Beyond the Cabello-Severini-Winter framework: Making sense of contextuality without sharpness of measurements

Ravi Kunjwal

Perimeter Institute for Theoretical Physics,
31 Caroline Street North, Waterloo, Ontario, Canada, N2L 2Y5.
September 4, 2019

We develop a hypergraph-theoretic framework for Spekkens contextuality applied to Kochen-Specker (KS) type scenarios that goes beyond the Cabello-Severini-Winter (CSW) framework. To do this, we add new hypergraph-theoretic ingredients to the CSW framework. We then obtain noise-robust non-contextuality inequalities in this generalized framework by applying the assumption of (Spekkens) noncontextuality to both preparations and measurements. The resulting framework goes beyond the CSW framework in both senses, conceptual and technical. On the conceptual level: 1) as in any treatment based on the generalized notion of noncontextuality à la Spekkens, we relax the assumption of outcome determinism inherent to the Kochen-Specker theorem but retain measurement noncontextuality, besides introducing preparation non-contextuality, 2) we do not require the *exclusivity principle* – that pairwise exclusive measurement events must all be mutually exclusive – as a fundamental constraint on measurement events of interest in an experimental test of contextuality, given that this property is not true of general quantum measurements, and 3) as a result, we do not need to presume that measurement events of interest are “sharp” (for any definition of sharpness), where this notion of sharpness is meant to imply the exclusivity principle. On the technical level, we go beyond the CSW framework in the following senses: 1) we introduce a source events hypergraph – besides the measurement events hypergraph usually considered – and define a new operational quantity Corr that appears in our inequalities, 2) we define a new hypergraph invariant – the *weighted max-predictability* – that is necessary for our analysis and appears in our inequalities, and 3) our noise-robust noncontextuality inequalities quantify tradeoff relations between three operational quantities – Corr , R , and p_0 – only one of which (namely, R) corresponds to the Bell-

Kochen-Specker functionals appearing in the CSW framework; when $\text{Corr} = 1$, the inequalities formally reduce to CSW type bounds on R . Along the way, we also consider in detail the scope of our framework vis-à-vis the CSW framework, particularly the role of Specker’s principle in the CSW framework, i.e., what the principle means for an operational theory satisfying it and why we don’t impose it in our framework.

Contents

1	Introduction	2
2	Spekkens framework	5
2.1	Operational theory	5
2.2	Ontological model	6
2.3	Representation of coarse-graining . . .	7
2.3.1	Coarse-graining of measurements	7
2.3.2	Coarse-graining of preparations	8
2.4	Joint measurability (or compatibility)	9
2.5	Noncontextuality	9
2.6	An example of Spekkens contextuality: the fair coin flip inequality	10
2.7	Connection to Bell scenarios	12
3	Hypergraph approach to Kochen-Specker scenarios in the Spekkens framework	13
3.1	Measurements	13
3.1.1	Classification of probabilistic models	14
3.1.2	Distinguishing two consequences of Specker’s principle: Structural Specker’s principle vs. Statistical Specker’s principle	15
3.1.3	What does it mean for an operational theory to satisfy structural/statistical Specker’s principle?	16

Ravi Kunjwal: rkunjwal@perimeterinstitute.ca

3.1.4	Remark on the classification of probabilistic models: why we haven't defined "quantum models" as those obtained from projective measurements	20
3.1.5	Scope of this framework	20
3.2	Sources	21
4	A key hypergraph invariant: the weighted max-predictability	23
5	Noise-robust noncontextuality inequalities	24
5.1	Key notions from CSW	24
5.2	Key notion <i>not</i> from CSW: source-measurement correlation, Corr	25
5.3	Obtaining the noise-robust noncontextuality inequalities	25
5.3.1	Expressing operational quantities in ontological terms	25
5.3.2	Derivation of the noncontextual tradeoff for any graph G	26
5.3.3	When is the noncontextual tradeoff violated?	27
5.4	Example: KCBS scenario	27
6	Discussion	30
6.1	Measurement-measurement correlations vs. source-measurement correlations	30
6.2	Can our noise-robust noncontextuality inequalities be saturated by a noncontextual ontological model?	30
6.2.1	The special case of facet-defining Bell-KS inequalities: Corr=1	30
6.2.2	The general case: Corr < 1	30
6.3	Can trivial POVMs ever violate these noncontextuality inequalities?	31
6.3.1	The case $p \in \mathcal{C}(\Gamma_G)$	31
6.3.2	The case $p \in \text{ConvHull}(\mathcal{G}(\Gamma_G) _{\text{ind}})$	31
6.3.3	The general case $p \in \mathcal{G}(\Gamma_G)$	31
7	Conclusions	32
	Acknowledgments	33
A	Status of KS-contextuality as an experimentally testable notion of nonclassicality for POVMs in quantum theory	33
A.1	Limitations of KS-contextuality vis-à-vis POVMs	34
A.1.1	KS-contextuality for POVMs in the literature	34
A.1.2	Classifying probabilistic models: restriction of quantum models to PVMs	35

A.2	Robustness of Bell nonlocality vis-à-vis POVMs	36
B	Ontological models without respecting coarse-graining relations	37
B.1	How to construct a "KS-noncontextual" ontological model of the KCBS experiment [47] without coarse-graining relations	37
B.2	How to construct a "preparation and measurement noncontextual" ontological model without coarse-graining relations	37
C	Trivial POVMs	38
C.1	Bell-CHSH scenario	38
C.2	CHSH-type contextuality scenario: 4-cycle	38
D	The KS-uncolourable hypergraph Γ_{18}	40
	References	41

1 Introduction

To say that quantum theory is counterintuitive, or that it requires a revision of our classical intuitions, requires us to be mathematically precise in our definition of these classical intuitions. Once we have a precise formulation of such *classicality*, we can begin to investigate those features of quantum theory that power its *nonclassicality*, i.e., its departure from our classical intuitions, and thus prove *theorems* about such nonclassicality. To the extent that a physical theory is provisional, likely to be replaced by a better theory in the future, it also makes sense to articulate such notions of classicality in as operational a manner as possible. By 'operational', we refer to a formulation of the theory that takes the operations – preparations, measurements, transformations – that can be carried out in an experiment as primitives and which specifies the manner in which these operations combine to produce the data in the experiment. Such an operational formulation often suggests generalizations of the theory that can then be used to better understand its axiomatics [1–3]. At the same time, an operational formulation also lets us articulate our notions of nonclassicality in a manner that is experimentally testable and thus allows us to leverage this nonclassicality in applications of the theory. Indeed, a key area of research in quantum foundations and quantum information is the development of methods to assess nonclassicality in an experiment under minimal assumptions on the operational theory describing it. The paradigmatic example of this is the case of Bell's theorem and Bell experiments [4–11], where any operational theory that is non-signalling between the different spacelike separated wings of the experiment

is allowed. The notion of classicality at play in Bell’s theorem is the assumption of *local causality*: any non-signalling theory that violates the assumption of local causality is said to exhibit nonclassicality by the lights of Bell’s theorem.

More recently, much work [12–17] has been devoted to obtaining constraints on operational statistics that follow from a generalized notion of noncontextuality proposed by Spekkens [18]. This notion of classicality [18] has its roots in the Kochen-Specker (KS) theorem [19], a no-go theorem that rules out the possibility that a deterministic underlying ontological model [20] could reproduce the operational statistics of (projective) quantum measurements in a manner that satisfies the assumption of *KS-noncontextuality*. KS-noncontextuality is the notion of classicality at play in the Kochen-Specker theorem. The Spekkens framework abandons the assumption of outcome determinism [18] – the idea that the ontic state of a system fixes the outcome of any measurement deterministically – that is intrinsic to KS-noncontextuality. It also applies to general operational theories and extends the notion of noncontextuality to general experimental procedures – preparations, transformations, and measurements – rather than measurements alone.

Parallel to work along the lines of Spekkens [18], work seeking to directly operationalize the Kochen-Specker theorem (rather than revising the notion of noncontextuality at play) culminated in two recent approaches that classify theories by the degree to which they violate the assumption of KS-noncontextuality: the graph-theoretic framework of Cabello, Severini, and Winter (CSW) [21, 22], where a general approach to obtaining graph-theoretic bounds on linear Bell-KS functionals was proposed, and the related hypergraph framework of Acín, Fritz, Leverrier, and Sainz (AFLS) [23], where an approach to characterizing sets of correlations was proposed. The CSW framework relates well-known graph invariants to: 1) upper bounds on Bell-KS inequalities that follow from KS-noncontextuality, 2) upper bounds on maximum quantum violations of these inequalities that can be obtained from projective measurements, and 3) upper bounds on their violation in general probabilistic theories [24] – denoted *E1* – which satisfy the “exclusivity principle” [22]. Complementary to this, the AFLS framework uses graph invariants in the service of deciding whether a given assignment of probabilities to measurement outcomes in a KS-contextuality experiment belongs to a particular set of correlations; they showed that membership in the quantum set of correlations (defined *only* for projective measurements in quantum theory) cannot be witnessed by a graph invariant, cf. Theorem 5.1.3 of Ref. [23]. Another recent approach due to Abramsky and Brandenburger [25] employs sheaf-theoretic ideas to formulate KS-contextuality.

A key achievement of the frameworks of Refs. [22,

23, 25] is a formal unification of Bell scenarios with KS-contextuality scenarios, treating them on the same footing. Indeed, the perspective there is to consider Bell scenarios as a special case of KS-contextuality scenarios. What is lost in this mathematical unification, however, is the fact that Bell-locality and KS-noncontextuality have physically distinct, if related, motivations. The physical situation that Bell’s theorem refers to requires (at least) two spacelike separated labs (where local measurements are carried out) so that the assumption of local causality (or Bell-locality) can be applied.¹ On the other hand, the physical situation that the Kochen-Specker theorem refers to does not require spacelike separation as a necessary ingredient and one can therefore consider experiments in a single lab. However, the assumption of KS-noncontextuality entails outcome determinism [18], something not required by local causality in Bell scenarios.² This difference in the physical situation for the two kinds of experiments is one of the reasons for generalizing KS-noncontextuality to the notion of noncontextuality in the Spekkens framework [18] (so that outcome determinism is not assumed) while leaving Bell’s notion of local causality untouched.

In the present paper we build a bridge from the CSW approach, where KS-noncontextual correlations are bounded by Bell-KS inequalities, to noise-robust noncontextuality inequalities in the Spekkens framework [18]. That is, we show how the constraints from KS-noncontextuality in the framework of Ref. [22] translate to constraints from generalized noncontextuality in the framework of Ref. [18]. The resulting operational criteria for contextuality à la Spekkens

¹What do we mean by whether an assumption “can be applied”? Of course, mathematically, one can “apply” any assumption one wants in the service of proving a theorem. But insofar as the mathematics here is trying to model a real experiment, the consistency of those assumptions with some essential facts of the experiment is the minimal requirement for any no-go theorem derived from such assumptions to be physically interesting. Hence, in the presence of signalling (implying the absence of spacelike separation), it makes no sense to assume local causality in a Bell experiment and derive the resulting Bell inequalities: such an assumption on the ontological model is already in conflict with the fact of signalling across the labs and no Bell inequalities are needed to witness this fact. Bell inequalities only become physically interesting when the theories being compared relative to them are all non-signalling: if the experiment itself is signalling, any non-signalling description – locally causal, quantum, or in a general probabilistic theory (GPT) – is *ipso facto* ruled out.

²Note that this assumption of outcome determinism doesn’t affect the conclusions in a Bell scenario even if one adopted it because of Fine’s theorem [26]: a locally deterministic ontological model entails the same set of (Bell-local) correlations as a locally causal ontological model. Relaxing outcome determinism, however, doesn’t mean the same thing for the kinds of experiments envisaged by the Kochen-Specker theorem – in particular, it doesn’t mean that models satisfying factorizability à la Ref. [25] are the most general outcome-indeterministic models – and thus considerations parallel to Fine’s theorem [26] do not apply, cf. [27, 28].

are noise-robust and therefore applicable to arbitrary positive operator-valued measures (POVMs) and mixed states in quantum theory. Note that the insights gleaned from frameworks such as those of Refs. [22, 23, 25] regarding Bell nonlocality require no revision in our approach. It is *only* in the application of such frameworks (in particular, the CSW framework) to the question of contextuality that we seek to propose an alternative hypergraph framework (formalizing Spekkens contextuality [18]) that is more operationally motivated for experimental situations where one cannot appeal to spacelike separation to justify locality of the measurements.³ For Kochen-Specker type experimental scenarios, we will consider the twin notions of preparation noncontextuality and measurement noncontextuality – taken together as a notion of classicality – to obtain noise-robust noncontextuality inequalities that generalize the KS-noncontextuality inequalities of CSW. These inequalities witness nonclassicality even when quantum correlations arising from arbitrary (i.e., possibly nonprojective) quantum measurements on any quantum state are allowed. A key innovation of this approach is that it treats all measurements in an operational theory on an equal footing. No definition of “sharpness” [29–31] is needed to justify or derive noncontextuality inequalities in this approach. Furthermore, if certain idealizations are presumed about the operational statistics, then these inequalities formally recover the usual Bell-KS inequalities à la CSW. The Bell-KS inequalities can be viewed as an instance of the classical marginal problem [25–27, 32, 33], i.e., as constraints on the (marginal) probability distributions over subsets of a set of observables that follow from requiring the existence of global joint probability distribution over the set of all observables. Since the Bell-KS inequalities are only recovered under certain idealizations, but not otherwise, the noise-robust noncontextuality inequalities we obtain *cannot* in general be viewed as arising from a classical marginal problem. Hence, they cannot be understood within existing frameworks that rely on this (reduction to the classical marginal problem) property to formally unify the treatment of Bell-nonlocality and KS-contextuality [22, 23, 25]. This is a *crucial* distinction relative to the usual Bell-KS inequality type witnesses of KS-contextuality.

This paper is based on a previous contribution [16] that laid the conceptual groundwork for the progress we make here. Besides the noise-robust noncontextuality inequalities that generalize constraints from KS-noncontextuality in the CSW framework leveraging the graph invariants of CSW [22] (cf. Section 5),

³Nor the sharpness of the measurements to justify outcome determinism. We discuss these issues in detail – in particular, the physical basis of KS-noncontextuality vis-à-vis Bell-locality and how that influences our framework – in Appendix A for the interested reader.

the contributions of this paper also include:

- An exposition of Specker’s principle and how different implications of it (e.g., consistent exclusivity [23]) for a given operational theory arise in the hypergraph framework (cf. Sections 3.1.2 and 3.1.3), in particular the results in Theorems 1, 2, and Corollary 1.
- Introduction of a hypergraph invariant – the weighted max-predictability – that is key to our noise-robust noncontextuality inequalities, cf. Section 4. This invariant is also key to the hypergraph framework of Ref. [34] which is complementary to the present framework.
- A detailed discussion of how KS-noncontextuality for POVMs has been previously treated in the literature and the limitations of those treatments, cf. Appendices A and C. Also, unlike for the case of KS-noncontextuality inequalities, we show that trivial POVMs can never violate our noise-robust noncontextuality inequalities, cf. Section 6.3.
- A discussion of coarse-graining relations in Section 2.3 and their importance for contextuality no-go theorems, in particular a discussion of ontological models that do not respect coarse-graining relations in Appendix B. We show, in Appendix B, how relaxing the constraint from coarse-graining relations on an ontological model renders either notion of noncontextuality – whether Kochen-Specker [19] or Spekkens [18] – vacuous.
- A discussion, by example, of why our generalization of the CSW framework cannot accommodate contextuality scenarios that are KS-uncolourable in Appendix D and why one needs a distinct framework, i.e., the framework of Ref. [34], to treat KS-uncolourable scenarios.

The structure of this paper follows: Section 2 reviews the Spekkens framework for generalized noncontextuality [18]. Section 3 introduces a hypergraph framework that shares features of traditional frameworks for KS-contextuality [22, 23] but is also augmented (relative to these traditional frameworks) with the ingredients necessary for obtaining noise-robust noncontextuality inequalities. In particular, its subsections 3.1.2 and 3.1.3 discuss Specker’s principle [35] and define its different implications for contextuality scenarios à la Ref. [23]. Section 4 defines a new hypergraph invariant – the weighted max-predictability – that we need later on as a crucial new ingredient in our inequalities. Section 5 obtains noise-robust noncontextuality inequalities within the framework defined in Section 3 and using the hypergraph invariant of Section 4 in addition to two graph invariants from the CSW framework [22]. These inequalities can be

seen as special cases of the general approach outlined in Ref. [16]. In Section 6, we include discussions on various features of our noise-robust noncontextuality inequalities, in particular the fact that trivial POVMs can never violate them. Section 7 concludes with some open questions and directions for future research.

2 Spekkens framework

We concern ourselves with prepare-and-measure experiments. A schematic of such an experiment is shown in Figure 1 where, for the sake of simplicity, we imagine a single source device that can perform any preparation procedure of interest (rather than a collection of source devices, each implementing a particular preparation procedure) and a single measurement device that can perform any measurement procedure of interest (rather than a collection of measurement devices, each implementing a particular measurement procedure). Note that this is just a conceptual abstraction: in particular, the various possible measurement settings on the measurement device may, for example, correspond to incompatible measurement procedures in quantum theory. The fact that we represent the different measurement settings by choices of knob settings $M \in \mathbb{M}$ on a single measurement device *does not* mean that it's physically possible to implement all the measurement procedures represented by $\mathbb{M}$ jointly; it only means that the experimenter can choose to implement any of the measurements in the set $\mathbb{M}$ in a particular prepare-and-measure experiment. The same is true for our abstraction of preparation procedures to knob settings ($S \in \mathbb{S}$) and outcomes ($s \in V_S$) of a single source device: it's not that the same device can physically implement all possible preparation procedures; it's just that an experimenter can choose to implement any procedure in the set $\mathbb{S}$ in a particular prepare-and-measure experiment.

We will consider two levels of description of prepare-and-measure experiments represented by Fig. 1: *operational* and *ontological*. The operational description will be specified by an operational theory that takes source and measurement devices as primitives and describes the experiment solely in terms of the probabilities associated to their input/output behaviour. The ontological description will be specified by an ontological model that takes the system that passes between the source and measurement devices as primitive and describes the experiment in terms of probabilities associated to properties of this system, deriving the operational description as a consequence of coarse-graining over these properties. Let us look at each description in turn.

2.1 Operational theory

We now describe the various components of Fig. 1 in more detail. The source device has a *source setting*

Figure 1: A prepare-and-measure experiment.

labelled by S that can be chosen from a set $\mathbb{S}$. The set $\mathbb{S}$ represents, in general, some subset of the set of all source settings, $\mathcal{S}$, that are admissible in the operational theory, i.e., $\mathbb{S} \subseteq \mathcal{S}$. In a particular prepare-and-measure experiment, $\mathbb{S}$ will typically be a finite set of source settings. Choosing the setting S prepares a system according to an ensemble of preparation procedures, denoted $\{(p(s|S), P_{[s|S]})\}_{s \in V_S}$, where $\{p(s|S)\}_{s \in V_S}$ is a probability distribution over the preparation procedures $\{P_{[s|S]}\}_{s \in V_S}$ in the ensemble. This means that the source device has one classical input S and two outputs: one output is a classical label $s \in V_S$ identifying the preparation procedure (in the ensemble $\{(p(s|S), P_{[s|S]})\}_{s \in V_S}$) that is carried out when source outcome s is observed for source setting S (this *source event* is denoted $[s|S]$), and the other output is a system prepared according to the source event $[s|S]$, i.e., preparation procedure $P_{[s|S]}$, with probability $p(s|S)$. Thus, the assemblage of possible ensembles that the source device can prepare can be denoted by $\{(p(s|S), P_{[s|S]})\}_{s \in V_S} \}_{S \in \mathbb{S}}$.

On the other hand, the measurement device has two inputs, one a classical input $M \in \mathbb{M}$ specifying the choice of measurement setting to be implemented, and the other input receives the system prepared according to preparation procedure $P_{[s|S]}$ and on which this measurement M is carried out. The measurement device has one classical output $m \in V_M$ denoting the outcome of the measurement M implemented on a system prepared according to $P_{[s|S]}$, and which occurs with probability $p(m|M, S, s)$. The set $\mathbb{M}$ represents,

in general, some subset of the set of all measurement settings, $\mathcal{M}$, that are admissible in the operational theory, i.e., $\mathbb{M} \subseteq \mathcal{M}$. In a particular prepare-and-measure experiment, $\mathbb{M}$ will typically be a finite set of measurement settings.

We will be interested in the operational joint probability $p(m, s|M, S) \equiv p(m|M, S, s)p(s|S)$ for this prepare-and-measure experiment for various choices of $M \in \mathbb{M}$, $S \in \mathbb{S}$. Note how this operational description takes as primitive the *operations* carried out in the lab and restricts itself to specifying the probabilities of classical outcomes (i.e., m, s) given some interventions (i.e., classical inputs, M, S). So far, we haven't assumed any structure on the operational theory describing the schematic of Fig. 1 beyond the fact that it is a catalogue of input/output probabilities

$$\{\{p(m, s|M, S) \in [0, 1]\}_{m \in V_M, s \in V_S}\}_{M \in \mathbb{M}, S \in \mathbb{S}}$$

for various interventions $S \in \mathbb{S}$ and $M \in \mathbb{M}$ that we will consider in a prepare-and-measure experiment. We now require more structure in the operational theory underlying this experiment, beyond a mere specification of these probabilities.

We require that the operational theory admits *equivalence relations* that partition experimental procedures of any type, whether preparations or measurements, into equivalence classes of that type. These equivalence relations are defined relative to the operational probabilities (not necessarily restricted to a particular prepare-and-measure experiment) that are admissible in the theory. We will call these equivalence relations “operational equivalences”, in keeping with standard terminology [18]. This means that any distinctions of labels between procedures in an equivalence class of procedures do not affect the operational probabilities associated with the procedures. We specify these equivalence relations for measurement and preparation procedures below.

Two measurement events $[m|M]$ and $[m'|M']$ are said to be operationally equivalent, denoted $[m|M] \simeq [m'|M']$, if there exists no source event in the operational theory that can distinguish them, i.e.,

$$p(m, s|M, S) = p(m', s|M', S) \quad \forall [s|S], s \in V_S, S \in \mathcal{S}. \quad (1)$$

Note that the statistical indistinguishability of $[m|M]$ and $[m'|M']$ must hold for all possible source settings $\mathcal{S}$ in the operational theory, not merely the source settings $\mathbb{S}$ that are of direct interest in a particular prepare-and-measure experiment. Hence, the “distinction of labels”, $[m|M]$ or $[m'|M']$, is empirically inconsequential since the two procedures are, in principle, indistinguishable by the lights of the operational theory.

Similarly, two source events $[s|S]$ and $[s'|S']$ are said to be operationally equivalent, denoted $[s|S] \simeq [s'|S']$, if there exists no measurement event in the opera-

tional theory that can distinguish them, i.e.,

$$p(m, s|M, S) = p(m, s'|M, S'), \quad \forall [m|M], m \in V_M, M \in \mathcal{M}. \quad (2)$$

Again, the statistical indistinguishability of $[s|S]$ and $[s'|S']$ must hold for all possible measurement settings $\mathcal{M}$, not merely those (i.e., $\mathbb{M}$) that are of direct interest in a particular prepare-and-measure experiment. Similar to measurement events, the “distinction of labels”, $[s|S]$ or $[s'|S']$, is empirically inconsequential since the two procedures are, in principle, indistinguishable by the lights of the operational theory.

Given this equivalence structure for preparation and measurement procedures in the operational theory, we can now formalize the notion of a *context*:

Definition 1. A context is any distinction of labels between operationally equivalent procedures in the operational theory.

To see concrete examples of the kinds of contexts that will be of interest to us in this paper, consider quantum theory. Any mixed quantum state admits multiple convex decompositions in terms of other quantum states, i.e., it can be prepared by coarse-graining over distinct ensembles of quantum states, each ensemble denoted by a different label. In this case, the “distinction of labels” between different decompositions denotes a distinction of preparation ensembles, which instantiates our notion of a preparation context. Similarly, a given positive operator can be implemented by different positive operator-valued measures (POVMs), and the distinction of labels denoting these different POVMs instantiates our notion of a measurement context.

2.2 Ontological model

Given the operational description of the experiment in terms of probabilities $p(m, s|M, S)$, we want to explore the properties of any underlying ontological model for this operational description. Any such ontological model, defined within the ontological models framework [20], takes as primitive the *physical system* (rather than *operations* on it) that passes between the source and measurement devices, i.e., its basic objects are *ontic states* of the system, denoted $\lambda \in \Lambda$, that represent intrinsic properties of the physical system. When a preparation procedure $[s|S]$ is carried out, the source device samples from the space of ontic states Λ according to a probability distribution $\{\mu(\lambda|S, s) \in [0, 1]\}_{\lambda \in \Lambda}$, where $\sum_{\lambda \in \Lambda} \mu(\lambda|S, s) = 1$, and the joint distribution over s and λ given S , i.e., $\{\mu(\lambda, s|S)\}_{\lambda \in \Lambda}$, is given by $\mu(\lambda, s|S) \equiv \mu(\lambda|S, s)p(s|S)$. On the other hand, when a system in ontic state λ is input to the measurement device with measurement setting $M \in \mathbb{M}$, the probability distribution over the measurement outcomes is given by $\{\xi(m|M, \lambda) \in [0, 1]\}_{m \in V_M}$, where

$\sum_{m \in V_M} \xi(m|M, \lambda) = 1$. The operational statistics

$$\{p(m, s|M, S) \in [0, 1]\}_{m \in V_M, s \in V_S} \}_{M \in \mathbb{M}, S \in \mathbb{S}}$$

results from a coarse-graining over λ , i.e.,

$$p(m, s|M, S) = \sum_{\lambda \in \Lambda} \xi(m|M, \lambda) \mu(\lambda, s|S), \quad (3)$$

for all $m \in V_M, s \in V_S, M \in \mathbb{M}, S \in \mathbb{S}$.

Note that the definition of an ontological model above extends to the definition of an ontological model of the operational *theory* (as opposed to a particular fragment of the theory representing the experiment) when we take $\mathbb{M} = \mathcal{M}$ and $\mathbb{S} = \mathcal{S}$.

