Therefore, unlike Ref. [24], we remain unconvinced that the assumption that Eq. (27) holds for every pair of quantum states is interesting in its own right.

[30] This is the sort of explanation one obtains in the toy theory model of the single qubit stabilizer subtheory of quantum theory [31] or the Kochen-Specker model of a single qubit [2]. Note that this is not the only way to explain the degree of confusability; the response function for $|\phi\rangle$ might be nontrivial outside the ontic support of the distribution representing $|\phi\rangle$ and even indeterministic in that region, and if so, one can have a nonzero confusability even though μ_ϕ and μ_ψ have disjoint ontic supports.

[31] R. W. Spekkens, *Evidence for the Epistemic View of Quantum States: A Toy Theory*, Phys. Rev. A **75**, 032110 (2007).

[32] R. W. Spekkens and T. Rudolph, *Degrees of Concealment and Bindingness in Quantum Bit Commitment Protocols*, Phys. Rev. A **65**, 012310 (2001).

[33] D. Mayers, *Unconditionally Secure Quantum Bit Commitment Is Impossible*, Phys. Rev. Lett. **78**, 3414 (1997).

[34] H.-K. Lo and H. F. Chau, *Is Quantum Bit Commitment Really Possible?*, Phys. Rev. Lett. **78**, 3410 (1997).

[35] A. Ambainis, *A New Protocol and Lower Bounds for Quantum Coin Flipping, Special Issue on STOC 2001*, J. Comput. Syst. Sci. **68**, 398 (2004).

[36] A. Chailloux, I. Kerenidis, and J. Sikora, *Lower Bounds for Quantum Oblivious Transfer*, Quantum Inf. Comput. **13**, 158 (2013).

[37] The reason there is no possibility of proving the Kochen-Specker theorem with projective measurements in dimension two is that no projector appears in more than a single context [2,18]. By contrast, it is known that there are proofs of the failure of preparation noncontextuality that hold even in two-dimensional Hilbert spaces [4], and the proof we present here is of this type.

[38] 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).

[39] M. D. Mazurek, M. F. Pusey, K. J. Resch, and R. W. Spekkens, *Experimentally Bounding Deviations from Quantum Theory in the Landscape of Generalized Probabilistic Theories*, arXiv:1710.05948.

[40] M. F. Pusey, *The Robust Noncontextuality Inequalities in the Simplest Scenario*, arXiv:1506.04178.

[41] L. Hardy, *Quantum Theory from Five Reasonable Axioms*, arXiv:quant-ph/0101012.

[42] J. Barrett, *Information Processing in Generalized Probabilistic Theories*, Phys. Rev. A **75**, 032304 (2007).

[43] 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).

[44] Faacets, http://www.faacets.com.

[45] M. S. Leifer and R. W. Spekkens, *Pre- and Post-Selection Paradoxes and Contextuality in Quantum Mechanics*, Phys. Rev. Lett. **95**, 200405 (2005).

[46] M. F. Pusey and M. S. Leifer, *Logical Pre- and Post-Selection Paradoxes Are Proofs of Contextuality*, EPTCS**195**, 295 (2015).

[47] These assumptions for B_d ensure that the relationship we derive between s_q and c_q will saturate the bound on s_q implied by any noncontextual model. If we had not used this argument, we would obtain the same relationship, but only as a bound on s_q, not as the saturating equality. This is easily verified explicitly, e.g., by taking $\epsilon = 0$ in Appendix B. However, including two more ontic states requires considerably more algebra.

[48] A. Krishna, Master's thesis, University of Waterloo, 2015.

[49] A. Krishna, R. W. Spekkens, and E. Wolfe, *Deriving Robust Noncontextuality Inequalities from Algebraic Proofs of the Kochen-Specker Theorem: The Peres-Mermin Square*, arXiv:1704.01153.

[50] C. Budroni and A. Cabello, *Bell Inequalities from Variable-Elimination Methods*, J. Phys. A **45**, 385304 (2012).

[51] R. Chaves, L. Luft, and D. Gross, *Causal Structures from Entropic Information: Geometry and Novel Scenarios*, New J. Phys. **16**, 043001 (2014).

[52] R. Chaves, *Polynomial Bell Inequalities*, Phys. Rev. Lett. **116**, 010402 (2016).

[53] E. Wolfe, R. W. Spekkens, and T. Fritz, *The Inflation Technique for Causal Inference with Latent Variables*, arXiv:1609.00672.

[54] C. M. Lee and R. W. Spekkens, *Causal Inference via Algebraic Geometry: Feasibility Tests for Functional Causal Structures with Two Binary Observed Variables*, J. Causal Infer. **5**, 20160013 (2017).

[55] M. Navascues, S. Pironio, and A. Acin, *A Convergent Hierarchy of Semidefinite Programs Characterizing the Set of Quantum Correlations*, New J. Phys. **10**, 073013 (2008).