2.3 Representation of coarse-graining

We will now specify how coarse-graining of procedures in a prepare-and-measure experiment is represented in its description, whether operational or ontological. Namely, if a procedure is defined as a coarse-graining of other procedures, then we require that the representation of such a procedure is defined by the same coarse-graining of the representation of the other procedures.⁴ Implicit in this discussion is the assumption that the operational theory allows one to define new procedures in the set $\mathcal{M}$ or $\mathcal{S}$ by coarse-graining other procedures in these sets, i.e., both $\mathcal{M}$ and $\mathcal{S}$ are closed under coarse-grainings. In particular, one can consider coarse-graining measurement and source settings (belonging to sets $\mathbb{M}$ and $\mathbb{S}$, respectively) actually implemented in the lab to define new measurement and source settings that belong to $\mathcal{M} \setminus \mathbb{M}$ and $\mathcal{S} \setminus \mathbb{S}$, respectively.⁵

2.3.1 Coarse-graining of measurements

Let us see how this works for the case of measurement procedures: if a measurement procedure M with measurement events $\{[m|M]\}_{m \in V_M}$ is *defined* as a coarse-graining of another measurement procedure $\tilde{M}$ with measurement events $\{[\tilde{m}|\tilde{M}]\}_{\tilde{m} \in V_{\tilde{M}}}$, symbolically denoted by

$$[m|M] \equiv \sum_{\tilde{m}} p(m|\tilde{m})[\tilde{m}|\tilde{M}],$$

$$\text{where } \forall m, \tilde{m} : p(m|\tilde{m}) \in \{0, 1\}, \sum_m p(m|\tilde{m}) = 1, \quad (4)$$

⁴Quantum theory is an example of an operational theory that satisfies this requirement because of the linearity of the Born rule with respect to both preparations and measurements. The same is true, more generally, of general probabilistic theories (GPTs) [1, 24]. We require this feature in any ontological model as well, regardless of its (non)contextuality.

⁵Similarly, we also allow probabilistic mixtures of (preparation or measurement) procedures in the operational theory to define new procedures, i.e., the theory is convex. See the last paragraph of Section 2.5 for the role of this convexity in experimental tests of contextuality and Section 2.6 for an example where a probabilistic mixture of measurement settings is required in a proof of contextuality.

then its representation in the operational description as well as in the ontological description satisfies this coarse-graining relation.⁶ More explicitly, the coarse-graining relation of Eq. (4) denotes the following post-processing of $\tilde{M}$: for each $m \in V_M$, relabel each outcome $\tilde{m} \in V_{\tilde{M}}$ to outcome m with probability $p(m|\tilde{m}) \in \{0, 1\}$; the logical disjunction of those $\tilde{m}$ which are relabelled to m with probability 1 then defines the measurement event $[m|M]$. Now, in the operational theory, this post-processing is represented by

$$\forall [s|S], \text{ where } s \in V_S, S \in \mathcal{S} :$$

$$p(m, s|M, S) \equiv \sum_{\tilde{m}} p(m|\tilde{m}) p(\tilde{m}, s|\tilde{M}, S), \quad (5)$$

and in the ontological model it is represented by

$$\forall \lambda \in \Lambda : \xi(m|M, \lambda) \equiv \sum_{\tilde{m}} p(m|\tilde{m}) \xi(\tilde{m}|\tilde{M}, \lambda). \quad (6)$$

As an example, consider a three-outcome measurement $\tilde{M}$ with outcomes $\tilde{m} \in \{1, 2, 3\}$, which can be classically post-processed to obtain a two-outcome measurement M with outcomes $m \in \{0, 1\}$, such that $p(m = 0|\tilde{m} = 1) = p(m = 0|\tilde{m} = 2) = 1$ and $p(m = 1|\tilde{m} = 3) = 1$. The measurement events of M are then just

$$[m = 0|M] \equiv [\tilde{m} = 1|\tilde{M}] + [\tilde{m} = 2|\tilde{M}], \quad (7)$$

$$[m = 1|M] \equiv [\tilde{m} = 3|\tilde{M}], \quad (8)$$

where the “+” sign denotes (just as the summation sign in the definition of $[m|M]$ in Eq. (4) did) logical disjunction, i.e., measurement event $[m = 0|M]$ is said to occur when $[\tilde{m} = 1|\tilde{M}]$ or $[\tilde{m} = 2|\tilde{M}]$ occurs. The operational and ontological representations of these measurement events are then given by

$$\forall [s|S], \text{ where } s \in V_S, S \in \mathcal{S} :$$

$$p(m = 0, s|M, S) \equiv \sum_{\tilde{m}=1}^2 p(\tilde{m}, s|\tilde{M}, S), \quad (9)$$

$$p(m = 1, s|M, S) \equiv p(\tilde{m} = 3, s|\tilde{M}, S), \quad (10)$$

$$\forall \lambda \in \Lambda :$$

$$\xi(m = 0|M, \lambda) \equiv \sum_{\tilde{m}=1}^2 \xi(\tilde{m}|\tilde{M}, \lambda), \quad (11)$$

$$\xi(m = 1|M, \lambda) \equiv \xi(\tilde{m} = 3|\tilde{M}, \lambda). \quad (12)$$

This requirement on the representation of coarse-graining of measurements is particularly important (and often implicit) when the notion of a measurement context is instantiated by compatibility

⁶Note that Eq. (4) is *not* an operational equivalence between independent procedures. It is a *definition* of a new procedure obtained by coarse-graining another procedure.

(or joint measurability), as in the case of KS-contextuality, where one needs to consider coarse-grainings of distinct measurements. For example, consider a measurement setting M_{12} with outcomes $(m_1, m_2) \in V_1 \times V_2$ that is coarse-grained over m_2 to define an effective measurement setting $M_1^{(2)}$ with measurement events $\{[m_1|M_1^{(2)}]\}_{m_1 \in V_1}$. Symbolically, $[m_1|M_1^{(2)}] \equiv \sum_{m_2} [(m_1, m_2)|M_{12}]$, which is represented in the operational theory as $\forall [s|S] : p(m_1, s|M_1^{(2)}, S) \equiv \sum_{m_2} p((m_1, m_2), s|M_{12}, S)$ and in the ontological model as $\forall \lambda : \xi(m_1|M_1^{(2)}, \lambda) \equiv \sum_{m_2} \xi((m_1, m_2)|M_{12}, \lambda)$. Similarly, consider another measurement setting M_{13} with outcomes $(m_1, m_3) \in V_1 \times V_3$ that is coarse-grained over m_3 to define an effective measurement setting $M_1^{(3)}$ with measurement events $\{[m_1|M_1^{(3)}]\}_{m_1 \in V_1}$. Symbolically, $[m_1|M_1^{(3)}] \equiv \sum_{m_3} [(m_1, m_3)|M_{13}]$, which is represented in the operational theory as $\forall [s|S] : p(m_1, s|M_1^{(3)}, S) \equiv \sum_{m_3} p((m_1, m_3), s|M_{13}, S)$ and in the ontological model as $\forall \lambda : \xi(m_1|M_1^{(3)}, \lambda) \equiv \sum_{m_3} \xi((m_1, m_3)|M_{13}, \lambda)$.

Now, imagine that the following operational equivalence holds at the operational level: $[m_1|M_1^{(2)}] \simeq [m_1|M_1^{(3)}]$. KS-noncontextuality is then the assumption that $\sum_{m_2} \xi((m_1, m_2)|M_{12}, \lambda) = \sum_{m_3} \xi((m_1, m_3)|M_{13}, \lambda)$ (i.e., $\xi(m_1|M_1^{(2)}, \lambda) = \xi(m_1|M_1^{(3)}, \lambda)$) for all λ and that $\xi((m_1, m_2)|M_{12}, \lambda), \xi((m_1, m_3)|M_{13}, \lambda) \in \{0, 1\}$ for all λ . This assumption applied to multiple (compatible) subsets of a set of carefully chosen measurements can then provide a proof of the KS theorem, i.e., there exist sets of measurements in quantum theory such that their operational statistics cannot be emulated by a KS-noncontextual ontological model.

The key point here is this: the requirement that coarse-graining relations between measurements be respected by their representations in the ontological model is independent of the KS-(non)contextuality of the ontological model.⁷ However, this requirement is *necessary* for the assumption of KS-noncontextuality to produce a contradiction with quantum theory; on the other hand, a KS-contextual ontological model (while respecting the coarse-graining relations) can always emulate quantum theory. In this sense, the representation of coarse-grainings is baked into an ontological model from the beginning (just as it is baked into an operational description), before any claims about its (non)contextuality.⁸

⁷In our example, this requirement has to do with the definitions of $\xi(m_1|M_1^{(2)}, \lambda)$ and $\xi(m_1|M_1^{(3)}, \lambda)$, not their ontological equivalence. The ontological equivalence only comes into play when invoking KS-noncontextuality.

⁸One could, of course, choose to not respect the coarse-graining relations and define a notion of an ontological model without them. In such a model, one could treat every mea-

2.3.2 Coarse-graining of preparations

Let us now consider the representation of coarse-grainings for preparation procedures. This works in a way similar to the case of measurement procedures which we have already outlined. If an ensemble of source events $\{[s|S]\}_{s \in V_S}$ is defined as a coarse-graining of another ensemble, $\{[\tilde{s}|\tilde{S}]\}_{\tilde{s} \in V_{\tilde{S}}}$, symbolically denoted as

$$[s|S] \equiv \sum_{\tilde{s}} p(s|\tilde{s})[\tilde{s}|\tilde{S}], \text{ where} \\ \forall s, \tilde{s} : p(s|\tilde{s}) \in \{0, 1\}, \sum_s p(s|\tilde{s}) = 1, \quad (13)$$

then its representation should satisfy the same coarse-graining relation in any description, operational or ontological. More explicitly, this coarse-graining denotes the following post-processing: for any $s \in V_S$, relabel each outcome $\tilde{s} \in V_{\tilde{S}}$ to outcome s with probability $p(s|\tilde{s}) \in \{0, 1\}$; the logical disjunction of those $\tilde{s}$ which are relabelled to s with probability 1 then defines the source event $[s|S]$. Now, in the operational theory, this coarse-graining is represented by

$$\forall [m|M], \text{ where } m \in V_M, M \in \mathcal{M} : \\ p(m, s|M, S) \equiv \sum_{\tilde{s}} p(s|\tilde{s})p(m, \tilde{s}|M, \tilde{S}), \quad (14)$$

and in the ontological model it is represented by

$$\forall \lambda \in \Lambda : \mu(\lambda, s|S) \equiv \sum_{\tilde{s}} p(s|\tilde{s})\mu(\lambda, \tilde{s}|\tilde{S}). \quad (15)$$

In this paper, we will focus on a specific type of coarse-graining: namely, completely coarse-graining over the outcomes of a source setting, say $\{[\tilde{s}|\tilde{S}]\}_{\tilde{s} \in V_{\tilde{S}}}$, to yield an effective one-outcome source-setting, denoted $\tilde{S}_\top$, associated with a single source event $\{[\top|\tilde{S}_\top]\}$, where $[\top|\tilde{S}_\top] \equiv \sum_{\tilde{s}} [\tilde{s}|\tilde{S}]$. In the operational theory, this coarse-graining is represented by

$$\forall [m|M], \text{ where } m \in V_M, M \in \mathcal{M} : \\ p(m, \top|M, \tilde{S}_\top) \equiv \sum_{\tilde{s}} p(m, \tilde{s}|M, \tilde{S}), \quad (16)$$

and in the ontological model it is represented by

$$\forall \lambda \in \Lambda : \mu(\lambda, \top|\tilde{S}_\top) \equiv \sum_{\tilde{s}} \mu(\lambda, \tilde{s}|\tilde{S}). \quad (17)$$

surement obtained by coarse-graining another (parent) measurement as a fundamentally new measurement with response functions *not* respecting the coarse-graining relations with the parent measurement's response functions, even if such coarse-graining relations are respected in the operational description. Such an ontological model, however, will not be able to articulate the ingredients needed for a proof of the KS theorem and we will not consider it here. Indeed, in the absence of the requirement that coarse-graining relations be respected in an ontological model, one can easily construct an ontological model that is "KS-noncontextual" for any operational theory. The interested reader may look at Appendix B for more details, perhaps after looking at Section 2.5 for the relevant definitions of noncontextuality.

Hence, we use the notation $[\top|\tilde{S}_\top]$ to denote the source event that “at least one of the source outcomes in the set $V_{\tilde{S}}$ occurs for source setting $\tilde{S}$ ” (i.e., the logical disjunction of $\tilde{s} \in V_{\tilde{S}}$), formally denoting the choice of $\tilde{S}$ and the subsequent coarse-graining over $\tilde{s}$ by the “source setting” $\tilde{S}_\top$ and the definite outcome of this source setting by “ $\top$ ”. This source event always occurs, i.e., $p(\top|\tilde{S}_\top) = 1$, so $p(m, \top|M, \tilde{S}_\top) = p(m|M, \tilde{S}_\top, \top)$ and $\mu(\lambda, \top|\tilde{S}_\top) = \mu(\lambda|\tilde{S}_\top, \top)$.

This notion of coarse-graining over all the outcomes of a source setting allows us to define a notion of operational equivalence between the source settings themselves. More precisely, two *source settings* S and S' are said to be operationally equivalent, denoted $[\top|S_\top] \simeq [\top|S'_\top]$, if no measurement event can distinguish them once all their outcomes are coarse-grained over, i.e.,

$$\sum_{s \in V_S} p(m, s|M, S) = \sum_{s' \in V_{S'}} p(m, s'|M, S') \quad \forall [m|M], m \in V_M, M \in \mathcal{M}. \quad (18)$$

In quantum theory, this would correspond to the operational equivalence $\sum_s p(s|S)\rho_{[s|S]} = \sum_{s'} p(s'|S')\rho_{[s'|S']}$ for the density operator obtained by completely coarse-graining over two distinct ensembles of quantum states, $\{(p(s|S), \rho_{[s|S]})\}_{s \in V_S}$ and $\{(p(s'|S'), \rho_{[s'|S']})\}_{s' \in V_{S'}}$ on some Hilbert space $\mathcal{H}$.

2.4 Joint measurability (or compatibility)

A given measurement procedure, $\{[m|M]\}_{m \in V_M}$ for some $M \in \mathcal{M}$, in the operational description can be coarse-grained in many different ways to define new effective measurement procedures. The coarse-grained measurement procedures thus obtained from $\{[m|M]\}_{m \in V_M}$ are then said to be jointly measurable (or compatible), i.e., they can be jointly implemented by the same measurement procedure $\{[m|M]\}_{m \in V_M}$ which we refer to as their parent or joint measurement. Formally, a set $\mathcal{C}$ of measurement procedures

$$\{[m_i|M_i]\}_{m_i \in V_{M_i}} | i \in \{1, 2, 3, \dots, |\mathcal{C}|\}$$

is said to be jointly measurable (or compatible) if it arises from coarse-grainings of a single measurement procedure $M \in \mathcal{M}$, i.e., for all $\{[m_i|M_i]\}_{m_i \in V_{M_i}} \in \mathcal{C}$

$$[m_i|M_i] \equiv \sum_{m \in V_M} p(m_i|m)[m|M], \quad (19)$$

where for all i, m, m_i : $p(m_i|m) \in \{0, 1\}$ and $\sum_{m_i \in V_{M_i}} p(m_i|m) = 1$. In terms of the operational probabilities, this means that

$$\begin{aligned} & \forall [s|S], s \in V_S, S \in \mathcal{S} \text{ and } \forall \{[m_i|M_i]\}_{m_i \in V_{M_i}} \in \mathcal{C} : \\ & p(m_i, s|M_i, S) \equiv \sum_{m \in V_M} p(m_i|m)p(m, s|M, S). \end{aligned} \quad (20)$$

If, on the other hand, a set of measurement procedures cannot arise from coarse-grainings of any single

measurement procedure, then the measurement procedures in the set are said to be *incompatible*, i.e., they cannot be jointly implemented.

Note that we will also often refer to a measurement procedure $\{[m_i|M_i]\}_{m_i \in V_{M_i}}$ by just its measurement setting, M_i , and thus speak of the (in)compatibility of measurement settings. Another notion that we will need to refer to is the joint measurability of measurement events: a set of measurement events that arise as outcomes of a single measurement setting are said to be jointly measurable, e.g., all the measurement events in $\{[m|M]\}_{m \in V_M}$ are jointly measurable since they arise as outcomes of a single measurement setting M .

As a quantum example, consider a commuting pair of projective measurements, say $\{\Pi_1, I - \Pi_1\}$ and $\{\Pi_2, I - \Pi_2\}$, where Π_1 and Π_2 are projectors on some Hilbert space $\mathcal{H}$ such that $\Pi_1\Pi_2 = \Pi_2\Pi_1$ and I is the identity operator on $\mathcal{H}$. This pair is jointly implementable since they can be obtained by coarse-graining the outcomes of the joint projective measurement given by $\{\Pi_1\Pi_2, \Pi_1(I - \Pi_2), (I - \Pi_1)\Pi_2, (I - \Pi_1)(I - \Pi_2)\}$.

2.5 Noncontextuality

It is always possible to build an ontological model reproducing the predictions of any operational theory, while respecting the coarse-graining relations.⁹ A trivial example of such an ontological model is one where ontic states λ are identified with the preparation procedures $P_{[s|S]}$ (where $s \in V_S$ and $S \in \mathcal{S}$) and we have $\mu(\lambda, s|S) \equiv \delta_{\lambda, \lambda_{[s|S]}} p(s|S)$, where ontic state $\lambda_{[s|S]}$ is the one deterministically sampled by the preparation procedure $P_{[s|S]}$. Further, the response functions are identified with operational probabilities as $\xi(m|M, \lambda_{[s|S]}) \equiv p(m|M, S, s)$. Then we have $\sum_{\lambda \in \Lambda} \xi(m|M, \lambda) \mu(\lambda, s|S) = \xi(m|M, \lambda_{[s|S]}) p(s|S) = p(m, s|M, S)$. Also, coarse-graining relations of the type $[\tilde{m}|\tilde{M}] \equiv \sum_m p(\tilde{m}|m)[m|M]$ and $[\tilde{s}|\tilde{S}] \equiv \sum_s p(\tilde{s}|s)[s|S]$ that are respected in the operational description are also respected in this ontological description: that is, we have $\forall \lambda \in \Lambda : \xi(\tilde{m}|\tilde{M}, \lambda) \equiv \sum_m p(\tilde{m}|m) \xi(m|M, \lambda)$ and $\forall \lambda \in \Lambda : \mu(\lambda, \tilde{s}|\tilde{S}) \equiv \sum_s p(\tilde{s}|s) \mu(\lambda, s|S)$.

Hence, it is only when additional assumptions are imposed on an ontological model that deciding its existence becomes a nontrivial problem. Such additional assumptions must, of course, play an explanatory role to be worth investigating. The assumption we are interested in is *noncontextuality*, applied to both preparation and measurement procedures. Motivated by the methodological principle of the identity of indiscernibles [18], noncontextuality is an inference from

⁹Note that we will always assume coarse-graining relations are respected in any ontological model. The exception is (some of) the discussion in Section 2.3 and Appendix B where we consider the alternative possibility.

the operational description to the ontological description of an experiment. It posits that the equivalence structure in the operational description is preserved in the ontological description, i.e., the reason one cannot distinguish two operationally equivalent representations of procedures based on their operational statistics is that there is, ontologically, no difference in their representations. We now formally define the notion of noncontextuality in its generalized form due to Spekkens [18].

Mathematically, the assumption of measurement noncontextuality entails that

$$\begin{aligned} [m|M] &\simeq [m'|M'] \\ \Rightarrow \xi(m|M, \lambda) &= \xi(m'|M', \lambda), \forall \lambda \in \Lambda, \end{aligned} \quad (21)$$

while the assumption of preparation noncontextuality entails that

$$\begin{aligned} [s|S] &\simeq [s'|S'] \Rightarrow \mu(\lambda, s|S) = \mu(\lambda, s'|S') \quad \forall \lambda \in \Lambda, \\ [\top|S_\top] &\simeq [\top|S'_\top] \Rightarrow \mu(\lambda|S) = \mu(\lambda|S') \quad \forall \lambda \in \Lambda. \end{aligned} \quad (22)$$

Here we denote $\mu(\lambda|S) \equiv \sum_{s \in V_S} \mu(\lambda, s|S)$, etc., for simplicity of notation, rather than use the notation $\mu(\lambda, \top|S_\top)$, etc., for these coarse-grained probability distributions. Note that since coarse-grainings are respected in any ontological model we consider, we indeed have that $\mu(\lambda, \top|S_\top) \equiv \sum_{s \in V_S} \mu(\lambda, s|S)$.

These are the assumptions of noncontextuality – termed *universal noncontextuality* – that form the basis of our approach to noise-robust noncontextuality inequalities [12–17, 45]. Note that the traditional notion of KS-noncontextuality entails, besides measurement noncontextuality above, the assumption of *outcome-determinism*, i.e., for any measurement event $[m|M]$, $\xi(m|M, \lambda) \in \{0, 1\}$ for all $\lambda \in \Lambda$.

It is important to note that in order for our notion of operational equivalence to be experimentally testable, we need that each of sets $\mathbb{M}$ and $\mathbb{S}$ includes a tomographically complete set of measurements and preparations, respectively. That is, the prepare-and-measure experiment testing contextuality can probe a tomographically complete set of preparations and measurements. Of course, the set of all possible measurements in a theory ($\mathcal{M}$) is (by definition) tomographically complete for any preparation in the theory and, similarly, the set of all possible preparations ($\mathcal{S}$) in a theory is tomographically complete for any measurement in the theory. However, there may exist smaller (finite) sets of preparations and measurements in the theory that are tomographically complete and in that case we require that $\mathbb{S}$ and $\mathbb{M}$ include such tomographically complete sets, even if they don't include all possible preparations and measurements in the theory. For example, when the operational theory is quantum theory for a qubit, the three spin measurements $\{\sigma_x, \sigma_y, \sigma_z\}$ are tomographically complete for any qubit preparation, so we require that $\mathbb{M}$ includes these three measurements even if it doesn't include every other possible measurement on a qubit. While the

requirement that $\mathbb{S}$ and $\mathbb{M}$ include tomographically complete sets doesn't directly reflect in our *theoretical* derivation of the noise-robust noncontextuality inequalities later, it is crucial for *experimentally* verifying the operational equivalences (cf. Eqs. (1),(2),(18)) we need to even invoke the assumption of noncontextuality (cf. Eqs. (21),(22)). Further, this assumption on $\mathbb{M}$ and $\mathbb{S}$ has so far been necessary to be able to implement an actual noise-robust contextuality experiment [13], besides the requirement that the operational theory be convex, i.e., probabilistic mixtures of procedures in the theory (whether preparations or measurements) are also valid procedures in the theory. We refer the reader to Refs. [13, 36, 37] for a discussion of what tomographic completeness entails for (convex) operational theories formalized as general probabilistic theories (GPTs). Although we will not discuss it in this paper, see Ref. [38] for some recent work towards relaxing the tomographic completeness requirement for the set of measurement settings.

2.6 An example of Spekkens contextuality: the fair coin flip inequality

We recap here an example of Spekkens contextuality that has been experimentally demonstrated [13] to give the reader a flavour of the general approach we are going to adopt in the rest of this paper with regard to Kochen-Specker type scenarios. We call the inequality tested in Ref. [13] the “fair coin flip” inequality.

Consider a prepare-and-measure scenario with three source settings, denoted $\mathbb{S} \equiv \{S_1, S_2, S_3\}$, such that $V_{S_i} \equiv \{0, 1\}$ and we have $p(s_i = 0|S_i) = p(s_i = 1|S_i) = 1/2$ for all $i \in \{1, 2, 3\}$. Each S_i thus corresponds to the ensemble of preparation procedures $\{(p(s_i|S_i), P_{[s_i|S_i]})\}_{s_i \in V_{S_i}}$ and we have the following operational equivalence among the source settings after coarse-graining:

$$[\top|S_{1\top}] \simeq [\top|S_{2\top}] \simeq [\top|S_{3\top}]. \quad (23)$$

There are four measurement settings in this scenario, denoted $\mathbb{M} \equiv \{M_1, M_2, M_3, M_{\text{fcf}}\}$, such that $V_{M_i} \in \{0, 1\}$ for all $i \in \{1, 2, 3, \text{fcf}\}$. The measurement setting M_{fcf} is a fair coin flip, i.e., it is insensitive to the preparation procedure preceding it and yields the outcome $m_{\text{fcf}} = 0$ or 1 with equal probability for any preparation procedure $P_{[s|S]}$, i.e., $p(m_{\text{fcf}} = 0|M_{\text{fcf}}, S, s) = p(m_{\text{fcf}} = 1|M_{\text{fcf}}, S, s) = 1/2$ for all $[s|S]$.

We also define a measurement procedure M_{mix} as a classical post-processing of M_1, M_2, M_3 , i.e., its measurement events $\{[m_{\text{mix}}|M_{\text{mix}}]\}_{m_{\text{mix}}=0}^1$ are defined by the classical post-processing relation

$$[m_{\text{mix}}|M_{\text{mix}}] \equiv \sum_{i=1}^3 p(i) \sum_{m_i=0}^1 p(m_{\text{mix}}|m_i)[m_i|M_i], \quad (24)$$

which symbolically denotes the following post-processing of measurements M_1, M_2, M_3 : consider a uniform probability distribution $\{p(i) = \frac{1}{3}\}_{i=1}^3$ over the measurement settings $\{M_i\}_{i=1}^3$ and relabel the respective measurement outcomes, i.e., $\{m_i \in \{0, 1\}\}_{i=1}^3$, to a measurement outcome $m_{\text{mix}} \in \{0, 1\}$ according to the probability distributions

$$\{p(m_{\text{mix}}|m_i) = \delta_{m_{\text{mix}}, m_i}\}_{m_{\text{mix}} \in \{0, 1\}}_{i=1}^3;$$

coarse-graining over m_i and i then yields the effective measurement setting M_{mix} with outcomes labelled by $m_{\text{mix}} \in \{0, 1\}$. In contrast to the kinds of coarse-graining (over measurement outcomes) that appear in KS-noncontextuality (which we discussed in Section 2.3), the (probabilistic) coarse-graining here is over the measurement settings themselves while retaining the outcome labels.¹⁰ We require that this coarse-graining relation be respected in the operational as well as the ontological description. In the operational description, this coarse-graining is represented by

$$\forall [s|S], b \in \{0, 1\} :$$

$$p(m_{\text{mix}} = b, s|M_{\text{mix}}, S) \equiv \frac{1}{3} \sum_{i=1}^3 p(m_i = b, s|M_i, S). \quad (25)$$

We require the following operational equivalence between measurement events of M_{mix} and M_{fct} with respect to which we invoke the assumption of measurement noncontextuality:

$$\forall b \in \{0, 1\} : [m_{\text{mix}} = b|M_{\text{mix}}] \simeq [m_{\text{fct}} = b|M_{\text{fct}}] \quad (26)$$

If we then look at an operational quantity quantifying source-measurement correlations, namely,

$$\text{Corr}_{\text{fct}} \equiv \sum_{i=1}^3 \frac{1}{3} \sum_{m_i, s_i} \delta_{m_i, s_i} p(m_i, s_i|M_i, S_i), \quad (27)$$

then the assumption of preparation noncontextuality applied to operational equivalence in Eq. (23) (so that $\mu(\lambda|S_1) = \mu(\lambda|S_2) = \mu(\lambda|S_3)$ for all $\lambda \in \Lambda$) and the assumption of measurement noncontextuality applied to the operational equivalence in Eq. (26) (so that $\frac{1}{3}\xi(0|M_1, \lambda) + \frac{1}{3}\xi(0|M_2, \lambda) + \frac{1}{3}\xi(0|M_3, \lambda) = \frac{1}{2}$ for all $\lambda \in \Lambda$) lead to the following constraint:

$$\text{Corr}_{\text{fct}} \leq \frac{5}{6}. \quad (28)$$

¹⁰We did not discuss these more general types of classical post-processing in Section 2.3 because they are not relevant to the treatment of Kochen-Specker type scenarios in the Spekkens framework. The example we present here is from Ref. [13], which is not of Kochen-Specker type. The general principle underlying the representation of such classical post-processings is, however, the same: they should be respected in the operational as well as the ontological description.

To see how this is obtained, note that

$$\begin{aligned} & \sum_{i=1}^3 \frac{1}{3} \sum_{m_i, s_i} \delta_{m_i, s_i} p(m_i, s_i|M_i, S_i) \\ &= \sum_{i=1}^3 \frac{1}{3} \sum_{m_i, s_i} \delta_{m_i, s_i} \sum_{\lambda \in \Lambda} \xi(m_i|M_i, \lambda) \mu(\lambda, s_i|S_i) \\ &\leq \sum_{i=1}^3 \frac{1}{3} \sum_{\lambda \in \Lambda} \max_{m_i} \xi(m_i|M_i, \lambda) \sum_{m_i, s_i} \delta_{m_i, s_i} \mu(\lambda, s_i|S_i) \\ &= \sum_{i=1}^3 \frac{1}{3} \sum_{\lambda \in \Lambda} \zeta(M_i, \lambda) \sum_{s_i} \mu(\lambda, s_i|S_i) \\ &= \sum_{\lambda \in \Lambda} \sum_{i=1}^3 \frac{1}{3} \zeta(M_i, \lambda) \nu(\lambda), \end{aligned} \quad (29)$$

where we have that $\zeta(M_i, \lambda) \equiv \max_{m_i} \xi(m_i|M_i, \lambda)$ and that $\nu(\lambda) \equiv \mu(\lambda|S_i)$ for all $i \in \{1, 2, 3\}$. This allows us to put the upper bound

$$\text{Corr}_{\text{fct}} \leq \max_{\lambda \in \Lambda} \frac{1}{3} \sum_{i=1}^3 \zeta(M_i, \lambda), \quad (30)$$

which, subject to the constraint (from measurement noncontextuality) that $\frac{1}{3}\xi(0|M_1, \lambda) + \frac{1}{3}\xi(0|M_2, \lambda) + \frac{1}{3}\xi(0|M_3, \lambda) = \frac{1}{2}$, yields Eq. (28).¹¹ It turns out that in quantum theory the sources and measurements required for this scenario can be realized on a qubit and they can, in principle, achieve the value $\text{Corr} = 1$. This can be achieved by taking the three preparations to be the trine preparations on an equatorial plane (say, the Z-X plane) of the Bloch sphere and the measurements $\{M_i\}_{i=1}^3$ to be the trine measurements, i.e.,

$$\begin{aligned} \rho_{[s_i=0|S_i]} &\equiv \frac{1}{2}(\mathbb{I} + \vec{\sigma} \cdot \vec{n}_i) \equiv \Pi_i^0, \\ \rho_{[s_i=1|S_i]} &\equiv \frac{1}{2}(\mathbb{I} - \vec{\sigma} \cdot \vec{n}_i) \equiv \Pi_i^1, \\ E_{[m_i=0|M_i]} &\equiv \Pi_i^0, \\ E_{[m_i=1|M_i]} &\equiv \Pi_i^1, \end{aligned} \quad (31)$$

where $\vec{n}_1 \equiv (0, 0, 1)$, $\vec{n}_2 \equiv (\frac{\sqrt{3}}{2}, 0, -\frac{1}{2})$, $\vec{n}_3 \equiv (-\frac{\sqrt{3}}{2}, 0, -\frac{1}{2})$, and $\vec{\sigma} \equiv (\sigma_x, \sigma_y, \sigma_z)$ denotes the three Pauli matrices $\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, $\sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$, and $\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$. The operational equivalences are then easy to verify:

$$\begin{aligned} \rho_{[\top|S_{i\top}]} &= \frac{\mathbb{I}}{2}, \quad \forall i \in \{1, 2, 3\}, \\ \frac{1}{3} \sum_{i=1}^3 \Pi_i^0 &= \frac{\mathbb{I}}{2}. \end{aligned} \quad (32)$$

¹¹The reader may look at Appendix B.1 of Ref. [13] to convince themselves that the maximum is achieved for an assignment of response functions of the type $\xi(0|M_1, \lambda) = 1$, $\xi(0|M_2, \lambda) = \frac{1}{2}$ and $\xi(0|M_3, \lambda) = 0$ for some λ .

The quantity $\text{Corr}_{\text{fcf}} = 1$ from this quantum realization. The experimental violation of the noise-robust noncontextuality inequality, Eq. (28), was demonstrated in Ref. [13], where more details may be found. Note that the fair coin flip inequality, Eq.(28), is not inspired by the kinds of operational equivalences that are relevant in a proof of the Kochen-Specker theorem, but employs other kinds of operational equivalences allowed in the Spekkens framework [18], i.e., the operational equivalences in Eqs. (23) and (26) do not arise from the same measurement outcome being shared by different measurements.

Our goal in the present paper is to provide a framework for noise-robust noncontextuality inequalities obtained from statistical proofs of the KS theorem, in particular those that are covered by the CSW framework [22], so that such inequalities can be put to an experimental test along the lines of Ref. [13] within the Spekkens framework. Hence, the operational equivalences between measurement events that will be of interest to us in this paper are precisely those which allow for a proof of the KS theorem, i.e., those which correspond to the same measurement outcome (e.g., a projector) being shared by different measurements (e.g., projective measurements).

2.7 Connection to Bell scenarios

As further motivation to study the questions we are posing, note that one can also view the general prepare-and-measure scenario we are considering in this paper (Fig. 1) as arising on one wing of a two-party Bell experiment: that is, given two parties – Alice and Bob – sharing an entangled state and performing local measurements in a Bell experiment, one can view each choice of measurement setting on Alice’s side as preparing an ensemble of states on Bob’s side; on account of no-signalling, the reduced state on Bob’s side will be the same regardless of Alice’s choice of measurement setting, i.e., all the ensembles corresponding to Alice’s measurement settings (hence, Bob’s source settings) will be operationally equivalent.

For example, consider a Bell experiment where Alice has two choices of measurement settings, $M_x^A \equiv \sigma_x$ or $M_z^A \equiv \sigma_z$, and she shares a Bell state with Bob: $|\psi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. Bob has access to some set of measurement settings $\mathbb{M}^B \equiv \{M_j^B\}_j$ on his system. When Alice measures M_x^A , she prepares the ensemble of states $S_x^A \equiv \{(1/2, \rho_{[s_x^A=0|S_x^A]} \equiv |+\rangle\langle+|), (1/2, \rho_{[s_x^A=1|S_x^A]} \equiv |-\rangle\langle-|)\}$ on Bob’s side and when she measures M_z^A she prepares the ensemble of states $S_z^A \equiv \{(1/2, \rho_{[s_z^A=0|S_z^A]} \equiv |0\rangle\langle 0|), (1/2, \rho_{[s_z^A=1|S_z^A]} \equiv |1\rangle\langle 1|)\}$. These ensembles are operationally equivalent, yielding the maximally mixed state on coarse-graining, i.e.,

$$\frac{1}{2}|0\rangle\langle 0| + \frac{1}{2}|1\rangle\langle 1| = \frac{1}{2}|+\rangle\langle+| + \frac{1}{2}|-\rangle\langle-| = \frac{\mathbb{I}}{2}. \quad (33)$$

The quantity of interest in a Bell experiment $p(m_i^A, m_j^B | M_i^A, M_j^B)$ ($i \in \{x, z\}$) is then formally the same as the quantity $p(s_i^A, m_j^B | S_i^A, M_j^B)$ that we are interested in our prepare-and-measure scenario. In the ontological model describing the effective prepare-and-measure experiment on Bob’s system, we have the following:

$$\begin{aligned} p(s_i^A, m_j^B | S_i^A, M_j^B) &= \sum_{\lambda} \Pr(m_j^B | M_j^B, \lambda) \Pr(\lambda, s_i^A | S_i^A) \\ &= \sum_{\lambda} \Pr(m_j^B | M_j^B, \lambda) \Pr(s_i^A | S_i^A, \lambda) \Pr(\lambda | S_i^A). \end{aligned} \quad (34)$$

Assuming preparation noncontextuality relative to the operational equivalence $[\top | S_x^A] \simeq [\top | S_z^A]$, we have $\Pr(\lambda | S_x^A) = \Pr(\lambda | S_z^A) \equiv \Pr(\lambda)$, so that

$$\begin{aligned} p(s_i^A, m_j^B | S_i^A, M_j^B) &= \sum_{\lambda} \Pr(s_i^A | S_i^A, \lambda) \Pr(m_j^B | M_j^B, \lambda) \Pr(\lambda), \end{aligned} \quad (35)$$

which formally resembles the expression for local causality when applied to the corresponding two-party Bell experiment:

$$\begin{aligned} p(m_i^A, m_j^B | M_i^A, M_j^B) &= \sum_{\lambda} \Pr(m_i^A | M_i^A, \lambda) \Pr(m_j^B | M_j^B, \lambda) \Pr(\lambda). \end{aligned} \quad (36)$$

If no other assumption of noncontextuality is invoked besides the one applied to the operational equivalence of source settings on Bob’s system, then the constraints on $p(s_i^A, m_j^B | S_i^A, M_j^B)$ will be the same as the constraints on $p(m_i^A, m_j^B | M_i^A, M_j^B)$ from Bell inequalities.

Note, however, that the response functions $\Pr(m_j^B | M_j^B, \lambda)$ and $\Pr(m_j^A | M_j^A, \lambda)$ can be completely arbitrary in a locally causal ontological model for the Bell experiment and the same applies to the distributions $\Pr(s_i^A | S_i^A, \lambda)$ and $\Pr(m_j^B | M_j^B, \lambda)$ in a preparation noncontextual model of the corresponding prepare-and-measure scenario on Bob’s side. We will be interested in imposing *additional* constraints on the response functions $\Pr(m_j^B | M_j^B, \lambda)$ of the prepare-and-measure scenario (on Bob’s side) that follow from the assumption of measurement noncontextuality applied to operational equivalences between measurement events on Bob’s side. In particular, we are interested in those operational equivalences between measurement events that are required by any statistical proof of the Kochen-Specker theorem [16, 47]. We develop this approach more carefully in the following sections.

3 Hypergraph approach to Kochen-Specker scenarios in the Spekkens framework

Having set up the framework needed to articulate the relevant notions in Section 2, we now proceed to consider Kochen-Specker type experimental scenarios in this framework. To do this, we will use the language of hypergraphs and their subgraphs to represent the operational equivalences between measurement events that are required in a Kochen-Specker argument as well as the operational equivalences between source settings that we will invoke in our generalization. The (hyper)graph-theoretic ingredients of our approach will represent those aspects of the general framework of Section 2 that are necessary to go from the CSW framework for KS-contextuality to a hypergraph framework for Spekkens contextuality applied to Kochen-Specker type experimental scenarios.

Our presentation will be a hybrid one, discussing features of the CSW framework [21, 22] in the notation of the AFLS framework [23], but extending both in ways appropriate for the purpose of this paper. Our goal is to demonstrate how the graph-theoretic invariants of CSW [22] can be repurposed towards obtaining noise-robust noncontextuality inequalities.

We do this in two parts: first, we define a representation of measurement events in the manner of Refs. [22, 23], and then we define a representation of source events in the spirit of Ref. [12].

3.1 Measurements

The basic object for representing measurements is a hypergraph, Γ , with a finite set of vertices $V(\Gamma)$ such that each vertex $v \in V(\Gamma)$ denotes a measurement outcome, and a set of hyperedges $E(\Gamma)$ such that each hyperedge $e \in E(\Gamma)$ is a subset of $V(\Gamma)$ and denotes a measurement consisting of outcomes in e . Here, $E(\Gamma) \subseteq 2^{V(\Gamma)}$ and $\bigcup_{e \in E(\Gamma)} e = V(\Gamma)$. Such a hypergraph satisfies the definition of a *contextuality scenario* à la AFLS [23]. We will further assume, unless specified otherwise, that the hypergraph is *simple*: that is, for all $e_1, e_2 \in E(\Gamma)$, $e_1 \subseteq e_2 \Rightarrow e_1 = e_2$, or that no hyperedge is a strict subset of another. Such hypergraphs are also called *Sperner families* [46]. Two measurement events are said to be *(mutually) exclusive* if the vertices denoting them appear in a common hyperedge, i.e., if they can be realized as outcomes of a single measurement setting.

The structure of a *contextuality scenario* Γ represents the operational equivalences between measurement events that are of interest in a Kochen-Specker argument. We emphasize here that we take the operational theory to be fundamental and the contextuality scenario for a particular Kochen-Specker argument to be derived from (and as a graphical representation of) the operational equivalences in the operational theory

(cf. Section 2). In particular, depending on the operational equivalences that an operational theory can exhibit (by virtue of (in)compatibility relations between measurements), it may or may not allow some contextuality scenario to be realized by measurement events in the theory. The fact that a given vertex, say $v \in V(\Gamma)$, appears in multiple hyperedges, say $E' \equiv \{e \in E(\Gamma) | v \in e\}$, means that the measurement events corresponding to this vertex, i.e., $\{[v|e]\}_{e \in E'}$, are operationally equivalent, and the equivalence class of these measurement events is denoted by the vertex v itself. In the case of quantum theory, for example, v can represent a positive operator that appears in different positive operator-valued measures (POVMs) represented by the hyperedges.

A *probabilistic model* on Γ is an assignment of probabilities to the vertices $v \in V(\Gamma)$ such that $p(v) \geq 0$ for all $v \in V(\Gamma)$ and $\sum_{v \in e} p(v) = 1$ for all $e \in E(\Gamma)$. As we have noted, every vertex v represents an *equivalence class* of measurement events, denoted $[m|\mathcal{M}]$, and every hyperedge e represents an *equivalence class* of measurement procedures, denoted $\mathcal{M}$.¹² The fact that each v represents an equivalence class of measurement events means that

1. any probabilistic model p on Γ , realized by *operational* probabilities for a given source event – that is, where for all $v \in V(\Gamma)$ and a given $[s|S]$, $p(v) \equiv p(v|S, s) \equiv p(m|\mathcal{M}, S, s)$ – is consistent with the operational equivalences represented by Γ , and
2. any probabilistic model on Γ , realized by *ontological* probabilities for a given ontic state – that is, where for all $v \in V(\Gamma)$ and a given ontic state λ , $p(v) \equiv p(v|\lambda) \equiv \xi(m|\mathcal{M}, \lambda)$ – respects (by definition) the assumption of measurement noncontextuality with respect to the presumed operational equivalences between measurement events.

We will therefore often write $p(m, s|\mathcal{M}, S)$ as $p(v, s|S)$ and $p(m|\mathcal{M}, S, s)$ as $p(v|S, s)$, where $[s|S]$ is a source event. Similarly, we will also write $\xi(m|\mathcal{M}, \lambda)$ as $p(v|\lambda)$, where λ is an ontic state.

Orthogonality graph of Γ , $O(\Gamma)$: Given the hypergraph Γ , we construct its orthogonality graph $O(\Gamma)$: that is, the vertices of $O(\Gamma)$ are given by $V(O(\Gamma)) \equiv V(\Gamma)$, and the edges of $O(\Gamma)$ are given by $E(O(\Gamma)) \equiv \{\{v, v'\} | v, v' \in e \text{ for some } e \in E(\Gamma)\}$.

¹²Note that two measurement procedures with measurement settings M and M' are operationally equivalent if every measurement event of one is operationally equivalent to a distinct measurement event of the other. That is, there is a bijective correspondence (of operational equivalence) between the two sets of measurement events. In quantum theory, for example, a given POVM (which is what a hyperedge would represent), say $\{E_k\}_k$, can be implemented in many possible ways, each such measurement procedure corresponding to different quantum instrument. Mathematically, these different procedures can be represented by different sets of operators $\{O_k\}_k$ such that $E_k = O_k^\dagger O_k$ for all k and $\sum_k O_k^\dagger O_k = \mathbb{I}$.

Figure 2: The KCBS scenario with 4-outcome joint measurements, visualized as a hypergraph Γ [16, 22, 47].

Each edge of $O(\Gamma)$ denotes the exclusivity of the two measurement events it connects, i.e., the fact that they can occur as outcomes of a single measurement.

For any Bell-KS inequality constraining correlations between measurement events from $O(\Gamma)$ (when all measurements are implemented on a given source event), we construct a subgraph G of $O(\Gamma)$ such that the vertices of G , i.e., $V(G)$, correspond to measurement events that appear in the inequality with nonzero coefficients, and two vertices share an edge in G if and only if they share an edge in $O(\Gamma)$. More explicitly, consider a Bell-KS expression

$$R([s|S]) \equiv \sum_{v \in V(G)} w_v p(v|S, s), \quad (37)$$

where $w_v > 0$ for all $v \in V(G)$. A Bell-KS inequality imposes a constraint of the form $R([s|S]) \leq R_{\text{KS}}$, where R_{KS} is the upper bound on the expression in any operational theory that admits a KS-noncontextual ontological model. Often, but not always, these inequalities are simply of the form where $w_v = 1$ for all $v \in V(G)$. In keeping with the CSW notation [22], we will denote the general situation by a *weighted* graph (G, w) , where w is a function that maps vertices $v \in V(G)$ to weights $w_v > 0$. See Figures 2 and 3 for an example from the Klyachko-Can-Binicioğlu-Shumovsky (KCBS) scenario [22, 47].

Below, we make some remarks clarifying the scope of the framework described above before we move to the case of sources.

3.1.1 Classification of probabilistic models

We classify the probabilistic models on a hypergraph Γ as follows:

- KS-noncontextual probabilistic models, $\mathcal{C}(\Gamma)$: a probabilistic model which is a convex combination of deterministic assignments $p : V(\Gamma) \rightarrow$

Figure 3: A subgraph of KCBS hypergraph Γ , representing orthogonality relations of the events of interest in the KCBS inequality [22, 47].

$\{0, 1\}$, where $\sum_{v \in e} p(v) = 1$ for all $e \in E(\Gamma)$. In Ref. [23], this is referred to as a “classical model”.¹³

Note that we call Γ *KS-colourable* if $\mathcal{C}(\Gamma) \neq \emptyset$ and we call it *KS-uncolourable* if $\mathcal{C}(\Gamma) = \emptyset$. Our terminology here is inspired by the traditional usage of the term “Kochen-Specker colouring” to refer to an assignment of two colours to vectors satisfying some orthogonality relations under the colouring constraints of the KS theorem [48].

- Consistent exclusivity satisfying probabilistic models, $\mathcal{CE}^1(\Gamma)$: a probabilistic model on Γ , $p : V(\Gamma) \rightarrow [0, 1]$, such that (in addition to satisfying the definition of a *probabilistic model*), $\sum_{v \in c} p(v) \leq 1$ for all cliques c in the orthogonality graph $O(\Gamma)$. This is the same as the set of *E1* probabilistic models of Ref. [22].

Note that a *clique* in the orthogonality graph $O(\Gamma)$ is a set of vertices that are pairwise exclusive (i.e., every vertex in this set shares an edge with every other vertex).

- General probabilistic models, $\mathcal{G}(\Gamma)$: Any p that satisfies the definition of a probabilistic model is a general probabilistic model, i.e., it can arise from measurements in some general probabilistic theory [1] that isn’t necessarily quantum.

The set of all probabilistic models $\mathcal{G}(\Gamma)$ (for any Γ) forms a polytope since it is defined by just the positivity and normalization constraints on the probabilities. The extremal points (or vertices) of this polytope fall into two categories that will interest us: deterministic and indeterministic. The deterministic extremal points are

¹³We use a different term because we are advocating a revision of the notion of classicality from KS-noncontextuality to generalized noncontextuality à la Spekkens.

the $p : V(\Gamma) \rightarrow \{0, 1\}$ such that $\sum_{v \in e} p(v) = 1$ for all $e \in E(\Gamma)$ and we denote the set of these points by $\mathcal{G}(\Gamma)|_{\text{det}}$. The indeterministic extremal points are the $p \in \mathcal{G}(\Gamma)$ which are not deterministic and which, furthermore, cannot be expressed as a convex mixture of other points in $\mathcal{G}(\Gamma)$. We denote the set of indeterministic extremal points by $\mathcal{G}(\Gamma)|_{\text{ind}}$. Clearly, $\mathcal{G}(\Gamma)|_{\text{det}} \subsetneq \mathcal{C}(\Gamma)$ and $\mathcal{G}(\Gamma)|_{\text{ind}} \subseteq \mathcal{G}(\Gamma) \setminus \mathcal{C}(\Gamma)$.

Overall, we have

$$\mathcal{C}(\Gamma) \subseteq \mathcal{CE}^1(\Gamma) \subseteq \mathcal{G}(\Gamma) \quad (38)$$

for any hypergraph Γ .

3.1.2 Distinguishing two consequences of Specker's principle: Structural Specker's principle vs. Statistical Specker's principle

The CSW framework [22] restricts the scope of probabilistic models on a hypergraph to those satisfying consistent exclusivity (the $E1$ probabilistic models), motivated by what is sometimes called *Specker's principle* [35]: that is,

“if you have several questions and you can answer any two of them, then you can also answer all of them”

If by “questions” we understand *measurement settings*, then the principle says that a set of pairwise jointly implementable measurement settings is itself jointly implementable. Note that when we say a set of *measurement settings* is “jointly implementable”, “jointly measurable”, or “compatible”, we mean that there exists another choice of a single measurement setting in the theory such that this measurement setting can reproduce the statistics of all the measurement settings in the set by coarse-graining.¹⁴ As such, in its application to measurement settings, Specker's principle is a constraint on the measurements allowed in a physical theory that respects it, e.g., measurement settings that correspond to PVMs (projection valued measures) in quantum theory. This is, for example, the reading adopted in Ref. [49], where the failure of Specker's principle in any almost quantum theory was demonstrated. On the other hand, we will often also refer to the “joint measurability” of a set of *measurement events*, by which we mean that this set of measurement events is a subset of the set of measurement outcomes for some choice of measurement setting. At the level of *measurement events*,¹⁵ then, there are two distinct ways to read Specker's principle

¹⁴The reader may recall from Section 2.4 the general definition of compatibility. Also, see Ref. [44] for an overview of joint measurability in quantum theory.

¹⁵Recall that a *measurement event* is a measurement outcome given a choice of measurement setting, e.g., a projector that appears in a particular PVM in quantum theory.

that one needs to keep in mind which we distinguish as *structural Specker's principle* vs. *statistical Specker's principle*. We define these two readings below:

- *Structural Specker's principle* imposes a *structural* constraint on a contextuality scenario Γ .

This (strong) reading of Specker's principle applies to any set of measurement events, say $\mathfrak{M} \subseteq V(\Gamma)$, where every pair of measurement events can arise as outcomes of a single measurement: that is, for each pair $\{v, v'\} \subseteq \mathfrak{M}$, there exists some $e \in E(\Gamma)$ such that $\{v, v'\} \subseteq e$. The principle then states:

Given a set $\mathfrak{M}$ of pairwise jointly measurable measurement events in some contextuality scenario Γ , all the measurement events in $\mathfrak{M}$ are jointly measurable, i.e., all the measurement events in the set can arise as outcomes of a single measurement: $\mathfrak{M} \subseteq e$ for some $e \in E(\Gamma)$.

Alternatively, the constraint of structural Specker's principle can be restated as:

Every clique in the orthogonality graph of Γ , $O(\Gamma)$, is a subset of some hyperedge in Γ .

Note that we haven't said anything directly about *probabilities* here: any Γ satisfying the above property is said to satisfy structural Specker's principle.

- *Statistical Specker's principle* (or *consistent exclusivity*) imposes a *statistical* constraint on probabilistic models on any contextuality scenario Γ representing measurement events in an operational theory.

This (weak) reading of Specker's principle imposes an additional constraint on a probabilistic model $p \in \mathcal{G}(\Gamma)$ (thus defining $\mathcal{CE}^1(\Gamma) \subseteq \mathcal{G}(\Gamma)$), namely:

Given a set $\mathfrak{M}$ of pairwise jointly measurable measurement events, p satisfies $\sum_{v \in \mathfrak{M}} p(v) \leq 1$.

This can also be expressed as:

A probabilistic model $p \in \mathcal{G}(\Gamma)$ is said to satisfy statistical Specker's principle if the sum of probabilities it assigns to the vertices of every clique in the orthogonality graph of Γ , $O(\Gamma)$, does not exceed 1, i.e., $\sum_{v \in c} p(v) \leq 1$ for all cliques c in $O(\Gamma)$.

All probabilistic models that satisfy this constraint define the set of probabilistic models $\mathcal{CE}^1(\Gamma)$ (or $E1$) for any contextuality scenario Γ regardless of whether Γ satisfies structural Specker's principle. Clearly, $\mathcal{CE}^1(\Gamma) \subseteq \mathcal{G}(\Gamma)$.

Any probabilistic model p on Γ such that $p \in \mathcal{CE}^1(\Gamma)$ is said to satisfy statistical Specker's principle or, equivalently, *consistent exclusivity* [23].

Probabilistic models on any hypergraph Γ which satisfies the (strong) structural Specker's principle obviously satisfy the (weak) statistical Specker's principle. This holds simply on account of the structure of such Γ : that is, for all Γ satisfying structural Specker's principle, we have $\mathcal{CE}^1(\Gamma) = \mathcal{G}(\Gamma)$. To see this, note that every clique c in $O(\Gamma)$ is a subset of some hyperedge in Γ , hence for every clique c , $\sum_{v \in c} p(v) \leq 1$ for all $p \in \mathcal{G}(\Gamma)$, i.e., $p \in \mathcal{CE}^1(\Gamma)$.¹⁶ On the other hand, it remains an *open question* whether the converse is true:

That is, given that $\mathcal{CE}^1(\Gamma) = \mathcal{G}(\Gamma)$ for some Γ , is it the case that Γ must then necessarily satisfy structural Specker's principle, namely, that every clique in $O(\Gamma)$ is a subset of some hyperedge in Γ ?

A positive answer to this question would answer Problem 7.2.3 of Ref. [23] asking for a characterization of Γ for which $\mathcal{CE}^1(\Gamma) = \mathcal{G}(\Gamma)$.

3.1.3 What does it mean for an operational theory to satisfy structural/statistical Specker's principle?

We have so far defined structural Specker's principle as a constraint on Γ and statistical Specker's principle as a constraint on a probabilistic model on any Γ . Any operational theory would typically allow many possible Γ to be realized by its measurement events as well as many possible probabilistic models to be realized on any Γ representing its measurement events. Note that when we say that a particular Γ is “realizable” or “allowed” by an operational theory, we mean that there exist measurement events in the operational theory that satisfy the operational equivalences required by Γ .¹⁷ Further, given such a Γ , the realizability of a probabilistic model on it by the operational theory means that there exists a source event in the operational theory that assigns probabilities to the measurement events in Γ according to the probabilistic model. It will be useful for our discussion to define what it means for an operational theory, say $\mathbb{T}$, to satisfy structural or statistical Specker's principle. But before we do that, let us formally specify what it means for $\mathbb{T}$ to satisfy Specker's principle:

$\mathbb{T}$ satisfies Specker's principle: *An operational theory $\mathbb{T}$ is said to satisfy Specker's principle if, for any set of measurement settings in $\mathbb{T}$ that are pairwise jointly implementable, it follows that they are all jointly implementable in $\mathbb{T}$.*¹⁸

¹⁶This partially answers the open Problem 7.2.3 of Ref. [23].

¹⁷Realizability of a particular Γ in an operational theory depends on the (in)compatibility relations that the operational theory allows between its measurements (cf. Section 2.4). Recall that incompatibility of measurements is *necessary* for KS-contextuality to be witnessed and the structure of Γ depends on this incompatibility.

¹⁸Recall from Section 2.4 the definition of joint implementability (or joint measurability) of some set of measurement settings.

We denote by $\mathcal{T}(\Gamma)$ the set of probabilistic models achievable on Γ by an operational theory $\mathbb{T}$, i.e., for any $p \in \mathcal{T}(\Gamma)$, we have that $\forall v \in V(\Gamma) : p(v) = p(v|S, s)$ for some source event $[s|S]$ possible in the operational theory $\mathbb{T}$.¹⁹ Since an operational theory can only put further constraints on probabilistic models in $\mathcal{G}(\Gamma)$, we obviously have: $\mathcal{T}(\Gamma) \subseteq \mathcal{G}(\Gamma)$.

1. $\mathbb{T}$ satisfies statistical Specker's principle:

*We say an operational theory $\mathbb{T}$ satisfies statistical Specker's principle if $\mathcal{T}(\Gamma) \subseteq \mathcal{CE}^1(\Gamma) \subseteq \mathcal{G}(\Gamma)$ for all Γ .*²⁰

Since the satisfaction of statistical Specker's principle is a constraint on the statistical predictions of $\mathbb{T}$, there must be some *fact* about the structure of theory $\mathbb{T}$ that leads to this constraint. This fact enforcing statistical Specker's principle could be some restriction arising from the structure of allowed measurement events and/or even the structure of allowed preparations in the operational theory $\mathbb{T}$. For instance, this is the case for quantum theory when one only considers projective measurements implemented on an arbitrary quantum state, i.e., $\mathcal{Q}(\Gamma) \subseteq \mathcal{CE}^1(\Gamma) \subseteq \mathcal{G}(\Gamma)$, where $\mathcal{Q}(\Gamma)$ denotes the set of probabilistic models that can be obtained in this way. More generally, one could relax the no-restriction hypothesis [3] in some particular way in $\mathbb{T}$ so that not all probabilistic models in $\mathcal{G}(\Gamma)$ are allowed in $\mathcal{T}(\Gamma)$. In the case of quantum theory, restricting attention to only projective measurements (as we just pointed out) rather than the more general case allowing arbitrary POVMs is one way of restricting the set of possible probabilistic models realizable with quantum states and measurements to a strict subset of $\mathcal{G}(\Gamma)$. Allowing arbitrary POVMs would lead to a violation of statistical Specker's principle by probabilistic models arising from quantum theory.²¹

Let us now define what it means for an operational theory $\mathbb{T}$ to satisfy structural Specker's principle.

2. $\mathbb{T}$ satisfies structural Specker's principle:

An operational theory $\mathbb{T}$ is said to satisfy structural Specker's principle if for any set of measurement events that are pairwise jointly measurable, i.e., measurement events in each pair arise

¹⁹Note that if the operational theory does not admit measurement events (represented by vertices) exhibiting the operational equivalences represented by Γ (that is, $\mathbb{T}$ does not allow Γ), then we have that $\mathcal{T}(\Gamma)$ is an empty set.

²⁰That is, instead of considering only a particular probabilistic model on a particular Γ , we now consider the satisfaction of statistical Specker's principle by a whole set of probabilistic models, namely, $\mathcal{T}(\Gamma)$, for all Γ .

²¹See Appendices A (specifically A.1.2) and C for other consequences of allowing arbitrary POVMs, in particular the trivial ‘classical’ ones.

as outcomes of some measurement in the theory, it is the case that all the measurement events in the set are jointly measurable, i.e., all the measurement events in the set arise as outcomes of a single measurement in the theory.

We now show that a theory $\mathbb{T}$ that satisfies Specker's principle also satisfies structural Specker's principle.

Theorem 1. *If an operational theory $\mathbb{T}$ satisfies Specker's principle, then it also satisfies structural Specker's principle.*

Proof. The argument here relies on the fact that the operational theory $\mathbb{T}$ is such that measurement settings can be coarse-grained to yield new measurement settings with fewer outcomes. Operationally, this just corresponds to binning some subsets of outcomes together in a measurement procedure. The operational theories we consider in this paper satisfy this property, as outlined in Section 2.3 on coarse-graining.

The argument proceeds, for any Γ realizable in $\mathbb{T}$, by constructing a set of binary-outcome measurement settings for any given set of pairwise jointly measurable vertices in Γ . These measurement settings are, by construction, pairwise jointly measurable, so Specker's principle applied to them implies that they are all jointly measurable. This in turn means that the pairwise jointly measurable vertices in the given set are also all realizable as outcomes of a single measurement setting. Hence, the theory $\mathbb{T}$ satisfies structural Specker's principle. We detail the argument below.

Consider a contextuality scenario Γ realizable in $\mathbb{T}$. To each vertex $v \in V(\Gamma)$, we can associate a measurement setting M_v with two possible outcomes labelled $\{0, 1\}$ such that $[1|M_v]$ denotes the occurrence of v and $[0|M_v]$ denotes the non-occurrence of v , i.e., $p(v|S, s) = p(1|M_v, S, s)$ and $1 - p(v|S, s) = p(0|M_v, S, s)$ for any probabilistic model on Γ induced by some source event $[s|S]$. The measurement setting M_v can be obtained in various (operationally equivalent) ways from the hyperedges that $v \in V(\Gamma)$ appears in: for each hyperedge $e \in E(\Gamma)$ such that $v \in e$, we have that the binary-outcome measurement setting consisting of the vertices $\{v, e \setminus v\}$ — where $e \setminus v$ denotes a coarse-graining over all the measurement outcomes of e except v — is operationally equivalent to M_v .

Now, for any pair of vertices $\{v, v'\}$ that appear in a common hyperedge of Γ , consider the two corresponding measurement settings $\{M_v, M_{v'}\}$ such that they are jointly measurable and their outcomes are mutually exclusive. The measurement events that can possibly occur in their joint measurement, denoted $M_{vv'}$, are $[10|M_{vv'}]$, $[01|M_{vv'}]$ and $[00|M_{vv'}]$. The probability of $[11|M_{vv'}]$ is always zero, reflecting the fact that v and v' are mutually exclusive. Here, the coarse-graining relations are: $[1|M_v] \equiv [10|M_{vv'}]$, $[1|M_{v'}] \equiv$

$$[01|M_{vv'}], [0|M_v] \equiv [00|M_{vv'}] + [01|M_{vv'}], [0|M_{v'}] \equiv [00|M_{vv'}] + [10|M_{vv'}].$$

The joint measurement $M_{vv'}$ can be constructed from any hyperedge that v and v' appear in: for any $e \in E(\Gamma)$ such that $\{v, v'\} \subseteq e$, we have that $[10|M_{vv'}]$ is a measurement event corresponding to v ,²² $[01|M_{vv'}]$ corresponds to v' , $[00|M_{vv'}]$ corresponds to $e \setminus \{v, v'\}$ (the coarse-graining of all measurement outcomes in e except v and v'), and $[11|M_{vv'}]$ denotes the null event $\emptyset \subseteq e$. This means $p(10|M_{vv'}, S, s) + p(01|M_{vv'}, S, s) + p(00|M_{vv'}, S, s) = p(v|S, s) + p(v'|S, s) + p(e \setminus \{v, v'\}|S, s) = 1$ and $p(11|M_{vv'}, S, s) = 0$ for any probabilistic model (induced by some source event $[s|S]$) on Γ .

Consider now any set of vertices in Γ that is pairwise jointly measurable, denoted $V_{2\text{JM}} \subseteq V(\Gamma)$. We need to show that any such set of vertices $V_{2\text{JM}}$ is jointly measurable, i.e., the theory $\mathbb{T}$ realizing Γ admits a single measurement such that all the vertices in $V_{2\text{JM}}$ arise as outcomes of this measurement.

Now, the two-outcome measurement settings $\{M_v|v \in V_{2\text{JM}}\}$ we have defined are pairwise jointly measurable and as such, following Specker's principle, they should all be jointly measurable in theory $\mathbb{T}$. The joint measurement corresponding to them can be defined as

$$M_{V_{2\text{JM}}} \equiv \{[\vec{b}|M_{V_{2\text{JM}}}]|\vec{b} \in \{0, 1\}^{V_{2\text{JM}}}\}, \quad (39)$$

where each event $[\vec{b}|M_{V_{2\text{JM}}}]$ in the joint measurement $M_{V_{2\text{JM}}}$ represents a particular set of outcomes for measurements in the set $\{M_v|v \in V_{2\text{JM}}\}$.

Denoting $V_{2\text{JM}} \equiv \{v_1, v_2, \dots, v_{|V_{2\text{JM}}|}\}$, we have that

$$\begin{aligned} [(10 \dots 0)|M_{V_{2\text{JM}}}] &\equiv [1|M_{v_1}], \\ [(01 \dots 0)|M_{V_{2\text{JM}}}] &\equiv [1|M_{v_2}], \\ &\vdots \\ [(00 \dots 1)|M_{V_{2\text{JM}}}] &\equiv [1|M_{v_{|V_{2\text{JM}}|}}], \\ [(00 \dots 0)|M_{V_{2\text{JM}}}] &\equiv [0|M_{v_1}] + [0|M_{v_2}] + \dots + [0|M_{v_{|V_{2\text{JM}}|}}], \end{aligned} \quad (40)$$

where $[0|M_{v_1}] + [0|M_{v_2}] + \dots + [0|M_{v_{|V_{2\text{JM}}|}}]$ denotes the measurement event obtained by coarse-graining the measurement events in $\{[0|M_v]|v \in V_{2\text{JM}}\}$. All the other measurement events of $M_{V_{2\text{JM}}}$ are null events that never occur, i.e., they are assigned probability zero by every source event. Thus, using Specker's principle applied to the binary-outcome measurement settings defined for the vertices in $V_{2\text{JM}}$, we have that the pairwise jointly measurable vertices in $V_{2\text{JM}}$ are all jointly measurable, appearing as outcomes of a single measurement $M_{V_{2\text{JM}}}$. $\square$

²²Recall that every vertex $v \in V(\Gamma)$ is an *equivalence class* of measurement events $[v|e] \simeq [v|e']$ for all e, e' such that $v \in e$ and $v \in e'$.

Having established Theorem 1, we now proceed to show that a theory which satisfies structural Specker's principle also satisfies statistical Specker's principle. To do this, we consider a contextuality scenario Γ which may not satisfy structural Specker's principle and from it construct a contextuality scenario Γ' which does satisfy the principle. The construction proceeds as follows:

1. Construct $O(\Gamma)$.
2. Turn each clique in $O(\Gamma)$ that *is* a hyperedge in Γ to a hyperedge in a new hypergraph Γ' . That is, Γ' is such that $V(\Gamma) \subseteq V(\Gamma')$ and $E(\Gamma) \subseteq E(\Gamma')$.
3. Turn each maximal clique c in $O(\Gamma)$ that is *not* a hyperedge in Γ to a hyperedge in Γ' and include an additional vertex v_c in this hyperedge. Here, a *maximal clique* in a graph is a clique that is not a strict subset of another clique, i.e., there is no vertex outside the clique that shares an edge with each vertex in the clique.

We then have for the hyperedges of Γ' ,

$$E(\Gamma') = E(\Gamma) \cup \{c \cup \{v_c\}\}_{c \in C}, \quad (41)$$

where C is the set of maximal cliques in $O(\Gamma)$ that are not hyperedges in Γ .

Note that as long as a theory $\mathbb{T}$ satisfies structural Specker's principle, converting maximal cliques in $O(\Gamma)$ that are not hyperedges in Γ to hyperedges in Γ' is a valid move within the theory since the resulting hyperedge would indeed constitute a valid measurement in the theory.

If $C = \emptyset$ (i.e., Γ satisfies structural Specker's principle), then we just have $E(\Gamma') = E(\Gamma)$.

4. The resulting contextuality scenario Γ' is thus given by: $V(\Gamma') = V(\Gamma) \cup \{v_c\}_{c \in C}$ and $E(\Gamma') = E(\Gamma) \cup \{c \cup \{v_c\}\}_{c \in C}$.

If $C = \emptyset$ we just have $V(\Gamma') = V(\Gamma)$ and $E(\Gamma') = E(\Gamma)$ so that $\Gamma' = \Gamma$ (i.e., the two hypergraphs are isomorphic).

Our construction of Γ' leads to the following properties:

- Γ' satisfies structural Specker's principle (by construction) since every clique in $O(\Gamma')$ is a subset of some hyperedge in Γ' . Hence, it's also the case that statistical Specker's principle holds for probabilistic models on Γ' as $\mathcal{CE}^1(\Gamma') = \mathcal{G}(\Gamma')$.

Note that the construction of Γ' relied on the fact that the theory we are considering satisfies structural Specker's principle. If the theory doesn't satisfy this principle, but one goes ahead with the construction of Γ' , then the new hyperedges in Γ' may not constitute valid measurements in the theory.

- Probabilistic models in $\mathcal{G}(\Gamma')$ are in bijective correspondence with probabilistic models in $\mathcal{CE}^1(\Gamma)$: for any probabilistic model $p_\Gamma \in \mathcal{CE}^1(\Gamma)$, there exists a unique probabilistic model $p_{\Gamma'} \equiv f(p_\Gamma) \in \mathcal{G}(\Gamma')$, where the function f is given by $p_{\Gamma'}(v) \equiv f(p_\Gamma)(v) = p_\Gamma(v)$ for all $v \in V(\Gamma)$ and $p_{\Gamma'}(v_c) \equiv f(p_\Gamma)(v_c) = 1 - \sum_{v \in c} p_\Gamma(v)$ for all $c \in C$.²³ Similarly, for any $p_{\Gamma'} \in \mathcal{G}(\Gamma')$, there exists a unique probabilistic model $p_\Gamma \equiv g(p_{\Gamma'}) \in \mathcal{CE}^1(\Gamma)$ given by $p_\Gamma(v) \equiv g(p_{\Gamma'})(v) = p_{\Gamma'}(v)$ for all $v \in V(\Gamma)$, i.e., we simply ignore the probabilities assigned to the vertices $v_c \in V(\Gamma') \setminus V(\Gamma)$ which do not appear in Γ . Now note that the functions f and g are inverses of each other: $g(f(p_\Gamma)) = g(p_{\Gamma'}) = p_\Gamma$ and $f(g(p_{\Gamma'})) = f(p_\Gamma) = p_{\Gamma'}$. Hence, there is a bijective correspondence between $\mathcal{G}(\Gamma')$ and $\mathcal{CE}^1(\Gamma)$.

- Hence, the set of probabilistic models on Γ that satisfy statistical Specker's principle, i.e., $\mathcal{CE}^1(\Gamma)$, are in one-to-one correspondence with the set of probabilistic models on Γ' which (by construction) satisfies structural Specker's principle so that $\mathcal{CE}^1(\Gamma') = \mathcal{G}(\Gamma')$.

We therefore have that $\mathcal{CE}^1(\Gamma) = \mathcal{CE}^1(\Gamma')|_{V(\Gamma)}$, where $\mathcal{CE}^1(\Gamma')|_{V(\Gamma)}$ denotes the probabilistic models induced on Γ by those on Γ' (ignoring the probabilities assigned to vertices in $V(\Gamma') \setminus V(\Gamma)$).

It is conceivable that a particular Γ may not admit probabilistic models from an operational theory $\mathbb{T}$, i.e., $\mathcal{J}(\Gamma) = \emptyset$. On the other hand, if Γ admits a representation in terms of measurement events admissible in $\mathbb{T}$, so that $\mathcal{J}(\Gamma) \neq \emptyset$, then two possibilities arise: Γ satisfies structural Specker's principle or it doesn't. If Γ satisfies structural Specker's principle then any probabilistic model in $\mathcal{J}(\Gamma)$ will satisfy statistical Specker's principle and we have $\Gamma' = \Gamma$. If Γ does not satisfy structural Specker's principle, we consider its relation with the contextuality scenario Γ' constructed from it that does satisfy structural Specker's principle. Such a Γ' admits a representation in a theory $\mathbb{T}$ satisfying structural Specker's principle (that is, $\mathcal{J}(\Gamma') \neq \emptyset$) as long as Γ admits such a representation (that is, $\mathcal{J}(\Gamma) \neq \emptyset$). Indeed, it's the satisfaction of structural Specker's principle in $\mathbb{T}$ that renders the construction of Γ' from Γ physically allowed in $\mathbb{T}$.

Thus, in a theory $\mathbb{T}$ that satisfies structural Specker's principle, the following holds: for every probabilistic model $p_\Gamma \in \mathcal{J}(\Gamma)$ ($\subseteq \mathcal{CE}^1(\Gamma)$), Γ' admits a corresponding probabilistic model $p_{\Gamma'} \in \mathcal{J}(\Gamma')$ satisfying $p_{\Gamma'}(v) = p_\Gamma(v)$ for all $v \in V(\Gamma)$ and $p_{\Gamma'}(v_c) = 1 - \sum_{v \in c} p_\Gamma(v)$ for all $c \in C$, where C is the set of maximal cliques in $O(\Gamma)$ such that none of them is a hyperedge in Γ . Similarly, given $p_{\Gamma'} \in \mathcal{J}(\Gamma')$

²³Recall that $\{v_c\}_{c \in C} = V(\Gamma') \setminus V(\Gamma)$.

($\subseteq \mathcal{CE}^1(\Gamma')$), $p_\Gamma \in \mathcal{T}(\Gamma)$ is uniquely fixed: it's obtained by just neglecting the probabilities assigned by $p_{\Gamma'}$ to the vertices in $V(\Gamma') \setminus V(\Gamma)$.

We must therefore have $\mathcal{T}(\Gamma) = \mathcal{T}(\Gamma')|_{V(\Gamma)}$ for any Γ , where $\mathcal{T}(\Gamma')|_{V(\Gamma)}$ denotes the set of probabilistic models induced on Γ by the set of probabilistic models in $\mathcal{T}(\Gamma')$ under the correspondence we have already established above. We can now state and prove the following theorem:

Theorem 2. *If an operational theory $\mathbb{T}$ satisfies structural Specker's principle, then it also satisfies statistical Specker's principle.*

Proof. For any Γ that does not admit a probabilistic model in $\mathbb{T}$, i.e., $\mathcal{T}(\Gamma) = \emptyset$, statistical Specker's principle is trivially satisfied since $\mathcal{T}(\Gamma) = \emptyset \subseteq \mathcal{CE}^1(\Gamma) \subseteq \mathcal{G}(\Gamma)$.

For any Γ that does admit a probabilistic model in $\mathbb{T}$, i.e., $\mathcal{T}(\Gamma) \neq \emptyset$, we can have one of two possibilities: either it satisfies structural Specker's principle, in which case $\mathcal{T}(\Gamma) \subseteq \mathcal{CE}^1(\Gamma) = \mathcal{G}(\Gamma)$, or it doesn't, in which case we consider the Γ' constructed from it following the recipe we have already outlined so that we have:

$$\mathcal{T}(\Gamma')|_{V(\Gamma)} \subseteq \mathcal{G}(\Gamma')|_{V(\Gamma)} = \mathcal{CE}^1(\Gamma')|_{V(\Gamma)} = \mathcal{CE}^1(\Gamma).$$

Since $\mathbb{T}$ satisfies structural Specker's principle, we have $\mathcal{T}(\Gamma) = \mathcal{T}(\Gamma')|_{V(\Gamma)}$, which immediately implies that $\mathcal{T}(\Gamma) \subseteq \mathcal{CE}^1(\Gamma)$. That is, the theory $\mathbb{T}$ satisfies statistical Specker's principle on Γ : $\mathcal{T}(\Gamma) \subseteq \mathcal{CE}^1(\Gamma) \subseteq \mathcal{G}(\Gamma)$.

Overall, we have the desired result: $\mathbb{T}$ satisfies structural Specker's principle $\Rightarrow \mathcal{T}(\Gamma) \subseteq \mathcal{CE}^1(\Gamma) \subseteq \mathcal{G}(\Gamma)$ for all Γ , i.e., $\mathbb{T}$ satisfies statistical Specker's principle. $\square$

Thus, *one* way of enforcing that a particular operational theory $\mathbb{T}$ satisfies statistical Specker's principle — that is, $\mathcal{T}(\Gamma) \subseteq \mathcal{CE}^1(\Gamma) \subseteq \mathcal{G}(\Gamma)$ for all Γ — is to require that it satisfies structural Specker's principle, a constraint on the structure of measurement events in $\mathbb{T}$. This is, for example, what is achieved in Ref. [30] by invoking a notion of “sharpness” for measurement events in an operational theory such that any set of sharp measurement events that are pairwise jointly measurable are all jointly measurable. That is, structural Specker's principle is satisfied in a theory with such sharp measurement events and, consequently, statistical Specker's principle, or what is more conventionally called *consistent exclusivity* [23], is also satisfied. But it's conceivable that there may be *other* ways to ensure that only a subset of $\mathcal{CE}^1(\Gamma)$ probabilistic models are allowed in $\mathcal{T}(\Gamma)$ for any Γ . What we wish to emphasize here is that it is by no means obvious (or at least, it needs to be proven) that the *only* way to restrict the set of probabilistic models

$\mathcal{T}(\Gamma)$ to a subset of $\mathcal{CE}^1(\Gamma)$ for any Γ is to require that the theory $\mathbb{T}$ satisfy structural Specker's principle.²⁴

Corollary 1. *For any operational theory $\mathbb{T}$, the following implications hold:*

$$\begin{aligned} &\mathbb{T} \text{ satisfies Specker's principle} \\ \Rightarrow &\mathbb{T} \text{ satisfies structural Specker's principle} \end{aligned} \quad (42)$$

$$\begin{aligned} \Rightarrow &\mathbb{T} \text{ satisfies statistical Specker's principle,} \\ &\text{i.e., consistent exclusivity.} \end{aligned} \quad (43)$$

Proof. This follows from combining Theorems 1 and 2. $\square$

Note that statistical Specker's principle (or consistent exclusivity) is so intrinsic to the CSW approach [22] that they do not consider probabilistic models that do not satisfy this principle.²⁵ This will become important when we consider the fact that nonprojective measurements in quantum theory *do not* satisfy Specker's principle, structural or statistical (at the level of measurement *events*), and thus also fail to satisfy the stronger statement of Specker's principle for measurement *settings* (cf. Ref. [49]). Indeed, such measurements admit contextuality scenarios Γ that are not possible with projective measurements, such as the one from three binary-outcome POVMs that are pairwise jointly measurable but not triplewise so [39–41], and the probabilistic models they give rise to can only be accommodated in the most general set of probabilistic models, $\mathcal{G}(\Gamma)$, since trivial POVMs can realize any probabilistic model. Specker's principle, structural Specker's principle, and statistical Specker's principle were all motivated by the fact that projective measurements in quantum theory satisfy them. In particular, consistent exclusivity (or statistical Specker's principle) would be obeyed in any theory where measurement events satisfy structural Specker's principle, and indeed, the more recent approach [29] is to restrict attention to “sharp” measurements in such theories [30, 31], where the definition of “sharp” ensures the property of pairwise jointly measurable events being globally jointly measurable. This property forms the motivational basis

²⁴Indeed, any putative theory yielding the set of almost quantum correlations (which satisfy statistical Specker's principle) [50] *cannot* satisfy Specker's principle — that pairwise joint implementable measurement *settings* are all jointly implementable — for any notion of sharp measurements [49]. Whether structural Specker's principle, which is defined at the level of measurement *events*, can be upheld for an almost quantum theory — so that it falls in the category of operational theories with sharp measurements envisaged in Ref. [30] — remains an open question.

²⁵As we have already noted, a noise-robust noncontextuality inequality of the type in Ref. [12] that is based on a logical proof of the KS theorem is not even obtainable if one restricted attention to probabilistic models satisfying $\mathcal{CE}^1$. The upper bound on that inequality comes from a probabilistic model that *does not* satisfy $\mathcal{CE}^1$.

(and is *sufficient*) for statistical Specker’s principle to hold (cf. Theorem 2). That is, this approach [29, 30] regards statistical Specker’s principle as grounded in (and physically justified by) structural Specker’s principle. Theorem 2 is a precise statement of this intuition in the hypergraph formalism à la AFLS [23]. The work of Refs. [30, 31] can be understood as bridging the gap between structural Specker’s principle and statistical Specker’s principle by formally defining a notion of sharp measurements in an operational theory such that structural Specker’s principle holds for these sharp measurements.

On the other hand, and this is the key point for our purposes, if one wants to make no commitment about the representation of measurements in the operational theory (in particular, not requiring a notion of “sharpness”), then Specker’s principle is not a natural constraint to impose on probabilistic models and, indeed, one must deal with the full set of probabilistic models $\mathcal{G}(\Gamma)$ on any contextuality scenario Γ rather than restrict oneself to the set of probabilistic models $\mathcal{CE}^1(\Gamma)$. It is for this reason that we are translating the notions from CSW [22] to the notational conventions of AFLS [23], the latter being a more natural choice for our purposes, allowing the language needed to articulate the difference between $\mathcal{CE}^1(\Gamma)$ and $\mathcal{G}(\Gamma)$ rather than excluding the latter by fiat or, perhaps, by an appeal to structural Specker’s principle holding for sharp measurements in the landscape of operational theories under consideration (cf. Theorem 2). It is for all these reasons that the “exclusivity principle” à la CSW [22] is not enough to make sense of Spekkens contextuality applied to Kochen-Specker type scenarios. The framework we propose in this paper addresses this gap between the notions Spekkens contextuality (which applies to arbitrary measurements) requires in a hypergraph framework and those that the CSW framework [22] (which applies to “sharp” measurements) can provide in its graph-theoretic formulation.

3.1.4 Remark on the classification of probabilistic models: why we haven’t defined “quantum models” as those obtained from projective measurements

The reader may note that we haven’t tried to define any notion of a “quantum model” so far, having only adopted the definitions of Ref. [23] for KS-noncontextual models ($\mathcal{C}(\Gamma)$), for models satisfying consistent exclusivity ($\mathcal{CE}^1(\Gamma)$), and for general probabilistic models ($\mathcal{G}(\Gamma)$). The reason for this is that we do not wish to restrict ourselves to projective measurements in defining a “quantum model”, unlike the traditional Kochen-Specker approaches [22, 23]. In Ref. [23], a *quantum model* is defined as a probabilistic model that can be realized in the following manner: assign projectors $\{\Pi_v\}_{v \in V(\Gamma)}$ (defined on any Hilbert space) to all the vertices of Γ such that $\sum_{v \in e} \Pi_v = \mathbb{I}$ for all $e \in E(\Gamma)$, and we have $p(v) = \text{Tr}(\rho \Pi_v)$, for

some density operator ρ on the Hilbert space, $\mathbb{I}$ being the identity operator.

On the other hand, allowing arbitrary positive operator-valued measures (POVMs) in a definition of a quantum model (as we would rather prefer) means that, in fact, quantum models on a hypergraph Γ are as general as the general probabilistic models $\mathcal{G}(\Gamma)$, rendering such a definition *redundant*. This can be seen by noting that for any probabilistic model $p \in \mathcal{G}(\Gamma)$, one can associate positive operators to the vertices of Γ given by $p(v)\mathbb{I}$ such that for any quantum state ρ on some Hilbert space, we have $p(v) = \text{Tr}(\rho p(v)\mathbb{I})$, where $\mathbb{I}$ is the identity operator.

Our focus in this paper is not on quantum theory, in particular, even though the need to be able to handle noisy measurements and preparations (particularly, trivial POVMs) in quantum theory can be taken as a motivation for this work. Rather, our focus is on delineating the boundary between operational theories that admit noncontextual ontological models (for Kochen-Specker type experiments, suitably augmented with multiple preparation procedures, as outlined in this paper) and those that don’t by obtaining noise-robust noncontextuality inequalities. In particular, we want these inequalities to indicate the noise thresholds beyond which an experiment cannot rule out the existence of a noncontextual ontological model with respect to the quantities of interest. This also means that making sense of quantum correlations in this approach requires one to pay attention not only to the measurements involved in an experiment but *also* the preparations; indeed, this shift of focus from measurements alone, to include multiple preparations (or source settings), is a fundamental conceptual difference between our approach and that of traditional Kochen-Specker contextuality frameworks [22, 23, 25].

3.1.5 Scope of this framework

Note that whenever we refer to the “CSW framework”, we mean the framework of Ref. [22], which often differs from the framework of Ref. [21] in some respects, e.g., the normalization of probabilities in a given hyperedge, assumed in [22], but not in [21]. In Ref. [21], the authors write:

Notice that in all of the above we never require that any particular context should be associated to a complete measurement: the conditions only make sure that each context is a subset of outcomes of a measurement and that they are mutually exclusive. Thus, unlike the original KS theorem, it is clear that every context hypergraph Γ has always a classical noncontextual model, besides possibly quantum and generalized models.

On the other hand, in Ref. [22], they write:

The fact that the sum of probabilities of outcomes of a test is 1 can be used to

Figure 4: The KS-uncolourable hypergraph from Ref. [51] that is not covered by our generalization of the CSW framework. We denote this hypergraph as Γ_{18} .

express these correlations as a positive linear combination of probabilities of events, $S = \sum_i w_i P(e_i)$, with $w_i > 0$.

The latter presentation [22] is more in line with the “original KS theorem” [19], as well as the presentation in Ref. [23]. Since normalization of probabilities is thus presumed in Ref. [22], in keeping with the definition of a probabilistic model we have presented (following [23]), the graph invariants of CSW [22] refer, specifically, to subgraphs G of *those* hypergraphs Γ on which the set of KS-noncontextual probabilistic models is non-empty. In particular, our generalization of the CSW framework [22] in this paper says nothing about noise-robust noncontextuality inequalities from logical proofs of the Kochen-Specker theorem [19], which rely on hypergraphs Γ that admit no KS-noncontextual probabilistic models, i.e., KS-uncolourable hypergraphs. It also says nothing for the hypergraphs Γ that do not satisfy the property $\mathcal{CE}^1(\Gamma) = \mathcal{G}(\Gamma)$. An example of such a hypergraph, which is not covered by our generalization of the CSW framework on both counts, is the 18 ray hypergraph first presented in Ref. [51], denoted Γ_{18} (see Fig. 4 and Appendix D). Indeed, the study of noise-robust noncontextuality inequalities from such KS-uncolourable hypergraphs was initiated in Ref. [12], and a more exhaustive hypergraph-theoretic treatment of it is presented in Ref. [34]. In this paper, we will restrict ourselves to KS-colourable hypergraphs, the study of which was initiated in Ref. [16], and, of these, *only* those KS-colourable hypergraphs Γ which satisfy $\mathcal{CE}^1(\Gamma) = \mathcal{G}(\Gamma)$. Note that this is not a limitation of our general approach, which is based on Ref. [16] and applies to *any* KS-colourable hypergraph, but rather a limitation we inherit from the CSW framework [22]²⁶

²⁶Ref. [22] takes Specker’s principle to be fundamental and identifies $\mathcal{CE}^1(\Gamma_{18})$ as the most general set of probabilistic mod-

els, which is *not* the case for Γ_{18} (for example). See Appendix D for a detailed discussion of this point.

since we want to leverage their graph invariants in obtaining our noise-robust noncontextuality inequalities. The study of other KS-colourable hypergraphs, in particular those which arise *only* with nonprojective measurements in quantum theory [39–41] and are outside the scope of traditional frameworks [22, 23, 25], will be taken up in future work.

To summarize, the measurement events hypergraphs Γ where the present framework (and the CSW framework [22]) applies *must* satisfy two properties: $\mathcal{C}(\Gamma) \neq \emptyset$ (that is, *KS-colourability*) and $\mathcal{CE}^1(\Gamma) = \mathcal{G}(\Gamma)$.²⁷

In the next subsection, we define additional notions necessary to obtain noise-robust noncontextuality inequalities that make use of graph invariants from the CSW framework. These notions correspond to *source events* that are an integral part of our framework.

3.2 Sources

Having introduced the (hyper)graph-theoretic elements that we need to talk about measurement events, we are now in a position to introduce features of source events that are relevant in the Spekkens framework. This part of our framework has no precedent in the literature on KS-noncontextuality, in particular the CSW framework [22]. We introduce these source events in order to benchmark the measurement events against them, i.e., for every measurement event, we seek to identify in the operational theory a corresponding source event that makes this measurement event as likely as possible. This helps us deal with cases where a measurement device may be implementing very noisy measurements by explicitly accounting for this noise in our noise-robust noncontextuality inequalities. Further, while we do not assume outcome determinism (which is essential to KS-noncontextuality), we will invoke preparation noncontextuality with respect to these source events in the Spekkens framework [18]. As an example of what we mean by “benchmarking” a measurement event against a source event, consider the case of quantum

²⁷As we have shown, when the operational theory $\mathbb{T}$ under consideration satisfies structural Specker’s principle, we can always turn a hypergraph Γ that doesn’t satisfy structural Specker’s principle into a hypergraph Γ' that satisfies it and for which, therefore, $\mathcal{CE}^1(\Gamma') = \mathcal{G}(\Gamma')$ holds. This can be seen as justification for restricting oneself to probabilistic models satisfying consistent exclusivity in the CSW framework [22]: such a restriction is not really a restriction if the theory satisfies structural Specker’s principle. On the other hand, we restrict ourselves to hypergraphs for which $\mathcal{CE}^1(\Gamma) = \mathcal{G}(\Gamma)$ without assuming that $\mathbb{T}$ satisfies structural Specker’s principle. The justification for this seemingly *ad hoc* restriction is simply that it is necessary in order to meaningfully leverage the graph invariants of CSW [22] – in particular, the fractional packing number – in our noise-robust noncontextuality inequalities. This will become clear when we obtain our noise-robust noncontextuality inequalities.

theory, where any measurement event represented by a projector occurs with probability 1 for any source event that is represented by an eigenstate of this projector; on the other hand, a positive operator that isn't projective cannot occur with a probability greater than its largest eigenvalue (< 1) for any source event. We now proceed to describe the necessary hypergraph-theoretic ingredients we need to accommodate source events in our framework.

As we have argued previously, we require the measurement events hypergraph Γ to be such that $\mathcal{C}(\Gamma) \neq \emptyset$ and $\mathcal{C}^1(\Gamma) = \mathcal{G}(\Gamma)$ to be able to obtain noise-robust noncontextuality inequalities that use graph invariants from the CSW framework [22]. Hence, we will restrict ourselves to experiments that realize the operational equivalences represented by this class of Γ . Now, in the CSW framework [22], every Bell-KS expression picks out a particular subgraph G of the orthogonality graph $O(\Gamma)$ of the contextuality scenario Γ of interest. This amounts to focussing on a restricted set of probabilities (for the vertices of G) rather than probabilities for all the measurement events (represented by vertices of Γ) in the experiment. Hence, the vertices of G denote the measurement events of interest in a given Bell-KS expression and we have the following:

- A general probabilistic model $p \in \mathcal{G}(\Gamma)$ will assign probabilities to vertices in G such that: $p(v) \geq 0$ for all $v \in V(G)$ and $p(v) + p(v') \leq 1$ for every edge $\{v, v'\} \in E(G)$.
- A probabilistic model $p \in \mathcal{C}^1(\Gamma)$ will assign probabilities to vertices in G such that: $p(v) \geq 0$ for all $v \in V(G)$ and

$$\sum_{v \in c} p(v) \leq 1, \quad (44)$$

for every clique $c \subseteq V(G)$.

- A probabilistic model $p \in \mathcal{C}(\Gamma)$ will assign probabilities to vertices in G such that: $p(v) = \sum_k \Pr(k) p_k(v)$, where $\Pr(k) \geq 0$, $\sum_k \Pr(k) = 1$, and for each k , p_k is a deterministic assignment $p_k(v) \in \{0, 1\}$ for all $v \in V(G)$, and $p_k(v) + p_k(v') \leq 1$ for every edge $\{v, v'\} \in E(G)$.

Since Γ is such that $\mathcal{C}^1(\Gamma) = \mathcal{G}(\Gamma)$, the condition

$$\sum_{v \in c} p(v) \leq 1 \text{ for every clique } c \subseteq V(G)$$

on the probabilities assigned to vertices in G is redundant. We now obtain a simplified hypergraph, Γ_G , from G as follows: convert all maximal cliques in G to hyperedges and add an extra (no-detection) vertex to each such hyperedge.²⁸

²⁸Physically, a “no-detection” vertex denotes the case when none of the measurement events of interest (here, the events in G) for a given measurement setting occur.

Figure 5: The hypergraph Γ_G obtained from G by adding a no-detection vertex (represented by a hollow circle) to every maximal clique in G .

This Γ_G , for any G , will satisfy the property that $\mathcal{C}^1(\Gamma_G) = \mathcal{G}(\Gamma_G)$ and any probabilistic model on Γ assigning probabilities to measurement events in G will correspond to a probabilistic model on Γ_G which also assigns the same probabilities to measurement events in G . Formally: $V(\Gamma_G) \equiv V(G) \sqcup \{v_c | c \text{ is a maximal clique in } G\}$, and $E(\Gamma_G) \equiv \{c \sqcup \{v_c\} | c \text{ is a maximal clique in } G\}$, where v_c is the extra no-detection vertex added to the hyperedge corresponding to maximal clique c in G .

We have the following probabilistic model on Γ_G , given a probabilistic model $p \in \mathcal{G}(\Gamma)$: the probabilities assigned to the vertices in $V(G) \subseteq V(\Gamma_G)$ are the same as specified by $p \in \mathcal{G}(\Gamma)$ and the probabilities assigned to the remaining vertices in $V(\Gamma_G) \setminus V(G)$ are given by $p(v_c) = 1 - \sum_{v \in c} p(v)$, for every maximal clique c in G . Consider, for example, the KCBS scenario [16, 22, 47]: the 20-vertex Γ representing measurement events from five 4-outcome joint measurements (Fig. 2), its 5 vertices G involved in the KCBS inequality (Fig. 3), and 10-vertex hypergraph Γ_G constructed from G (Fig. 5).

Given Γ_G , constructed from G , we now require that the operational theory that realizes measurement events in Γ_G also admits preparations that can be represented by a hypergraph Σ_G of source events as follows: for every hyperedge $e \in E(\Gamma_G)$, corresponding to the choice of measurement setting M_e , we define a hyperedge $e \in E(\Sigma_G)$ denoting a corresponding choice of source setting S_e . And for every vertex $v \in e (\in E(\Gamma_G))$, we define a vertex $v_e \in e (\in E(\Sigma_G))$.²⁹ Hence, every measurement event

²⁹Recall from the discussion at the beginning of Section 3.2 that we seek to benchmark the measurement events against those source events in the operational theory that (ideally) make them as predictable as possible. The source setting against which the predictability of a particular measurement

Figure 6: The source events hypergraph with the operational equivalences between the source settings separately specified.

$[v|e]$ in Γ_G corresponds to a vertex v_e of Σ_G , and the number of such vertices in $V(\Sigma_G)$ is $|V(\Gamma_G)||E(\Gamma_G)|$. This means that the operational equivalences between the measurement events that are implicit in Γ_G — such as $[v|e]$ is operationally equivalent to $[v|e']$, where $e, e' \in E(\Gamma_G)$ are distinct hyperedges that share the vertex (representing an equivalence class of measurement events) $v \in V(\Gamma_G)$ — are not carried over to the source events, where none is presumed to be operationally equivalent to any other, hence $v_e \in V(\Sigma_G)$ is a different vertex from $v_{e'} \in V(\Sigma_G)$. Here v_e ($v_{e'}$) represents a source event $[s_e|S_e]$ ($[s_{e'}|S_{e'}]$), rather than an equivalence class of source events.

Besides these $|V(\Gamma_G)||E(\Gamma_G)|$ vertices in $V(\Sigma_G)$ and the associated hyperedges $e \in E(\Sigma_G)$, we require that the operational theory admits an additional hyperedge $e_* \in E(\Sigma_G)$, representing a source setting S_{e_*} , containing two new vertices $v_{e_*}^0, v_{e_*}^1 \in V(\Sigma_G)$. Here $v_{e_*}^0$ represents the source event $[s_{e_*} = 0|S_{e_*}]$ and $v_{e_*}^1$ represents the source event $[s_{e_*} = 1|S_{e_*}]$. Hence, we have $|V(\Sigma_G)| = |V(\Gamma_G)||E(\Gamma_G)| + 2$ and $|E(\Sigma_G)| = |E(\Gamma_G)| + 1$.

The operational equivalence we *do* require for Σ_G (in any operational theory that admits source events represented by Σ_G) applies to the source settings: all source settings, each represented by coarse-graining the source events in a hyperedge $e \in E(\Sigma_G)$, are operationally equivalent, i.e., $[\top|S_{e_\top}] \simeq [\top|S_{e'_\top}]$ for all $e, e' \in E(\Sigma_G)$, i.e., $\forall [m|M] : \sum_{s_e} p(m, s_e|M, S_e) = \sum_{s_{e'}} p(m, s_{e'}|M, S_{e'})$, for all $e, e' \in E(\Sigma_G)$.

An example of such a source events hypergraph was considered in Ref. [12], albeit without the additional

setting is tested — that is the predictability of each measurement event (e.g., $v \in e \in E(\Gamma_G)$) for this measurement setting (e.g., M_e) is benchmarked against some source event (e.g., $v_e \in e \in E(\Sigma_G)$) for the source setting (e.g., S_e) — is the “corresponding choice of source setting S_e ”. In Section 5.2 we will see how these pairs of source and measurement settings are used to compute an operational quantity relevant for our noise-robust noncontextuality inequalities.

source labelled by e_* here [16]. We illustrate it here in Fig. 6 for the KCBS scenario.

4 A key hypergraph invariant: the weighted max-predictability

We now define a hypergraph invariant that will be relevant for our noise-robust noncontextuality inequalities:

$$\beta(\Gamma_G, q) \equiv \max_{p \in \mathcal{G}(\Gamma_G)|_{\text{ind}}} \sum_{e \in E(\Gamma_G)} q_e \zeta(M_e, p), \quad (45)$$

where $q_e \geq 0$ for all $e \in E(\Gamma_G)$, $\sum_{e \in E(\Gamma_G)} q_e = 1$, and

$$\zeta(M_e, p) \equiv \max_{v \in e} p(v)$$

is the maximum probability assigned to a vertex in $e \in E(\Gamma_G)$ by an extremal indeterministic probabilistic model $p \in \mathcal{G}(\Gamma_G)|_{\text{ind}}$.³⁰

We call $\beta(\Gamma_G, q)$ the *weighted max-predictability* of the measurement settings (i.e., hyperedges) in Γ_G , where the hyperedges $e \in E(\Gamma_G)$ are weighted according to the probability distribution $q \equiv \{q_e\}_{e \in E(\Gamma_G)}$.

We now outline how this quantity is related to properties of an operational theory $\mathbb{T}$ admitting a measurement noncontextual ontological model. Γ_G represents a particular configuration of operational equivalences that a set of measurement events in $\mathbb{T}$ may realize. The probabilistic models on Γ_G that can be realized by $\mathbb{T}$ are, as earlier, denoted by $\mathcal{T}(\Gamma_G)$. Since $\mathbb{T}$ admits a measurement noncontextual ontological model,³¹ its predictions for the specific case of Γ_G can be reproduced by such a model. But since, in keeping with the CSW approach [22], we will look at witnesses of contextuality tailored to particular experiments (Γ_G representing features of one such experiment), we do not need an ontological model for the full theory $\mathbb{T}$ to reproduce its predictions for a particular experiment. Indeed, to construct a measurement noncontextual ontological model for the set of probabilistic models $\mathcal{T}(\Gamma_G)$, it suffices to assume (without loss of generality) that the extremal probabilistic models on Γ_G — given by $\mathcal{G}(\Gamma_G)|_{\text{det}} \sqcup \mathcal{G}(\Gamma_G)|_{\text{ind}}$

³⁰An extremal *indeterministic* probabilistic model refers to those extremal $p \in \mathcal{G}(\Gamma_G)$ for which $\zeta(M_e, p) < 1$ for some $e \in E(\Gamma_G)$.

³¹This will always be the case for any operational theory we consider: the assumption of measurement noncontextuality on its own can always be satisfied by a trivial ontological model of the type we outlined in Section 2.5. Indeed, quantum theory satisfies it, the Beltrametti-Bugajski model [52] that was discussed in Ref. [18] being an example of a measurement noncontextual ontological model of quantum theory. It is only when this assumption is supplemented with something else — outcome determinism in the case of KS-noncontextuality and preparation noncontextuality in the case of generalized noncontextuality [18] — that it can produce a contradiction with the predictions of an operational theory.

– are in bijective correspondence with the ontic states (Λ) of the physical system on which the measurements are carried out. This is because, firstly, *any* probabilistic model in $\mathcal{G}(\Gamma_G)$ can be expressed as a convex mixture of extremal probabilistic models in $\mathcal{G}(\Gamma_G)|_{\text{det}} \sqcup \mathcal{G}(\Gamma_G)|_{\text{ind}}$, and, secondly, associating each ontic state in the ontological model with an extremal probabilistic model³² in $\mathcal{G}(\Gamma_G)|_{\text{det}} \sqcup \mathcal{G}(\Gamma_G)|_{\text{ind}}$ means that *any* probabilistic model in $\mathcal{G}(\Gamma_G)$ corresponding to predictions of an operational theory (in particular, any $p \in \mathcal{T}(\Gamma_G) \subseteq \mathcal{G}(\Gamma_G)$) can be obtained by an appropriate probability distribution over this set of ontic states. Denoting the set of ontic states corresponding to $\mathcal{G}(\Gamma_G)|_{\text{det}}$ by Λ_{det} and the set of ontic states corresponding to $\mathcal{G}(\Gamma_G)|_{\text{ind}}$ by Λ_{ind} , we have that the measurement noncontextual ontological model given by $\Lambda \equiv \Lambda_{\text{det}} \sqcup \Lambda_{\text{ind}}$ reproduces the predictions $\mathcal{T}(\Gamma_G)$ of any operational theory $\mathbb{T}$ that admits a measurement noncontextual ontological model: that is, for every $p \in \mathcal{G}(\Gamma_G)$ (and therefore also $p \in \mathcal{T}(\Gamma_G)$),

$$p(v) = \sum_{\lambda \in \Lambda} \xi(v|\lambda) \mu(\lambda)$$

for all $v \in V(\Gamma_G)$, for some probability distribution $\mu : \Lambda \rightarrow [0, 1]$ such that $\sum_{\lambda \in \Lambda} \mu(\lambda) = 1$.³³ We can also then rewrite $\beta(\Gamma_G, q)$ as

$$\beta(\Gamma_G, q) = \max_{\lambda \in \Lambda_{\text{ind}}} \sum_{e \in E(\Gamma_G)} q_e \zeta(M_e, \lambda), \quad (46)$$

where $\zeta(M_e, \lambda) \equiv \max_{m_e} \xi(m_e|M_e, \lambda)$.

5 Noise-robust noncontextuality inequalities

We will now proceed to obtain our noise-robust noncontextuality inequalities following the ideas outlined in Ref. [16].

5.1 Key notions from CSW

We first recall some key notions from the CSW framework [22] before obtaining our inequalities.

Consider the positive linear combination of the probabilities of measurement events,

$$R([s|S]) \equiv \sum_{v \in V(G)} w_v p(v|S, s), \quad (47)$$

³²Representing response functions for the ontic state, i.e., $p(v) = \xi(v|\lambda), \forall v \in V(\Gamma_G)$

³³As a corollary, note that as long as the polytope $\mathcal{G}(\Gamma_G)$ has a finite number of extreme points, we can take the ontic state space to consist of a finite number of ontic states (as we have done) without any loss of generality. The hypergraphs Γ_G we study – representing the measurement events of interest in a contextuality experiment – have this property because of their finiteness.

where $w_v > 0$ for all $v \in V(G)$.

The fundamental result of CSW is that this quantity is bounded for different sets of correlations — KS-noncontextual, those realizable by projective quantum measurements, and those satisfying consistent exclusivity — by graph-theoretic invariants as follows:

$$\forall [s|S] : R([s|S]) \stackrel{\text{KS}}{\leq} \alpha(G, w) \stackrel{\text{Q}}{\leq} \theta(G, w) \stackrel{\text{CE}^1}{\leq} \alpha^*(G, w), \quad (48)$$

where KS denotes operational theories that admit KS-noncontextual ontological models and thus realize probabilistic models on Γ_G that fall in the set $\mathcal{C}(\Gamma_G)$, Q denotes quantum theory with projective measurements which assigns probabilistic models on Γ_G denoted by $\mathcal{Q}(\Gamma_G)$, and CE^1 denotes operational theories satisfying consistent exclusivity and thus realizing the set of probabilistic models $\mathcal{CE}^1(\Gamma_G)$ on Γ_G . The graph invariants of the weighted graph (G, w) , namely, $\alpha(G, w)$, $\theta(G, w)$, and $\alpha^*(G, w)$ are defined as follows:

1. Independence number $\alpha(G, w)$:

$$\alpha(G, w) \equiv \max_{\mathcal{I}} \sum_{v \in \mathcal{I}} w_v, \quad (49)$$

where $\mathcal{I} \subseteq V(G)$ is an *independent set* of vertices of G , i.e., a set of nonadjacent vertices of G , so that none of the vertices in this set shares an edge with any other vertex in the set.

2. Lovasz theta number $\theta(G, w)$:

$$\theta(G, w) \equiv \max_{\{|u_v\rangle\}_{v \in V(G)}, |\psi\rangle} \sum_{v \in V(G)} w_v |\langle \psi | u_v \rangle|^2, \quad (50)$$

where $\{|u_v\rangle\}_{v \in V(G)} = \{|u_v\rangle\}_{v \in V(\bar{G})}$ (each $|u_v\rangle$ a unit vector in $\mathbb{R}^d$) is an *orthonormal representation* (OR) of the complement of G , namely, $\bar{G}$, and the unit vector $|\psi\rangle \in \mathbb{R}^d$ is called a *handle*.

Here $V(\bar{G}) \equiv V(G)$ and $E(\bar{G}) \equiv \{(v, v') | v, v' \in V(G), (v, v') \notin E(G)\}$, and we have in an orthonormal representation that $\langle u_{v''} | u_{v'''} \rangle = 0$ for all pairs of nonadjacent vertices, (v'', v''') , in $\bar{G}$, or equivalently, for all $(v'', v''') \in E(G)$.

3. Fractional packing number $\alpha^*(G, w)$:

$$\alpha^*(G, w) \equiv \max_{\{p_v\}_{v \in V(G)}} \sum_{v \in V(G)} w_v p_v, \quad (51)$$

where $\{p_v\}_{v \in V(G)}$ is such that $p_v \geq 0$ for all $v \in V(G)$ and $\sum_{v \in c} p_v \leq 1$ for all cliques c in G .

Note that since we are always considering Γ_G such that $\mathcal{CE}^1(\Gamma_G) = \mathcal{G}(\Gamma_G)$, we, in fact, have the bounds

$$\forall [s|S] : R([s|S]) \stackrel{\text{KS}}{\leq} \alpha(G, w) \stackrel{\text{Q}}{\leq} \theta(G, w) \stackrel{\text{GPT}}{\leq} \alpha^*(G, w), \quad (52)$$

where “GPT” denotes the full set of probabilistic models on Γ_G , i.e., $\mathcal{G}(\Gamma_G)$.

In terms of the notation we have already introduced, where $R([s|S]) \leq R_{\text{KS}}$ was a Bell-KS inequality, we now have — from CSW [22] — that $R_{\text{KS}} = \alpha(G, w)$.

5.2 Key notion *not* from CSW: source-measurement correlation, Corr

We need to define a new quantity not in the CSW framework, namely,

$$\text{Corr} \equiv \sum_{e \in E(\Gamma_G)} q_e \sum_{m_e, s_e} \delta_{m_e, s_e} p(m_e, s_e | M_e, S_e), \quad (53)$$

where $\{q_e\}_{e \in E(\Gamma_G)}$ is a probability distribution, i.e., $q_e \geq 0$ for all $e \in E(\Gamma_G)$ and $\sum_{e \in E(\Gamma_G)} q_e = 1$, such that $\beta(\Gamma_G, q) < 1$ holds.³⁴ In previous work [12, 16], we have taken q to be the uniform distribution $q_e = \frac{1}{|E(\Gamma_G)|}$, but the derivation of the noncontextuality inequalities is independent of that choice (as we’ll see here). Also, note that we have chosen the following labelling convention for outcomes of source setting S_e (namely, s_e) and measurement setting M_e (namely, m_e): the source outcomes s_e for source setting S_e take values in the same set as measurement outcomes m_e for measurement setting M_e , i.e., $V_{S_e} = V_{M_e}$ (recalling notation from Section 2). In particular, outcomes corresponding to the measurement event $[v|e]$ (representing $[m_e|M_e]$) and its corresponding source event v_e (representing $[s_e|S_e]$) are both denoted by the same label, so that $m_e = s_e$ for them. An example of this from Figs. 5 and 6 would be to, say, denote the outcomes of a particular $e \in E(\Gamma_G)$ (measurement setting M_e) by $m_e \in V_{M_e} \equiv \{0, 1, 2\}$ and corresponding outcomes of $e \in E(\Sigma_G)$ (source setting S_e) by $s_e \in V_{S_e} \equiv \{0, 1, 2\}$; so if $[v|e]$ denotes $[m_e = 0|M_e]$, then v_e will denote $[s_e = 0|S_e]$, etc.

5.3 Obtaining the noise-robust noncontextuality inequalities

5.3.1 Expressing operational quantities in ontological terms

We begin with expressing the operational quantities of interest in terms of a noncontextual ontological model. In an ontological model, $R([s|S])$ is given by

$$R([s|S]) = \sum_{\lambda \in \Lambda} \sum_{v \in V(G)} w_v p(v|\lambda) \mu(\lambda|S, s). \quad (54)$$

Defining $R(\lambda) \equiv \sum_{v \in V(G)} w_v p(v|\lambda)$, we have that

$$R([s|S]) = \sum_{\lambda \in \Lambda} R(\lambda) \mu(\lambda|S, s). \quad (55)$$

³⁴Indeed, for the strongest possible constraint on Corr, one must pick q such that $\beta(\Gamma_G, q)$ is minimized.

Similarly, Corr is given by

$$\begin{aligned} \text{Corr} &= \sum_{\lambda \in \Lambda} \sum_{e \in E(\Gamma_G)} q_e \sum_{m_e, s_e} \delta_{m_e, s_e} \xi(m_e|M_e, \lambda) \mu(\lambda, s_e|S_e) \\ &= \sum_{\lambda \in \Lambda} \sum_{e \in E(\Gamma_G)} q_e \\ &\quad \sum_{m_e, s_e} \delta_{m_e, s_e} \xi(m_e|M_e, \lambda) \mu(s_e|S_e, \lambda) \mu(\lambda|S_e). \end{aligned} \quad (56)$$

Here, we have used the fact that

$$\mu(\lambda, s_e|S_e) = \mu(s_e|S_e, \lambda) \mu(\lambda|S_e)$$

to express Corr in a way that treats sources and measurements similarly.

Using *preparation noncontextuality* (cf. Eq. (22)), we have that

$$\begin{aligned} \forall e, e' \in E(\Sigma_G) : [\top|S_{e\top}] &\simeq [\top|S_{e'\top}] \\ \Rightarrow \mu(\lambda|S_e) &= \mu(\lambda|S_{e'}) \equiv \nu(\lambda), \forall \lambda \in \Lambda. \end{aligned} \quad (57)$$

Then we can rewrite Corr as

$$\begin{aligned} \text{Corr} &= \sum_{\lambda \in \Lambda} \sum_{e \in E(\Gamma_G)} q_e \sum_{m_e, s_e} \delta_{m_e, s_e} \xi(m_e|M_e, \lambda) \mu(s_e|S_e, \lambda) \nu(\lambda). \end{aligned} \quad (58)$$

Note that the only λ that contribute to Corr are those for which $\nu(\lambda) > 0$. Also, $\mu(s_e|S_e, \lambda)$ and $\mu(\lambda|S_e, s_e)$ satisfy the condition $\mu(s_e|S_e, \lambda) \nu(\lambda) = \mu(\lambda|S_e, s_e) p(s_e|S_e)$, so that $\mu(s_e|S_e, \lambda)$ is well-defined whenever $\nu(\lambda) > 0$.

Defining

$$\text{Corr}(\lambda) \equiv \sum_{e \in E(\Gamma_G)} q_e \sum_{m_e, s_e} \delta_{m_e, s_e} \xi(m_e|M_e, \lambda) \mu(s_e|S_e, \lambda), \quad (59)$$

we have that

$$\text{Corr} = \sum_{\lambda \in \Lambda} \text{Corr}(\lambda) \nu(\lambda), \quad (60)$$

Recalling that $\zeta(M_e, \lambda) = \max_{m_e} \xi(m_e|M_e, \lambda)$, note that $\text{Corr}(\lambda)$ is upper bounded as follows (for any $\lambda \in \Lambda$):

$$\begin{aligned} \text{Corr}(\lambda) &= \sum_{e \in E(\Gamma_G)} q_e \sum_{m_e, s_e} \delta_{m_e, s_e} \xi(m_e|M_e, \lambda) \mu(s_e|S_e, \lambda) \\ &\leq \sum_{e \in E(\Gamma_G)} q_e \zeta(M_e, \lambda) \sum_{s_e} \mu(s_e|S_e, \lambda) \\ &= \sum_{e \in E(\Gamma_G)} q_e \zeta(M_e, \lambda). \end{aligned} \quad (61)$$

If $\lambda \in \Lambda_{\text{det}}$, then this upper bound is trivial, i.e., $\text{Corr}(\lambda) \leq 1$, since every measurement has deterministic response functions. On the other hand, for all $\lambda \in \Lambda_{\text{ind}}$, we have (from Eq. (46))

$$\text{Corr}(\lambda) \leq \beta(\Gamma_G, q). \quad (62)$$

Similarly, for $\lambda \in \Lambda_{\text{det}}$ we have $R(\lambda) \leq \alpha(G, w)$, while for $\lambda \in \Lambda_{\text{ind}}$ we have $R(\lambda) \leq \alpha^*(G, w)$.

Using the fact that

$$\nu(\lambda) = \mu(\lambda|S) = \sum_s \mu(\lambda|S, s)p(s|S),$$

for any $S \equiv S_{e_*}, e_* \in E(\Sigma_G)$, we have

$$\begin{aligned} & \text{Corr} \\ &= \sum_s \left(\sum_{\lambda} \text{Corr}(\lambda) \mu(\lambda|S, s) \right) p(s|S) \\ &= \sum_s \text{Corr}_s p(s|S). \end{aligned} \quad (63)$$

where we have defined $\text{Corr}_s \equiv \sum_{\lambda} \text{Corr}(\lambda) \mu(\lambda|S, s)$.

5.3.2 Derivation of the noncontextual tradeoff for any graph G

We are now in a position to express our general noise-robust noncontextuality inequality as a tradeoff between three operational quantities: Corr , $R([s_{e_*} = 0|S_{e_*}])$, and $p(s_{e_*} = 0|S_{e_*})$.

First, note that KS-contextuality is witnessed when for some choice of $[s|S]$, here given by $[s_{e_*} = 0|S_{e_*}]$, we have

$$R([s_{e_*} = 0|S_{e_*}]) > \alpha(G, w).$$

This means that for some set of ontic states in the support of $[s_{e_*} = 0|S_{e_*}]$, i.e.,

$$\begin{aligned} & \lambda \in \text{Supp}\{\mu(\cdot|S_{e_*}, s_{e_*} = 0)\} \\ & \equiv \{\lambda \in \Lambda : \mu(\lambda|S_{e_*}, s_{e_*} = 0) > 0\}, \end{aligned} \quad (64)$$

we have $R(\lambda) > \alpha(G, w)$. For such a set of ontic states one must then have $\text{Corr}(\lambda) < 1$ (because these $\lambda \in \Lambda_{\text{ind}}$ and we have Eq. (62)), which in turn implies that $\text{Corr}_{s_{e_*}=0} < 1$. On the other hand, for $s_{e_*} = 1$, we have no constraints: $\text{Corr}_{s_{e_*}=1} \leq 1$. Thus,

$$\begin{aligned} & \text{Corr} \\ &= \text{Corr}_{s_{e_*}=0} p(s_{e_*} = 0|S_{e_*}) + \text{Corr}_{s_{e_*}=1} p(s_{e_*} = 1|S_{e_*}) \\ &\leq p_0 \text{Corr}_{s_{e_*}=0} + 1 - p_0, \end{aligned} \quad (65)$$

where $p_0 \equiv p(s_{e_*} = 0|S_{e_*})$.

Defining $\mu_{\text{det}} \equiv \sum_{\lambda \in \Lambda_{\text{det}}} \mu(\lambda|S_{e_*}, s_{e_*} = 0)$ and $\mu_{\text{ind}} \equiv \sum_{\lambda \in \Lambda_{\text{ind}}} \mu(\lambda|S_{e_*}, s_{e_*} = 0)$, we now have

$$\mu_{\text{det}} + \mu_{\text{ind}} = 1, \quad (66)$$

$$\text{Corr}_{s_{e_*}=0} \leq \mu_{\text{det}} + \beta(\Gamma_G, q) \mu_{\text{ind}}, \quad (67)$$

$$R \leq \alpha(G, w) \mu_{\text{det}} + \alpha^*(G, w) \mu_{\text{ind}}. \quad (68)$$

Note that assuming $\mu_{\text{det}} = 1$ would reduce these constraints to a standard Bell-KS inequality, $R \leq \alpha(G, w)$. However, since we are not assuming this,

simply eliminating μ_{det} and μ_{ind} from these constraints leads us to³⁵

$$\begin{aligned} & \text{Corr}_{s_{e_*}=0} \\ &\leq 1 - (1 - \beta(\Gamma_G, q)) \frac{R - \alpha(G, w)}{\alpha^*(G, w) - \alpha(G, w)} \end{aligned} \quad (69)$$

where the upper bound is nontrivial if and only if $\beta(\Gamma_G, q) < 1$ and $R - \alpha(G, w) > 0$.

If we are given that $\beta(\Gamma_G, q) < 1$, then we have a trivial upper bound on $\text{Corr}_{s_{e_*}=0}$ for the remaining cases: the upper bound is 1 for $R = \alpha(G, w)$ and greater than 1 for $R < \alpha(G, w)$.

Thus, our noise-robust noncontextuality inequality now reads:

$$\text{Corr} \leq 1 - p_0(1 - \beta(\Gamma_G, q)) \frac{R - \alpha(G, w)}{\alpha^*(G, w) - \alpha(G, w)}, \quad (70)$$

which can be rewritten as

$$R \leq \alpha(G, w) + \frac{\alpha^*(G, w) - \alpha(G, w)}{p_0} \frac{1 - \text{Corr}}{1 - \beta(\Gamma_G, q)}. \quad (71)$$

Note that Eq. (70) expresses the constraint from noncontextuality as an upper bound on the source-measurement correlations Corr , reminiscent of the noise-robust noncontextuality inequality first derived in Ref. [12] (and later treated in hypergraph-theoretic terms in Ref. [34]), except here the upper bound on Corr depends not only on the hypergraph invariant $\beta(\Gamma_G, q)$ but also two of the graph invariants from the CSW framework [22], namely, $\alpha(G, w)$ and $\alpha^*(G, w)$, besides also the operational quantity R , which is the figure-of-merit for KS-contextuality ($R > \alpha(G, w)$ witnesses KS-contextuality) in the CSW framework. Eq. (70) indicates that the source-measurement correlations would fail to be perfect (i.e., $\text{Corr} < 1$) in an operational theory admitting a noncontextual ontological model if and only if $R > \alpha(G, w)$ and $\beta(\Gamma_G, q) < 1$. Contextuality would be witnessed when the source-measurement correlations are stronger than the constraint from Eq. (70). For $R \leq \alpha(G, w)$, in particular, there is no constraint from noncontextuality on Corr .

On the other hand, rewriting the constraint from noncontextuality as Eq. (71), one is reminded of the CSW framework [22], where R is taken to be the quantity that is upper bounded by KS-noncontextuality. Here, instead, we have that R is upper bounded by a term that includes the source-measurement correlations Corr that can be achieved for the measurements and thus penalizes for measurements that cannot be made highly predictable with respect to some preparations, i.e., $\text{Corr} < 1$ makes it harder to violate the

³⁵To see this explicitly, just use Eq. (66) to make the substitution $\mu_{\text{ind}} = 1 - \mu_{\text{det}}$ in Eqs. (67) and (68), then eliminate μ_{det} from Eq. (67) by using the upper bound on it from Eq. (68).

upper bound on R . When the upper bound reaches $\alpha^*(G, w)$, it becomes trivial and R is no longer constrained by noncontextuality on account of noise in the measurements. Indeed, trivial POVMs (cf. Appendices A.1.2 and C) never violate such a noncontextuality inequality because of the penalty incurred via Corr , as we later show in Section 6.3.

5.3.3 When is the noncontextual tradeoff violated?

The inequality of Eq. (71) can be rewritten as the following tradeoff between Corr , p_0 , and R :

$$\text{Corr} + p_0(1 - \beta(\Gamma_G, q)) \frac{R - \alpha(G, w)}{\alpha^*(G, w) - \alpha(G, w)} \leq 1. \quad (72)$$

Writing the constraint from noncontextuality in the form of Eq. (72) (in contrast to Eqs. (70) and (71)) makes it more even-handed in its treatment of the two operational quantities R (which is key in the CSW framework [22]) and Corr (which is key in noise-robust noncontextuality inequalities inspired by logical proofs of the KS theorem [12, 34]) and emphasizes that noise-robust noncontextuality inequalities inspired by statistical proofs of the KS theorem [16] are tradeoffs between R (which is about the strength of correlations *between* measurements) and Corr (which is about the predictability of measurements) that must be satisfied by any operational theory admitting a noncontextual ontological model. Roughly speaking, a high degree of predictability for measurements (e.g., $\text{Corr} = 1$) cannot coexist with very strong correlations *between* the measurements (e.g., $R = \alpha^*(G, w)$) when the operational theory admits a noncontextual ontological model.

For a nontrivial constraint – and hence, the possibility of witnessing contextuality via violation of this inequality (Eq. (72)) – the upper bound on Corr (the right-hand-side of Eq. (70)) should be strictly bounded above by 1, and the upper bound on R (the right-hand-side of Eq. (71)) should be strictly bounded above by $\alpha^*(G, w)$ (the algebraic upper bound on R), that is

$$\begin{aligned} p_0 &> 0 \text{ and } \beta(\Gamma_G, q) < 1, \\ R &> \alpha(G, w), \\ \text{Corr} &> 1 - p_0(1 - \beta(\Gamma_G, q)). \end{aligned} \quad (73)$$

These are the minimal benchmarks necessary — besides the requirement of tomographic completeness of a finite set of procedures and the possibility of inferring secondary procedures with exact operational equivalences using convexity of the operational theory [13] — to witness contextuality in a Kochen-Specker type experiment adapted to our framework following Spekkens [18].

Suppose one achieves, by some means, a value of $R = \theta(G, w)$, the upper bound on the quantum value with projective measurements. When would

this value be an evidence of contextuality? For this to be the case, we must have:

$$\text{Corr} > 1 - p_0(1 - \beta(\Gamma_G, q)) \frac{\theta(G, w) - \alpha(G, w)}{\alpha^*(G, w) - \alpha(G, w)}. \quad (74)$$

Now, for the ideal quantum realization where measurement events are projectors, and the corresponding source events are eigenstates, it is always the case that $\text{Corr} = 1$, hence contextuality is witnessed. However, it's possible to witness contextuality even if $\text{Corr} < 1$, as long as it exceeds the lower bound specified above. In a sense, for quantum theory, this allows for a quantitative accounting of the effect of nonprojectiveness in the measurements (or mixedness in preparations) on the possibility of witnessing contextuality, a feature that is absent in traditional Kochen-Specker approaches [21–23, 25]. Indeed, as long as one achieves any value of $R > \alpha(G, w)$, it is possible to witness contextuality for a sufficiently high value of Corr (see Eq. (70)).

5.4 Example: KCBS scenario

We will now illustrate our hypergraph framework by applying it to the KCBS scenario to make differences with respect to the CSW graph-theoretic framework [22] explicit.

The graph G for the KCBS scenario is given in Fig. 3, the measurement events hypergraph Γ_G is given in Fig. 5, and the source events hypergraph Σ_G is given in Fig. 6. We then have

$$R([s|S]) = \sum_{v \in V(G)} p(v|S, s), \quad (75)$$

where the (vertex) weights $w_v = 1$ for all $v \in V(G)$, i.e., it's an unweighted graph and we will use $\alpha(G)$ and $\alpha^*(G)$ to denote its independence number and the fractional packing number, respectively. These are given by

$$\alpha(G) = 2 \quad \text{and} \quad \alpha^*(G) = 5/2. \quad (76)$$

The source-measurement correlation term is given by

$$\text{Corr} = \sum_{e \in E(\Gamma_G)} q_e \sum_{m_e, s_e} \delta_{m_e, s_e} p(m_e, s_e | M_e, S_e) \quad (77)$$

for any choice of probability distribution $q \equiv \{q_e\}_{e \in E(\Gamma_G)}$. For simplicity, we will just take this probability distribution to be uniform, i.e., $q_e = \frac{1}{5}$ for all $e \in E(\Gamma_G)$. Note that the only extremal probabilistic model on Γ_G corresponding to an indeterministic assignment (in Λ_{ind}) assigns $\xi(v|\lambda) = \frac{1}{2}$ for all $v \in V(G)$. This means

$$\beta(\Gamma_G, q) = \frac{1}{2} \quad \forall q. \quad (78)$$

Figure 7: Geometric configuration of the vectors appearing in the KCBS construction [47].

The noncontextuality inequality of Eq. (71)

$$R \leq \alpha(G, w) + \frac{\alpha^*(G, w) - \alpha(G, w)}{p_0} \frac{1 - \text{Corr}}{1 - \beta(\Gamma_G, q)} \quad (79)$$

then becomes (in the KCBS scenario)

$$R \leq 2 + \frac{1/2}{p_0} \frac{1 - \text{Corr}}{1/2}, \quad (80)$$

or

$$R \leq 2 + \frac{1 - \text{Corr}}{p_0}. \quad (81)$$

Recall that the KCBS inequality [22, 47] reads $R \leq 2$ and it would be a valid noncontextuality inequality in our framework if and only if one can find measurements and preparations such that $\text{Corr} = 1$. In the standard KCBS construction [47] that violates the inequality $R \leq 2$, we have the five vertices in G (say v_i , $i \in \{1, 2, 3, 4, 5\}$, labelled cyclically) associated with five projectors $\Pi_i = |l_i\rangle\langle l_i|$, $i \in \{1, 2, 3, 4, 5\}$, on a qutrit Hilbert space, given by the vectors $|l_i\rangle = (\sin \theta \cos \phi_i, \sin \theta \sin \phi_i, \cos \theta)$, $\phi_i = \frac{4\pi i}{5}$, and $\cos \theta = \frac{1}{\sqrt{5}}$. The special source event $[s_{e_*} = 0|S_{e_*}]$ is associated with the quantum state $|\psi\rangle = (0, 0, 1)$, so that

$$R = \sum_{i=1}^5 |\langle l_i | \psi \rangle|^2 = \sqrt{5} > 2. \quad (82)$$

See Fig. 7 for a depiction of the geometric configuration of these vectors.

To turn this KCBS construction into an argument against noncontextuality in our approach, we need additional ingredients beyond the graph G . Firstly, for both the measurement events hypergraph Γ_G and the source events hypergraph Σ_G , we denote the hyperedges by e_i , $i \in \{1, 2, 3, 4, 5\}$. In Γ_G , the measurement events for the setting M_{e_i} are given by $\{[m_{e_i} = 0|M_{e_i}] = |l_i\rangle\langle l_i|, [m_{e_i} = 1|M_{e_i}] = \mathbb{I} - |l_i\rangle\langle l_i| -$

$|l_{i+1}\rangle\langle l_{i+1}|, [m_{e_i} = 2|M_{e_i}] = |l_{i+1}\rangle\langle l_{i+1}|\}$, where for $i = 5$, $i + 1 = 1$ (addition modulo 5). Similarly, in Σ_G , the source events corresponding to source setting S_{e_i} are given by $\{[s_{e_i} = 0|S_{e_i}] = |l_i\rangle\langle l_i|, [s_{e_i} = 2|S_{e_i}] = |l_{i+1}\rangle\langle l_{i+1}|, \text{ and } [s_{e_i} = 1|S_{e_i}] = \mathbb{I} - |l_i\rangle\langle l_i| - |l_{i+1}\rangle\langle l_{i+1}|\}$, where $p(s_{e_i} = b|S_{e_i}) = \frac{1}{3}$ for all $b \in \{0, 1, 2\}$. The special source setting S_{e_*} consists of source events $\{[s_{e_*} = 0|S_{e_*}] = |\psi\rangle\langle\psi|, [s_{e_*} = 1|S_{e_*}] = \frac{\mathbb{I} - |\psi\rangle\langle\psi|}{2}\}$, where $p(s_{e_*} = 0|S_{e_*}) = \frac{1}{3}$ and $p(s_{e_*} = 1|S_{e_*}) = \frac{2}{3}$. We thus have the operational equivalences we need between the source settings:

$$[\top|S_{e_\top}] \simeq [\top|S_{e'_\top}] = \frac{\mathbb{I}}{3}, \quad \forall e, e' \in E(\Sigma_G). \quad (83)$$

This choice of representation for Γ_G and Σ_G yields $p_0 = \frac{1}{3}$, $\text{Corr} = 1$, and $R([s_{e_*} = 0|S_{e_*}]) = \sqrt{5}$, so that the inequality

$$R \leq 2 + \frac{1 - \text{Corr}}{p_0} \quad (84)$$

is violated. However, note that this is an idealization (under which $\text{Corr} = 1$) and, typically, the source events and measurement events will not be perfectly correlated ($\text{Corr} < 1$) and the operational equivalences between the source settings need not correspond to the maximally mixed state. All that is required for a test of noncontextuality using this inequality is that the operational equivalences hold for *some* choice of preparations and measurements which need not be the same as that in the ideal KCBS construction.

To illustrate what happens when $\text{Corr} < 1$, we consider the effect of a depolarizing channel on the states and measurements in the ideal KCBS construction. The channel is given by

$$\mathcal{D}_r(\cdot) = r\mathbb{I}(\cdot) + (1 - r)\frac{1}{3}\mathbb{I}\text{Tr}(\cdot), \quad r \in [0, 1]. \quad (85)$$

The action of this channel – with parameter $r_1 \in [0, 1]$, say – on the pure states $\{|l_i\rangle\langle l_i|\}_{i=1}^5, |\psi\rangle\langle\psi|\}$ yields the noisy states given by

$$\mathcal{D}_{r_1}(|l_i\rangle\langle l_i|) = r_1|l_i\rangle\langle l_i| + (1 - r_1)\frac{\mathbb{I}}{3}, \quad \forall i \in [5], \quad (86)$$

$$\mathcal{D}_{r_1}(|\psi\rangle\langle\psi|) = r_1|\psi\rangle\langle\psi| + (1 - r_1)\frac{\mathbb{I}}{3}, \quad (87)$$

and the action of its adjoint – with parameter $r_2 \in [0, 1]$, say – on the ideal projectors, $\{|l_i\rangle\langle l_i|\}_{i=1}^5$, involved in the measurements correspondingly yields the POVM elements given by

$$\mathcal{D}_{r_2}^\dagger(|l_i\rangle\langle l_i|) = r_2|l_i\rangle\langle l_i| + (1 - r_2)\frac{\mathbb{I}}{3}, \quad \forall i \in [5]. \quad (88)$$

Hence, we are imagining a situation where the preparation procedures are affected by depolarizing noise with parameter r_1 and measurement procedures are affected by depolarizing noise with parameter r_2 , similar to the situation considered previously in Section

II of the Supplemental material of Ref. [12]. The operational equivalences required for our argument from preparation and measurement noncontextuality are satisfied by these noisy preparations and measurements. That is, in Γ_G , we can represent the measurement events for the setting M_{e_i} (where $i \in [5]$, $i + 1 = 1$, i.e., addition modulo 5) by

$$[m_{e_i} = 0|M_{e_i}] = \mathcal{D}_{r_2}^\dagger(|l_i\rangle\langle l_i|), \quad (89)$$

$$[m_{e_i} = 1|M_{e_i}] = \mathcal{D}_{r_2}^\dagger(\mathbb{I} - |l_i\rangle\langle l_i| - |l_{i+1}\rangle\langle l_{i+1}|), \quad (90)$$

$$[m_{e_i} = 2|M_{e_i}] = \mathcal{D}_{r_2}^\dagger(|l_{i+1}\rangle\langle l_{i+1}|). \quad (91)$$

It is easy to verify that these form elements of a valid POVM denoted by the measurement setting M_{e_i} and that the operational equivalences between the measurement events (represented by Γ_G) are indeed respected. On the other hand, in Σ_G , the source events corresponding to source setting S_{e_i} can be represented by

$$[s_{e_i} = 0|S_{e_i}] = \mathcal{D}_{r_1}(|l_i\rangle\langle l_i|), \quad (92)$$

$$[s_{e_i} = 1|S_{e_i}] = \mathcal{D}_{r_1}(\mathbb{I} - |l_i\rangle\langle l_i| - |l_{i+1}\rangle\langle l_{i+1}|) \quad (93)$$

$$[s_{e_i} = 2|S_{e_i}] = \mathcal{D}_{r_1}(|l_{i+1}\rangle\langle l_{i+1}|), \quad (94)$$

where $p(s_{e_i} = b|S_{e_i}) = \frac{1}{3}$ for all $b \in \{0, 1, 2\}$, while the source events for source setting S_{e_*} can be represented by

$$[s_{e_*} = 0|S_{e_*}] = \mathcal{D}_{r_1}(|\psi\rangle\langle\psi|), \quad (95)$$

$$[s_{e_*} = 1|S_{e_*}] = \mathcal{D}_{r_1}\left(\frac{\mathbb{I} - |\psi\rangle\langle\psi|}{2}\right), \quad (96)$$

where $p(s_{e_*} = 0|S_{e_*}) = \frac{1}{3}$ and $p(s_{e_*} = 1|S_{e_*}) = \frac{2}{3}$.

These satisfy the operational equivalences

$$[\top|S_{e_\top}] \simeq [\top|S_{e'_\top}] = \frac{\mathbb{I}}{3}, \quad \forall e, e' \in E(\Sigma_G). \quad (97)$$

We then have

$$\text{Corr} = \frac{1}{5} \sum_{e \in E(\Gamma_G)} \frac{1}{3} \sum_{b \in \{0,1,2\}} p(m_e = b|M_e, S_e, s_e = b). \quad (98)$$

Noting that for any qutrit pure state $|\phi\rangle$ and its corresponding projector $|\phi\rangle\langle\phi|$, each affected by depolarizing noise with parameters r_1 and r_2 , respectively, we have

$$\begin{aligned} & \text{Tr}(\mathcal{D}_{r_1}(|\phi\rangle\langle\phi|)\mathcal{D}_{r_2}^\dagger(|\phi\rangle\langle\phi|)) \\ &= \frac{1}{3} + \frac{2}{3}r_1r_2. \end{aligned} \quad (99)$$

Now, each term in the summation defining Corr, namely, $p(m_e = b|M_e, S_e, s_e = b)$, is obtained from a calculation of the type in Eq. (99). Hence, we have for each such term,

$$\begin{aligned} & p(m_e = b|M_e, S_e, s_e = b) \\ &= \frac{1}{3} + \frac{2}{3}r_1r_2, \end{aligned} \quad (100)$$

so that

$$\text{Corr} = \frac{1}{3} + \frac{2}{3}r_1r_2. \quad (101)$$

In the noiseless regime, i.e., $r_1 = r_2 = 1$, this reduces to the ideal KCBS scenario. On the other hand, we have

$$\begin{aligned} & R([s_{e_*} = 0|S_{e_*}]) \\ &= \sum_{v \in V(G)} p(v|S_{e_*}, s_{e_*} = 0) \\ &= \sum_{i=1}^5 \text{Tr}(\mathcal{D}_{r_2}^\dagger(|l_i\rangle\langle l_i|)\mathcal{D}_{r_1}(|\psi\rangle\langle\psi|)) \\ &= \sum_{i=1}^5 \left(r_1r_2|\langle l_i|\psi\rangle|^2 + \frac{r_2(1-r_1)}{3} + \frac{r_1(1-r_2)}{3} + \frac{(1-r_1)(1-r_2)}{3} \right) \\ &= r_1r_2 \sum_{i=1}^5 |\langle l_i|\psi\rangle|^2 + \frac{5}{3}(1-r_1r_2). \end{aligned} \quad (102)$$

Recall that violation of the noncontextuality inequality requires that

$$R > 2 + \frac{1 - \text{Corr}}{p_0}. \quad (103)$$

That is,

$$\begin{aligned} & r_1r_2 \sum_{i=1}^5 |\langle l_i|\psi\rangle|^2 + \frac{5}{3}(1-r_1r_2) \\ & > 2 + 3 \left(1 - \left(\frac{1}{3} + \frac{2}{3}r_1r_2 \right) \right). \end{aligned} \quad (104)$$

Given that $\sum_{i=1}^5 |\langle l_i | \psi \rangle|^2 = \sqrt{5}$, this becomes

$$r_1 r_2 \sqrt{5} + \frac{5}{3}(1 - r_1 r_2) > 2 + 2(1 - r_1 r_2). \quad (105)$$

Rewriting this, we obtain

$$r_1 r_2 > 1 - \frac{\sqrt{5} - 2}{\sqrt{5} + \frac{1}{3}} \approx 0.908, \quad (106)$$

that is, the noncontextuality inequality can be violated only when the depolarizing noise is below a certain threshold given by $r_1 r_2 > 0.908$. In terms of Corr, this requires $\text{Corr} > 0.939$. The noiseless case $r_1 = r_2 = 1$ takes us back to the $\text{Corr} = 1$ regime that we previously discussed.

6 Discussion

6.1 Measurement-measurement correlations vs. source-measurement correlations

Note that the usual Kochen-Specker experiment, as conceptualized in Refs. [21–23, 25], for example, involves only the quantity $R([s|S])$, representing correlations between various measurement events when all the measurements are implemented on a system prepared according to the same preparation procedure, denoted by the source event $[s|S]$. Thus, R represents measurement-measurement correlations on a system prepared according to a fixed choice of preparation procedure.

On the other hand, the experiment we have conceptualized in this paper involves, besides the quantity R , a quantity Corr representing source-measurement correlations, characterizing the quality of the measurements in terms of their response to corresponding preparations.

Our noncontextuality inequalities represent a trade-off relation that must hold between R and Corr in an operational theory that admits a noncontextual ontological model. Here we note that the first example of such a tradeoff relation, albeit only for the case of operational quantum theory with unsharp measurements, appeared in Ref. [39] as the Liang-Spekkens-Wiseman (LSW) inequality [40] which has been shown to be experimentally violated in Ref. [53].³⁶ And, indeed, the developments reported in Ref. [16] and the present paper have their origins in the idea of such a trade-off relation that first appeared in Ref. [39].

³⁶This experiment, however, is not in a position to make claims about contextuality without presuming the operational theory is quantum theory simply because the LSW inequality presumes operational quantum theory. The noncontextuality inequalities in this paper do not require the operational theory to be quantum theory and can therefore be experimentally tested using techniques from Refs. [13, 37, 54].

6.2 Can our noise-robust noncontextuality inequalities be saturated by a noncontextual ontological model?

A natural question concerns the tightness of these noncontextuality inequalities, i.e., can Eq. (72) be saturated by a noncontextual ontological model? This requires one to specify a noncontextual ontological model reproducing the operational equivalences between the measurement events and between the source settings, such that

$$\text{Corr} + p_0(1 - \beta(\Gamma_G, q)) \frac{R - \alpha(G, w)}{\alpha^*(G, w) - \alpha(G, w)} = 1. \quad (107)$$

The assumption of measurement noncontextuality is already implicit in our characterization of the response functions $\xi(m_e | M_e, \lambda)$, and for this reason it is, indeed, trivial to satisfy measurement noncontextuality while saturating these noncontextuality inequalities. Measurement noncontextuality, alone, in fact even allows a violation of the inequality (when no preparation noncontextuality is imposed), the extreme case being $R = \alpha^*(G, w)$ and $1 \geq \text{Corr} > 1 - p_0(1 - \beta(\Gamma_G, q))$. It's the assumption of preparation noncontextuality that is nontrivial to satisfy and we do not know if there exists a general construction of a noncontextual ontological model saturating our noncontextuality inequalities. We outline the general situation below.

6.2.1 The special case of facet-defining Bell-KS inequalities: Corr=1

If outcome determinism is presumed (as in traditional Bell-KS type treatments), then we know that there exists a necessary and sufficient set of Bell-KS inequalities (each corresponding to a particular choice of $R([s|S])$) that are satisfied by *any* operational theory admitting a KS-noncontextual ontological model. In particular, each such (facet) Bell-KS inequality can be saturated by KS-noncontextual ontological models that yield probabilities (from $\mathcal{G}(\Gamma_G)$) corresponding to the facet-defining Bell-KS inequality, i.e., which satisfy $R([s|S]) = \alpha(G, w)$ for such a Bell-KS inequality. Indeed, our noise-robust noncontextuality inequalities corresponding to these choices of $R([s|S])$ (i.e., facet-defining Bell-KS inequalities of the Bell-KS polytope which is given by the convex hull of points in $\mathcal{G}(\Gamma_G)|_{\text{det}}$) can always be saturated when $\text{Corr} = 1$, because in that case outcome determinism is justified by preparation noncontextuality (cf. Ref. [16]) and our inequalities are identical to the Bell-KS inequalities (saturated by $R = \alpha(G, w)$).

6.2.2 The general case: Corr < 1

Since we do not want to assume outcome determinism, nor necessarily the idealization of $\text{Corr} = 1$,

what is at stake here is the assumption of preparation noncontextuality. This assumption must be satisfied while saturating the noise-robust noncontextuality inequality in order for a measurement noncontextual ontological model to be universally noncontextual. Constructing such a noncontextual ontological model amounts to specifying the distributions $\mu(s_e|S_e, \lambda)$ and $\nu(\lambda)$ such that

$$\forall \lambda \in \Lambda : \mu(\lambda|S_e) = \nu(\lambda), \forall e \in E(\Sigma_G), \quad (108)$$

i.e., preparation noncontextuality holds, and we have (rewriting the saturation condition from Eq. (107))

$$\begin{aligned} & (\alpha^*(G, w) - \alpha(G, w))\text{Corr} + p_0(1 - \beta(\Gamma_G, q))R \\ &= (\alpha^*(G, w) - \alpha(G, w)) + p_0\alpha(G, w)(1 - \beta(\Gamma_G, q)), \end{aligned} \quad (109)$$

where

$$\text{Corr} = \sum_{s_{e_*}} p(s_{e_*}|S_{e_*}) \text{Corr}_{s_{e_*}}, \quad (110)$$

$$\text{Corr}_{s_{e_*}} = \sum_{\lambda \in \Lambda} \text{Corr}(\lambda) \mu(\lambda|S_{e_*}, s_{e_*}), \quad (111)$$

$$\begin{aligned} & \text{Corr}(\lambda) \\ & \equiv \sum_{e \in E(\Gamma_G)} q_e \sum_{m_e, s_e} \delta_{m_e, s_e} \xi(m_e|M_e, \lambda) \mu(s_e|S_e, \lambda), \end{aligned} \quad (112)$$

and

$$R = \sum_{\lambda \in \Lambda} R(\lambda) \mu(\lambda|S_{e_*}, s_{e_*} = 0). \quad (113)$$

Unfortunately, we do not have a general construction that can show this to be possible for any noise-robust noncontextuality inequality obtained according to the approach we have outlined. We therefore leave it as an open question whether such an inequality can (always?) be saturated by a noncontextual ontological model.

6.3 Can trivial POVMs ever violate these noncontextuality inequalities?

No.

Recall that a trivial POVM is defined as an assignment of positive operators $p(v)\mathbb{I}$ to the vertices of Γ_G , where $\mathbb{I}$ is the identity operator on some Hilbert space and $p : V(\Gamma_G) \rightarrow [0, 1]$, such that $\sum_{v \in e} p(v) = 1$ for all $e \in E(\Gamma_G)$, is a probabilistic model on Γ_G .

6.3.1 The case $p \in \mathcal{C}(\Gamma_G)$

Consider trivial POVMs corresponding to any KS-noncontextual probabilistic model, i.e., $p \in \mathcal{C}(\Gamma_G)$ is a convex mixture of deterministic vertices, $\mathcal{G}(\Gamma_G)|_{\text{det}}$, or equivalently, of ontic states in Λ_{det} . In other words, $\mathcal{C}(\Gamma_G) \equiv \text{ConvHull}(\mathcal{G}(\Gamma_G)|_{\text{det}})$, the convex hull of points in $\mathcal{G}(\Gamma_G)|_{\text{det}}$. The largest value Corr

can take in this case is less than or equal to 1. This means that the upper bound on R from our noncontextuality inequality, Eq. (71), will be greater than or equal to $\alpha(G, w)$, whereas we know that for a KS-noncontextual probabilistic model, $R \leq \alpha(G, w)$. Hence, there is no violation of our noncontextuality inequality for such trivial POVMs.

6.3.2 The case $p \in \text{ConvHull}(\mathcal{G}(\Gamma_G)|_{\text{ind}})$

Now consider trivial POVMs that correspond to the indeterministic vertices, $\mathcal{G}(\Gamma_G)|_{\text{ind}}$ (correspondingly, Λ_{ind}), or their convex mixtures. We know that for these trivial POVMs, $\text{Corr} \leq \beta(\Gamma_G, q)$. For any $R \leq \alpha^*(G, w)$ that is achieved by these trivial POVMs, our noncontextuality inequality reads

$$\text{Corr} \leq 1 - p_0(1 - \beta(\Gamma_G, q)) \frac{R - \alpha(G, w)}{\alpha^*(G, w) - \alpha(G, w)}, \quad (114)$$

A sufficient condition for this inequality to be satisfied is that

$$\beta(\Gamma_G, q) \leq 1 - p_0(1 - \beta(\Gamma_G, q)) \frac{R - \alpha(G, w)}{\alpha^*(G, w) - \alpha(G, w)}, \quad (115)$$

which reduces, for $R > \alpha(G, w)$, to

$$p_0 \leq \frac{\alpha^*(G, w) - \alpha(G, w)}{R - \alpha(G, w)}, \quad (116)$$

where the upper bound is greater than or equal to 1, since $\alpha(G, w) < R \leq \alpha^*(G, w)$. This is trivially satisfied since $p_0 \leq 1$.

For $R < \alpha(G, w)$, the sufficient condition of Eq. (115) is again trivially satisfied since it reduces to

$$p_0 \geq -\frac{\alpha^*(G, w) - \alpha(G, w)}{\alpha(G, w) - R}, \quad (117)$$

and we must anyway have $p_0 \geq 0$.

For $R = \alpha(G, w)$, the sufficient condition reduces to $\beta(\Gamma_G, q) \leq 1$, which is again trivially satisfied since $\beta(\Gamma_G, q) < 1$ by definition.

6.3.3 The general case $p \in \mathcal{G}(\Gamma_G)$

In general, a probabilistic model achieved by trivial POVMs can be in the convex hull of both deterministic (Λ_{det}) and indeterministic (Λ_{ind}) ontic states, with the total weight on deterministic ontic states denoted by $\text{Pr}(\Lambda_{\text{det}})$ and that on indeterministic ontic states by $\text{Pr}(\Lambda_{\text{ind}})$, so that $\text{Pr}(\Lambda_{\text{det}}) + \text{Pr}(\Lambda_{\text{ind}}) = 1$. We then have

$$\begin{aligned} \text{Corr} & \leq \text{Pr}(\Lambda_{\text{det}}) + \text{Pr}(\Lambda_{\text{ind}})\beta(\Gamma_G, q), \\ R & \leq \text{Pr}(\Lambda_{\text{det}})\alpha(G, w) + \text{Pr}(\Lambda_{\text{ind}})\alpha^*(G, w). \end{aligned} \quad (118)$$

A sufficient condition for satisfaction of the noncontextuality inequality is then

$$\begin{aligned} & 1 - \Pr(\Lambda_{\text{ind}})(1 - \beta(\Gamma_G, q)) \\ & \leq 1 - p_0(1 - \beta(\Gamma_G, q)) \frac{R - \alpha(G, w)}{\alpha^*(G, w) - \alpha(G, w)}, \end{aligned} \quad (119)$$

which becomes

$$p_0 \leq \frac{\alpha^*(G, w) - \alpha(G, w)}{R - \alpha(G, w)} \Pr(\Lambda_{\text{ind}}) \quad (120)$$

when $R > \alpha(G, w)$. Noting that

$$R \leq \alpha(G, w) + \Pr(\Lambda_{\text{ind}})(\alpha^*(G, w) - \alpha(G, w)),$$

we have

$$\Pr(\Lambda_{\text{ind}}) \geq \frac{R - \alpha(G, w)}{\alpha^*(G, w) - \alpha(G, w)}, \quad (121)$$

so that the sufficient condition for satisfaction of the noncontextuality inequality becomes $p_0 \leq 1$, which is trivially satisfied.

When $R = \alpha(G, w)$, the sufficient condition becomes $\beta(\Gamma_G, q) \leq 1$, which is again trivially satisfied.

Finally, when $R < \alpha(G, w)$, the sufficient condition becomes

$$p_0 \geq -\frac{\alpha^*(G, w) - \alpha(G, w)}{\alpha(G, w) - R} \Pr(\Lambda_{\text{ind}}), \quad (122)$$

which is again trivially satisfied since $p_0 \geq 0$.

Hence trivial POVMs cannot yield a violation of our noncontextuality inequalities. This is the sense in which trivial POVMs cannot lead to nonclassicality in our approach, unlike the case of traditional Kochen-Specker approaches [21–23, 25] applied to the case of POVMs [73]. To violate our noncontextuality inequalities, the POVMs must necessarily have some nontrivial projective component (that is not the identity operator or zero) but they *need not* be projectors. Further, we do not rely on restricting the notion of joint measurability [44] (cf. Section 2.4) to commutativity for POVMs. Taking joint measurability to be just commutativity is the approach adopted in, for example, Ref. [25]. We refer to Appendix A and Appendix C for more discussion on these issues, in particular Appendix C for the role of commutativity vs. joint measurability.

7 Conclusions

We have obtained a hypergraph framework for obtaining noise-robust noncontextuality inequalities corresponding to KS-colourable scenarios, suitably augmented with preparation procedures in the spirit of Spekkens contextuality [18]. The inequalities take the form of a noncontextual tradeoff between the three operational quantities Corr , R , and p_0 , cf. Eq. (72).

This framework leverages the graph invariants from the graph-theoretic framework of CSW for doing this, in addition to a new hypergraph invariant (Eq. (45)) that we call the *weighted max-predictability*. Our approach is general enough to be applicable to any situation involving noisy preparations and measurements that arises from a KS-colourable contextuality scenario.

We conclude with a list of open questions raised in this paper and other directions for future research:

1. *Characterizing structural Specker's principle from probabilistic models on a hypergraph Γ :*

Given that $\mathcal{CE}^1(\Gamma) = \mathcal{G}(\Gamma)$ for some Γ , is it the case that Γ must then necessarily satisfy structural Specker's principle, namely, that every clique in $O(\Gamma)$ is a subset of some hyperedge in Γ ? Or is it the case that there exists a hypergraph Γ' for which $\mathcal{CE}^1(\Gamma') = \mathcal{G}(\Gamma')$ but structural Specker's principle fails?

More generally, is there *any* characterization of a hypergraph satisfying structural Specker's principle entirely in terms of the probabilistic models on it?

As already pointed out earlier, this open question relates to the open Problem 7.2.3 of Ref. [23] of characterizing Γ for which $\mathcal{CE}^1(\Gamma) = \mathcal{G}(\Gamma)$. It is known that Γ representing bipartite Bell scenarios [55] satisfy the property $\mathcal{CE}^1(\Gamma) = \mathcal{G}(\Gamma)$ and we have provided a generic recipe for converting any Γ that does not satisfy structural Specker's principle to a Γ' that does satisfy it so that $\mathcal{CE}^1(\Gamma') = \mathcal{G}(\Gamma')$. The question is if there are any *other* Γ that also satisfy $\mathcal{CE}^1(\Gamma) = \mathcal{G}(\Gamma)$.

2. *Almost quantum theory:* We know that an almost quantum theory cannot satisfy Specker's principle [49] but it satisfies statistical Specker's principle (or consistent exclusivity). An open question that remains is:

Can an almost quantum theory satisfy structural Specker's principle?

If not, this would render the satisfaction of consistent exclusivity by an almost quantum theory unexplained by a natural structural feature of measurements in the theory, namely, the satisfaction of structural Specker's principle, i.e., almost quantum theory would not fall in the category of operational theories envisaged in Ref. [30].

3. *Conditions for saturating the noise-robust noncontextuality inequalities:*

As mentioned in Section 6.2, it is an open question whether the noise-robust noncontextuality inequalities of Eq. (72) based on our generalization of the CSW framework [22] can be saturated by a noncontextual ontological model.

More generally, the status of these noise-robust noncontextuality inequalities vis-à-vis the algorithmic approach of Ref. [17] for finding necessary and sufficient conditions for noncontextuality in a general prepare-and-measure scenario remains to be explored. One would suspect that the algorithmic approach of Ref. [17] when adapted to the kind of situation considered in this paper would yield nontrivial noncontextuality inequalities that aren't merely generalizations of the ones obtained in the CSW framework [22]. It would be interesting to investigate the full structure of this set of inequalities and compare it with the facet-defining Bell-KS inequalities of the CSW framework.

4. *Properties of the weighted max-predictability, $\beta(\Gamma_G, q)$:*

Since the crucial new hypergraph-theoretic ingredient in our inequalities is the weighted max-predictability, it would be interesting to understand properties of this hypergraph invariant on both counts: as a new mathematical object in its own right, one we haven't been able to find a reference to in the hypergraph theory literature, as well as an important parameter of a hypergraph relevant for noise-robustness of a noise-robust noncontextuality inequality. Indeed, as we point out in Footnote 34, identifying a distribution q (in the definition of Corr, Eq. (53)) that minimizes $\beta(\Gamma_G, q)$ for a given Γ_G would lead to better noise-robustness in the inequalities of Eqs. (70) or (71).

5. *Noise-robust applications of quantum protocols based on KS-contextuality:*

A general research direction is to construct noise-robust versions of applications that have previously been suggested for KS-contextuality. Our approach provides a recipe for doing this for any Bell-KS inequality appearing in such applications. Besides serving as a witness for strong nonclassicality [56] (i.e., Spekkens contextuality),³⁷ noise-robust versions of these applications can help benchmark the experiments in terms of the noise that can be tolerated while still witnessing nonclassicality. Examples of such applications include those from Refs. [58–63].

Acknowledgments

I would like to thank Andreas Winter for his comments on an earlier version of some of these ideas, Tobias Fritz for the ping-pong and the sing-song in which

³⁷As opposed to weak nonclassicality that can arise in epistemically restricted classical theories [57]. See also the talk at Ref. [56], 41:43 minutes, for a short discussion.

we often talked about hypergraphs, Rob Spekkens for the often argumentative – but always productive – conversations over lunch, and participants at the Contextuality conference (CCIOSA) at Perimeter Institute, during July 24 - 28, 2017, for very stimulating discussions that fed into the narrative of this paper. I would also like to thank David Schmid, Ana Belén Sainz, Elie Wolfe, and Tomáš Gonda for helping me better articulate the difference between structural vs. statistical readings of Specker's principle, and Eric Cavalcanti for comments on the manuscript. Theorem 1 owes its origin to a discussion with Tomáš Gonda. I would also like to thank anonymous referees for suggestions that immensely improved the presentation of these results. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development Canada, and by the Province of Ontario through the Ministry of Research, Innovation and Science.

A Status of KS-contextuality as an experimentally testable notion of nonclassicality for POVMs in quantum theory

The purpose of this section is to emphasize how the progression from KS-contextuality to Spekkens contextuality for KS-type contextuality experiments is a natural one rather than an *ad hoc* move from one framework to another. That is, Spekkens contextuality is not just another notion of nonclassicality that is incomparable with KS-contextuality, but is indeed intimately connected in its motivations to the limitations of KS-contextuality [18]. In particular, we will focus on the role of KS-contextuality with respect to POVMs and why allowing arbitrary POVMs poses a difficulty for KS-contextuality as a notion of nonclassicality that is experimentally testable, i.e., a notion that applies to noisy measurements (POVMs) typically implemented in a laboratory experiment.³⁸ While one may be tempted to reject this premise for assessing the suitability of KS-contextuality as a notion of nonclassicality – claiming instead that KS-contextuality was never meant for POVMs and applies only to “purified” experiments (namely, ones with only PVMs and pure states) – the reasons for doing so are rooted in the literature on KS-contextuality where POVMs have indeed been considered and (at least) two kinds of conclusions drawn: one, that there exists a Kochen-Specker contradiction for POVMs, even on a qubit, so KS-contextuality for POVMs is interesting [64] and two, that allowing arbitrary POVMs in assessing nonclassicality would make the research program of identifying

³⁸And how a rather compelling way to arrive at a notion that is experimentally testable is Spekkens contextuality.

ing device-independent principles for quantum correlations in KS-contextuality experiments ill-defined, so quantum correlations allowing arbitrary POVMs are “pathological” [73]. We will look at these arguments in turn and use the latter, in particular, to segue into our motivations for the framework proposed in this paper.

A.1 Limitations of KS-contextuality vis-à-vis POVMs

A.1.1 KS-contextuality for POVMs in the literature

The first paper that applied KS-contextuality to the case of POVMs was by Cabello [64] where a KS-uncolourability argument for POVMs on a single qubit was proposed. This was motivated by the Gleason-type derivation of the Born rule starting with the structure of POVMs due to Busch [66] and Caves *et al.* [67], analogous to the case of the Kochen-Specker theorem [19] which can be seen as motivated by Gleason’s theorem [68]. Insofar as there exists a Gleason-type theorem for POVMs [66, 67], one could motivate KS-contextuality as a reasonable notion of nonclassicality for POVMs, as was presumably the case in Ref. [64]. The role of this notion of nonclassicality is then just to argue – using a finite set of POVM elements – that no KS-noncontextual assignment of outcomes is possible for certain finite sets of POVMs in quantum theory. Should we, however, assume that it is reasonable to demand deterministic assignment of outcomes to POVM elements in an ontological model, just as we do for PVM elements? The argument of Ref. [64] was later criticized on various counts [18, 28, 65] and we refer the reader to Ref. [28] for criticisms pertinent to this paper, namely, that *outcome determinism for all unsharp measurements* (ODUM in Ref. [28]) in quantum theory is untenable.³⁹ Other works in the literature where KS-contextuality for POVMs has been explored include Refs. [69–72].

Besides, doubts about the experimental testability of the KS theorem were raised in the late ‘90s in a series of papers by Meyer, Clifton, and Kent [74–76]. A review can be found in Ref. [77]. These doubts were premised on the idea that the set of KS-colourable projectors (or PVMs) on any given Hilbert space is dense in the set of all projectors (or PVMs) on that Hilbert space. That is, for any given set of PVMs yielding a KS contradiction, it is always possible to find PVMs which are arbitrarily “close” to the PVMs required for a KS contradiction (for any finite precision) but which do not themselves lead to a KS contradiction. The property of denseness of KS-colourable

sets of measurements in the set of all measurements in fact extends to even the most general case when the measurements are POVMs on any Hilbert space. So, even a KS contradiction for POVMs (such as the one in Ref. [64]) falls prey to the Meyer-Clifton-Kent argument [77]. As Ref. [77] notes:

Dealing with projective measurements is arguably not enough. One quite popular view of quantum theory holds that a correct version of the measurement rules would take POV measurements as fundamental, with projective measurements either as special cases or as idealisations which are never precisely realised in practice. In order to define an NCHV theory catering for this line of thought, Kent constructed a KS-colourable dense set of positive operators in a complex Hilbert space of arbitrary dimension, with the feature that it gives rise to a dense set of POV decompositions of the identity (Kent, 1999). Clifton and Kent constructed a dense set of positive operators in complex Hilbert space of arbitrary dimension with the special feature that no positive operator in the set belongs to more than one decomposition of the identity (Clifton & Kent, 2000). Again, the resulting set of POV decompositions is dense, and the special feature ensures that one can average over hidden states to recover quantum predictions.

Hence, in any finite precision experiment it would be impossible to test the Kochen-Specker theorem, i.e., such an experimental test would require an infinitely precise measurement and measurements in a real-world laboratory are never infinitely precise. Although there was a lively debate along these lines (see the references in [77]), the resolutions that were proposed all involved modifying the notion of KS-noncontextuality by adding auxiliary assumptions that seek to exclude the Meyer-Clifton-Kent type arguments. A recent attempt in this direction can be found in Ref. [78] where a notion of “ontological faithfulness” is proposed. As such, it was already recognized – for reasons independent of Spekkens contextuality [18] – that the notion of KS-noncontextuality needs to be revised if one is to make it experimentally testable.⁴⁰ What Spekkens brought to the fore [18], besides generalizing the notion of contextuality to all experimental procedures rather than measurements alone, was the idea that an experimental test of noncontextuality should not rely on inequalities that presume outcome determinism, just as a test of local causality does not require the assumption of outcome

³⁹Ref. [28] is also a good resource for a detailed analysis of arguments concerning dilations of POVMs, which we will not get into here. Besides, it also provides a principled recipe for assigning response functions to POVMs.

⁴⁰Of course, this takes nothing away from the importance of the Kochen-Specker theorem [19] as a no-go theorem concerning the logical structure of quantum theory and the constraints it places on the ontological models possible for the theory.

determinism. Indeed, the assumption of outcome determinism for sharp measurements in quantum theory is *derived* in the Spekkens framework from the assumption of preparation noncontextuality rather than being assumed independently.

We will now consider the more modern approach to KS-contextuality along the lines of the frameworks in Refs.[22, 23, 25] to segue into our framework for Spekkens contextuality which we develop in this paper.

A.1.2 Classifying probabilistic models: restriction of quantum models to PVMs

Research on KS-contextuality took a different turn with the advent of the graph-theoretic framework of Cabello, Severini and Winter in 2010 [21] (revised slightly in 2014 [22]), the sheaf-theoretic framework of Abramsky and Brandenburger in 2011 [25], and the hypergraph based formalism of Acín, Fritz, Leverrier, and Sainz in 2012 [23]. The unifying theme of these contributions was that they took the key mathematical idea underlying KS-noncontextuality and Bell-locality — namely, that both are instances of the classical marginal problem [26, 32, 33] — and built frameworks that sought to distinguish between classical theories (namely, those admitting KS-noncontextual ontological models), quantum theory, and post-quantum general probabilistic theories by classifying their empirical predictions relative to a Kochen-Specker experiment into these categories. All these frameworks, motivated by the device-independence paradigm, eschewed the erstwhile restriction of the notion of KS-noncontextuality to quantum theory and sought to make their analysis theory-independent, relying only on empirical predictions relative to a KS experiment to classify theories. They separated the assumption of KS-noncontextuality from the operational theory — namely, quantum theory — to which it was originally meant to apply, allowing arbitrary operational theories in their analysis. However, there was a key distinction between Bell scenarios and KS-contextuality scenarios that was lost in this formal unification: namely, that while the definition of a quantum probabilistic model in a Bell scenario need not be restricted to (local) PVMs (and arbitrary local POVMs can be allowed without changing the set of quantum models), the same is not true of a KS-contextuality scenario. Indeed, as Henson and Sainz note in their work [73],⁴¹ reflecting on the question of allowing arbitrary POVMs in the definition of a quantum probabilistic model:

...if we allow general POVMs rather than projective measurements then no principle

⁴¹Proposing a principle bounding the KS-contextuality possible in quantum theory, namely, “Macroscopic Noncontextuality”.

that places a non-trivial restriction on correlations will be respected. Thus, this kind of “quantum model” is clearly pathological.

One way to motivate the present work is as a response to the pathology that Henson and Sainz allude to: that trivial POVMs can realize any probabilistic model, hence allowing arbitrary POVMs makes the problem of finding principles to identify quantum correlations in KS-contextuality scenarios trivial, i.e., all probabilistic models are quantum and there is nothing to be learnt about post-quantum probabilistic models. This is because any set of probabilities satisfying the “no-disturbance” or “no-signalling” condition (of which the $E1$ correlations of CSW [22] are a subset, in general) can be achieved by (trivial) POVMs by simply multiplying an identity operator with every probability in such an assignment of probabilities.⁴² By the lights of KS-noncontextuality as one’s notion of classicality, then, trivial POVMs saturating the general probabilistic bound on the correlations would seem to be maximally nonclassical (i.e., maximally KS-contextual). To avoid such “pathological” quantum models, they restrict the definition of a quantum model to allow only projective measurements. Indeed, with recent work on a sensible notion of “sharp” measurement in a general probabilistic theory [30, 31], an appeal to the “fundamental sharpness” of all measurements (see, e.g., [29]) is made to restrict attention to sharp measurements in both quantum theory and general probabilistic theories.

On the other hand, the approach in this paper is different. In particular, we want our approach to capture the intuition that trivial POVMs are “classical” (and not pathological), so we must go beyond KS-noncontextuality. A simple operational sense in which trivial POVMs are “classical” is that they reveal nothing about the quantum state on which they are measured, being incapable of distinguishing any pair of states whatsoever.⁴³ The correlations (denoted by $R([s|S])$) usually examined in a KS-contextuality experiment do not allow such experiments to witness the “triviality” of trivial POVMs, i.e., the fact that they correspond to a fixed probability distribution that doesn’t vary even as the choice of preparation is varied. Moreover, since all nonprojective measurements are excluded *by fiat* in traditional Kochen-Specker type approaches [22, 23] for reasons alluded to by Henson and Sainz [73], one loses out on the potential to explore the possibilities that nontrivial and nonprojective measurements offer with respect to con-

⁴²Trivial POVMs are, therefore, trivial resolutions of the identity, where every POVM element is proportional to identity, i.e., $\{a\mathbb{I}\}_a$, such that $a \in [0, 1]$ and $\sum_a a = 1$.

⁴³Indeed, any trivial POVM can be realized in the following operational manner: take the quantum system prepared in some state, throw it in the garbage, and then sample from the classical probability distribution corresponding to the trivial POVM.

textuality.⁴⁴ Our approach, therefore, is to allow arbitrary POVMs when considering probabilistic models arising from quantum theory (and not restricting to any notion of “sharp measurements” in general probabilistic theories) but examine more quantities than are examined in traditional approaches, i.e., besides the quantity R typical in a KS-contextuality scenario, we invoke the quantity Corr to account for noise in the measurements.

If one restricts attention to operational theories that can always achieve $\text{Corr} = 1$ for any KS-contextuality scenario, then the usual classification of probabilistic models following Refs. [22, 23] holds (Eq. (72)). What is of interest in our framework, however, is the tradeoff between R and Corr : how large can both R and Corr be in an operational theory? (See Eq. (72).)

A.2 Robustness of Bell nonlocality vis-à-vis POVMs

Note that whenever we refer to “Bell-KS” functionals or inequalities for Kochen-Specker type experiments, we are *not* thinking of experiments that are Bell experiments [4, 5, 7, 9–11], which have spacelike separation between multiple parties, each performing local measurements on a shared multipartite preparation. For the case of Bell experiments, trivial local POVMs assigned to each party in a Bell experiment do not lead to Bell violations for a simple reason: the trivial POVMs for each party are all compatible with each other, thereby admitting a joint probability distribution over their outcomes for each party; taking a product of these local joint probability distributions (one for each party) results in a joint distribution over all measurements of all parties, hence satisfying Bell inequalities. The fact that the POVMs are trivial ensures that the Bell inequalities are satisfied regardless of the choice of shared quantum state. On the other hand, forgetting the constraint of *local* POVMs, there always exist global trivial POVMs that can violate Bell inequalities: e.g., just take the Popescu-Rohrlich (PR) box distribution [43], and multiply an identity operator (on the joint Hilbert space of Alice and Bob) with each probability in the PR-box; this results in four trivial POVMs, defined over the joint Hilbert space, that together violate the CHSH inequality maximally. But, of course, this violation is uninteresting because it doesn’t obey the *locality* constraint on the measurements in a Bell experiment. This is mathematically reflected in the fact that the PR-box distribution cannot be written as a convex mixture of product distributions, one for each party, hence the corresponding trivial POVM cannot be un-

derstood in terms of trivial *local* POVMs. Hence, it is the *locality* of the trivial POVMs in a Bell experiment that prevents them from violating a Bell inequality and renders them non-pathological, unlike in the case of KS-contextuality. The fact that they are “trivial” in the sense of being unable to distinguish two quantum states plays a role in the sense that, *regardless* of the shared quantum state, these POVMs yield fixed distributions over the measurement outcomes, thus always allowing the construction of a fixed (that is, independent of the quantum state) global joint probability distribution over all measurements in a Bell scenario. Since there are no such locality constraints on the form of the POVM elements in a Kochen-Specker experiment, they can easily violate any KS-noncontextuality inequality, e.g., the two-party CHSH experiment considered as a Kochen-Specker experiment with four observables in a 4-cycle where adjacent pairs are jointly measurable allows for trivial POVMs (like the PR-box trivial POVM above) violating the CHSH-type Bell-KS inequality in this scenario maximally. By the lights of KS-noncontextuality, this violation would indicate the maximum possible KS-contextuality with respect to this CHSH-type inequality.⁴⁵ For all these reasons, our discussion of KS-noncontextuality as a notion of classicality — in an experiment with no *locality* constraints on the measurements — does not extend to the case of Bell-locality (or local causality) as a notion of classicality in a Bell experiment, where the experiment *must* respect *locality* constraints on the measurements for a Bell inequality violation to be meaningful.

The unification of Bell nonlocality and KS-contextuality à la Refs. [22, 23, 25] forces a certain dichotomy in these approaches: while in Bell scenarios, one *need not* restrict to any notion of a “sharp” measurement in the definition of probabilistic models (and thus claim “theory independence”), in Kochen-Specker scenarios, one must make some statement about the nature of the measurements (concerning their presumed sharpness [29], or that their joint measurability [42, 44] is restricted to commutativity [25]), rendering any putative “theory independence” claim (on a level at par with Bell nonlocality) unfounded.⁴⁶

⁴⁵See Appendix C for more discussion.

⁴⁶See Ref. [27] for how this lack of locality of measurements in a Kochen-Specker type experiment translates, at the ontological level, to the unreasonableness of assuming factorizability in the ontological model; this factorizability (or the stronger condition of outcome determinism) is invoked to justify the resulting derivation of Bell-KS inequalities as constraints from a classical marginal problem.

⁴⁴All trivial POVMs are nonprojective, but not all nonprojective POVMs are trivial. Indeed, see Refs. [39–42] for examples of generalized contextuality [18] with nonprojective measurements, albeit assuming operational quantum theory.

B Ontological models without respecting coarse-graining relations

Here we will construct explicit examples where the coarse-graining relations are not respected in an ontological model, in contrast to the requirement on the representation of coarse-grainings that we invoked in Section II.C of the main text. The goal is to emphasize that the requirement of Section II.C is *necessary* not only for the treatment of Spekkens contextuality but also for Kochen-Specker contextuality. Below, we first demonstrate how a “KS-noncontextual” model can be constructed for any scenario that proves the KS theorem by using the example of the KCBS setup [47]. We then proceed to demonstrate how a “preparation and measurement noncontextual” model can be constructed in a similar way when considering generalized noncontextuality [18].

B.1 How to construct a “KS-noncontextual” ontological model of the KCBS experiment [47] without coarse-graining relations

Here we have that $\mathbb{M}$ contains at least the following measurement settings: $\{M_i\}_{i=1}^5$, each with three possible outcomes, $m_i \in \{0, 1, 2\}$. The measurement events for each measurement setting M_i can be coarse-grained into two different ways, defining new measurement settings M'_i (with outcomes $m'_i \in \{0, \bar{0}\}$) and M''_i (with outcomes $m''_i \in \{2, \bar{2}\}$), where the coarse-graining relations are given by

$$[0|M'_i] \equiv [0|M_i], \quad (123)$$

$$[\bar{0}|M'_i] \equiv [1|M_i] + [2|M_i], \quad (124)$$

$$[2|M''_i] \equiv [2|M_i], \quad (125)$$

$$[\bar{2}|M''_i] \equiv [0|M_i] + [1|M_i]. \quad (126)$$

In the operational theory, these coarse-graining relations are respected, i.e., for all $[s|S]$, $s \in V_S$, $S \in \mathbb{S}$,

$$p(0, s|M'_i, S) \equiv p(0, s|M_i, S), \quad (127)$$

$$p(\bar{0}, s|M'_i, S) \equiv p(1, s|M_i, S) + p(2, s|M_i, S), \quad (128)$$

$$p(2, s|M''_i, S) \equiv p(2, s|M_i, S), \quad (129)$$

$$p(\bar{2}, s|M''_i, S) \equiv p(0, s|M_i, S) + p(1, s|M_i, S). \quad (130)$$

However, we do not require that these relations be respected in an ontological model. Now, the KCBS argument requires the following operational equivalences,

$$[2|M''_i] \simeq [0|M'_{i+1}], \quad (131)$$

for all $i \in \{1, 2, 3, 4, 5\}$, where addition is modulo 5, so that $i + 1 = 1$ for $i = 5$. A KS-noncontextual ontological model for this experiment requires that

$$\xi(2|M''_i, \lambda) = \xi(0|M'_{i+1}, \lambda) \in \{0, 1\}, \quad \forall \lambda \in \Lambda. \quad (132)$$

Constructing such a model requires one to specify response functions for the measurements

$\{M_i, M'_i, M''_i\}_{i=1}^5$. However, since there are no constraints from coarse-graining relations on these response functions, there is no obstruction to the construction of a “KS-noncontextual model” of this type for any set of operational statistics. In particular, since we do not require that $\forall \lambda \in \Lambda : \xi(0|M'_{i+1}, \lambda) \equiv \xi(0|M_{i+1}, \lambda)$, nor that $\forall \lambda \in \Lambda : \xi(2|M''_i, \lambda) \equiv \xi(2|M_i, \lambda)$, we can assign arbitrary response functions to $\{M'_i, M''_i\}_{i=1}^5$, subject only to the condition from KS-noncontextuality that $\forall \lambda \in \Lambda : \xi(2|M''_i, \lambda) = \xi(0|M'_{i+1}, \lambda) \in \{0, 1\}$.⁴⁷ Note that, because coarse-graining relations are not respected, this does not imply that $\forall \lambda \in \Lambda : \xi(2|M_i, \lambda) = \xi(0|M_{i+1}, \lambda) \in \{0, 1\}$, which is the usual constraint we would have presumed from KS-noncontextuality when coarse-graining relations are respected in the ontological model. In the absence of any such constraints on the response functions for $\{M_i\}_{i=1}^5$, one can always reproduce their operational statistics, in particular the operational equivalences of the type $[2|M_i] \simeq [0|M_{i+1}]$, which follow from Eqs. (123), (125), and (131).

B.2 How to construct a “preparation and measurement noncontextual” ontological model without coarse-graining relations

Just as for measurements in the case of KS-noncontextuality, abandoning the coarse-graining relations for preparations in the case of generalized noncontextuality [18] makes possible the existence of a “preparation and measurement noncontextual” ontological model for any set of operational statistics. For the kinds of proofs of contextuality relevant to this article, the relevant notion of coarse-graining is that of complete coarse-graining: that is, consider two source settings S and S' with (respective) source events $\{[s|S]\}_{s \in V_S}$ and $\{[s'|S']\}_{s' \in V_{S'}}$, that can be completely coarse-grained to yield the operational equivalence $[\top|S_\top] \simeq [\top|S'_\top]$, cf. Eq. (18). In the operational description, where we assume the coarse-graining relation is respected, this is represented by

$$\forall [m|M], m \in V_M, M \in \mathbb{M} : \sum_s p(m, s|M, S) = \sum_{s'} p(m, s'|M, S'). \quad (133)$$

In the ontological description, however, we do *not* impose the coarse-graining relations $\mu(\lambda, \top|S_\top) \equiv \sum_s \mu(\lambda, s|S)$ and $\mu(\lambda, \top|S'_\top) \equiv \sum_{s'} \mu(\lambda, s'|S')$, which makes it trivial to write down probability distributions $\mu(\lambda, \top|S_\top)$ and $\mu(\lambda, \top|S'_\top)$ such that $\mu(\lambda, \top|S_\top) = \mu(\lambda, \top|S'_\top)$ (as required by preparation noncontextuality applied to $[\top|S_\top] \simeq [\top|S'_\top]$) but where we do not require that $\sum_s \mu(\lambda, s|S) =$

⁴⁷This “KS-noncontextual” ontological model will thus reproduce operational equivalences of the type $[2|M''_i] \simeq [0|M'_{i+1}]$ (cf. Eq. (131)).

$\sum_s \mu(\lambda, s|S)$ (which is not required by preparation noncontextuality). Note how the refusal to respect the coarse-graining relations, i.e., identifying $\mu(\lambda, \top|S_\top)$ with $\sum_s \mu(\lambda, s|S)$ and $\mu(\lambda, \top|S'_\top)$ with $\sum_{s'} \mu(\lambda, s'|S')$, lifts the constraint from preparation noncontextuality that would have been in place if the coarse-graining relations were respected. The same refusal for the case of measurements lifts any constraints (just as in the case of KS-noncontextuality above) from measurement noncontextuality on the ontological model. It thus becomes trivial to construct a “preparation and measurement noncontextual” ontological model without coarse-graining relations.

C Trivial POVMs

C.1 Bell-CHSH scenario

We have the Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$ for Alice ($\mathcal{H}_A$) and Bob ($\mathcal{H}_B$). Consider four binary-outcome POVMs, $\{A^{(0)}, A^{(1)}, B^{(0)}, B^{(1)}\}$, where

$$\begin{aligned} A^{(0)} &\equiv \{A_0^{(0)}, A_1^{(0)}\}, \\ A^{(1)} &\equiv \{A_0^{(1)}, A_1^{(1)}\}, \\ B^{(0)} &\equiv \{B_0^{(0)}, B_1^{(0)}\}, \\ B^{(1)} &\equiv \{B_0^{(1)}, B_1^{(1)}\}, \end{aligned} \quad (134)$$

$0 \leq A_0^{(0)}, A_0^{(1)} \leq \mathbb{I}_{\mathcal{H}_A}$, $0 \leq B_0^{(0)}, B_0^{(1)} \leq \mathbb{I}_{\mathcal{H}_B}$, $A_0^{(0)} + A_1^{(0)} = A_0^{(1)} + A_1^{(1)} = \mathbb{I}_{\mathcal{H}_A}$, and $B_0^{(0)} + B_1^{(0)} = B_0^{(1)} + B_1^{(1)} = \mathbb{I}_{\mathcal{H}_B}$. The quantum probability, given a shared quantum state ρ_{AB} defined on $\mathcal{H}_A \otimes \mathcal{H}_B$, is given by

$$p(a, b|x, y) = \text{Tr}(\rho_{AB} A_a^{(x)} \otimes B_b^{(y)}), \quad (135)$$

for $a, b, x, y \in \{0, 1\}$. Here $A^{(x)} \otimes \mathbb{I}_{\mathcal{H}_B}$ is jointly measurable with $\mathbb{I}_{\mathcal{H}_A} \otimes B^{(y)}$, just because of the commutativity of their respective POVM elements. The joint observable being measured is $A^{(x)} \otimes B^{(y)}$. Now, consider the case when all the POVM elements are trivial, i.e., $A_a^{(x)} = q_a^{(x)} \mathbb{I}_{\mathcal{H}_A}$ and $B_b^{(y)} = r_b^{(y)} \mathbb{I}_{\mathcal{H}_B}$, for some $q_a^{(x)}, r_b^{(y)} \in [0, 1]$ for all $a, b, x, y \in \{0, 1\}$. We then have

$$p(a, b|x, y) = q_a^{(x)} r_b^{(y)}, \forall a, b, x, y \in \{0, 1\}. \quad (136)$$

A global joint probability distribution which reproduces the above as marginals is simply given by their product:

$$p(a^{(0)}, a^{(1)}, b^{(0)}, b^{(1)}) \equiv q_{a^{(0)}}^{(0)} q_{a^{(1)}}^{(1)} r_{b^{(0)}}^{(0)} r_{b^{(1)}}^{(1)}. \quad (137)$$

Hence, trivial POVMs never violate any Bell-CHSH inequality for this scenario.

C.2 CHSH-type contextuality scenario: 4-cycle

We now consider the Bell-CHSH scenario without the constraint of spacelike separation. What the lack of

spacelike separation means from the quantum perspective is that one no longer needs to model this spacelike separation by requiring a tensor product structure, or (more generally) by requiring the commutativity of the observables that are jointly measured [25, 79, 80]. That is, there is no physical justification for imposing the tensor product structure or the commutativity of jointly measured observables.⁴⁸

Thus, we have the Hilbert space $\mathcal{H}$ and we consider four binary-outcome POVMs, $\{A^{(0)}, A^{(1)}, B^{(0)}, B^{(1)}\}$, on $\mathcal{H}$, where

$$\begin{aligned} A^{(0)} &\equiv \{A_0^{(0)}, A_1^{(0)}\}, \\ A^{(1)} &\equiv \{A_0^{(1)}, A_1^{(1)}\}, \\ B^{(0)} &\equiv \{B_0^{(0)}, B_1^{(0)}\}, \\ B^{(1)} &\equiv \{B_0^{(1)}, B_1^{(1)}\}, \end{aligned} \quad (138)$$

$0 \leq A_0^{(0)}, A_0^{(1)}, B_0^{(0)}, B_0^{(1)} \leq \mathbb{I}_{\mathcal{H}}$, $A_0^{(0)} + A_1^{(0)} = A_0^{(1)} + A_1^{(1)} = B_0^{(0)} + B_1^{(0)} = B_0^{(1)} + B_1^{(1)} = \mathbb{I}_{\mathcal{H}}$. Further, the following sets of POVMs are jointly measurable: $\{A^{(0)}, B^{(0)}\}$, $\{A^{(0)}, B^{(1)}\}$, $\{A^{(1)}, B^{(0)}\}$, $\{A^{(1)}, B^{(1)}\}$. The most general joint observable for a pair of compatible POVMs $\{A^{(x)}, B^{(y)}\}$ is given by a POVM $G^{(xy)} \equiv \{G_{00}^{(xy)}, G_{01}^{(xy)}, G_{10}^{(xy)}, G_{11}^{(xy)}\}$ (that isn't necessarily unique [42]) such that: $G_{00}^{(xy)} + G_{01}^{(xy)} = A_0^{(x)}$, $G_{01}^{(xy)} + G_{11}^{(xy)} = A_1^{(x)}$, $G_{00}^{(xy)} + G_{10}^{(xy)} = B_0^{(y)}$, $G_{10}^{(xy)} + G_{11}^{(xy)} = B_1^{(y)}$. In particular, if (and only if) the POVMs $A^{(x)}$ and $B^{(y)}$ commute, we can construct the joint POVM as a product: $G_{ab}^{(xy)} = A_a^{(x)} B_b^{(y)}$ for all $a, b, x, y \in \{0, 1\}$. In the absence of such commutativity, the joint POVM cannot be written as a product.

The quantum probability, given a quantum state ρ on $\mathcal{H}$, is given by

$$p(a, b|x, y) = \text{Tr}(\rho G_{ab}^{(xy)}), \quad (139)$$

for $a, b, x, y \in \{0, 1\}$. Note that this probability depends on the joint measurement $G^{(xy)}$ implementing $A^{(x)}$ and $B^{(y)}$ together, and that, in general, there may be multiple choices of $G^{(xy)}$ possible. This is easy to see since there is one undetermined positive operator in the joint measurement that is not fixed by $A^{(x)}$ or $B^{(y)}$, i.e., we can write the POVM elements of $G^{(xy)}$ as: $G_{01}^{(xy)} = A_0^{(x)} - G_{00}^{(xy)}$, $G_{10}^{(xy)} = B_0^{(y)} - G_{00}^{(xy)}$,

⁴⁸On the other hand, what this lack of spacelike separation means from the perspective of an ontological model is that one no longer has a justification for assuming factorizability [25] and, consequently, the generalization of Fine's theorem [26] fails to prove that there is no loss of generality in assuming outcome determinism in discussions of KS-contextuality (unlike the case of Bell scenarios, where factorizability is justified by spacelike separation); there is a definite loss of generality, in that measurement noncontextual and outcome-indeterministic ontological models that are non-factorizable are not empirically equivalent to measurement noncontextual and outcome-deterministic (or KS-noncontextual) ontological models. See Ref. [27] for a discussion of this aspect.

$G_{11}^{(xy)} = \mathbb{I} - A_0^{(x)} - B_0^{(y)} + G_{00}^{(xy)}$, where $G_{00}^{(xy)}$ is a positive semidefinite operator satisfying $A_0^{(x)} + B_0^{(y)} - \mathbb{I}_{\mathcal{H}} \leq G_{00}^{(xy)} \leq A_0^{(x)}, B_0^{(y)}$. Here $G_{00}^{(xy)}$ represents the freedom in the choice of how the joint measurement might be implemented within quantum theory. This freedom reflects the fact that since the jointly measured observables are no longer spacelike separated, it is possible to introduce correlations between them that are stronger than what is allowed in the corresponding Bell scenario in quantum theory. The strength of these correlations is only limited by the constraints on $G_{00}^{(xy)}$ imposed by the marginal observables $A^{(x)}$ and $B^{(y)}$. This is in contrast to the case where $A^{(x)}$ and $B^{(y)}$ are spacelike separated observables and the *only* choice of joint POVM consistent with spacelike separation is fixed by $G_{00}^{(xy)} = A_0^{(x)} B_0^{(y)}$, i.e., the strength of correlations between $A^{(x)}$ and $B^{(y)}$ is fixed entirely by them and there is no freedom in choosing $G^{(xy)}$.

Thus, we have that $A^{(x)}$ is jointly measurable with $B^{(y)}$ and $G^{(xy)}$ denotes a joint POVM of $A^{(x)}$ and $B^{(y)}$. Now, consider the case when all the POVM elements are trivial, i.e., $A_a^{(x)} = q_a^{(x)} \mathbb{I}_{\mathcal{H}}$ and $B_b^{(y)} = r_b^{(y)} \mathbb{I}_{\mathcal{H}}$, for some $q_a^{(x)}, r_b^{(y)} \in [0, 1]$ for all $a, b, x, y \in \{0, 1\}$.

In particular, consider the case where $q_a^{(x)} = r_b^{(y)} = \frac{1}{2}$ for all $a, b, x, y \in \{0, 1\}$. A possible joint POVM for these trivial POVMs is then the product POVM:

$$G_{ab}^{(xy)} = A_a^{(x)} B_b^{(y)} = \frac{1}{4} \mathbb{I}_{\mathcal{H}}. \quad (140)$$

If one restricted joint measurability of $A^{(x)}$ and $B^{(y)}$ to just *commutativity* — a sufficient but not necessary condition for joint measurability⁴⁹ [44] — we would take the above choice of the product POVM as a “natural” one. Being a product of trivial POVMs, this choice will never lead to a violation of the CHSH-type inequality for this scenario. Indeed, the structure of a Bell scenario — requiring the decomposition of the Hilbert space as $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$ (*tensor product paradigm*), or more generally, imposing the commutativity requirement $[A_a^{(x)}, B_b^{(y)}] = 0$ (*commutativity paradigm*) — is such that the only possible choice of joint measurement that can be implemented by spacelike separated parties is the one that corresponds to the product POVM, given by operators $G_{ab}^{(xy)} = A_a^{(x)} B_b^{(y)}$.

However, this is not the only allowed joint measurement for these trivial POVMs, particularly when there is no locality constraint on the measurements from spacelike separation.⁵⁰ An extreme choice of

⁴⁹Particularly in the absence of spacelike separation. It is the need to model spacelike separation in a quantum Bell experiment that makes commutativity a necessary (and sufficient) condition for joint measurability of spacelike separated observables in a Bell scenario

⁵⁰To incorporate such a constraint, spacelike separation needs to be modelled via either the tensor product paradigm

joint POVM is the following:

$$G_{ab}^{PR(xy)} = \frac{\mathbb{I}_{\mathcal{H}}}{2} \delta_{a \oplus b, xy}, \quad (141)$$

which leads to the probability distribution $p(a, b|x, y) = \frac{1}{2} \delta_{a \oplus b, xy}$ for any choice of quantum state. Hence, this joint POVM $G^{PR(xy)}$ always yields statistics corresponding to the PR-box, maximally violating the CHSH-type inequality for this scenario, namely,

$$\sum_{a, b, x, y | a \oplus b = xy} \frac{1}{4} p(a, b|x, y) \leq \frac{3}{4}. \quad (142)$$

Physically, it's possible to implement this (without requiring any quantum resources) by providing a box that always produces these correlations between measurement settings denoted by $(xy) \in \{0, 1\}^2$, regardless of the input state. Such a black-box would maximally violate the CHSH-type inequality (viewed as a Bell-KS inequality witnessing KS-contextuality), but that shouldn't be surprising in the absence of spacelike separation. Also, the trivial PR-box joint POVM $G_{ab}^{PR(xy)}$ is a perfectly valid way to implement the joint measurement of trivial POVMs $A^{(x)}$ and $B^{(y)}$ within the standard paradigm of operational quantum theory.⁵¹

To summarize, we note the following:

- Within the traditional framework of KS-noncontextuality, if one wants to go beyond projective measurements to arbitrary POVMs in a contextuality scenario, then one must — in order to avoid the pathology of trivial POVMs violating the Bell-KS inequalities maximally — *restrict* by fiat the notion of joint measurability to merely commutativity. This is, for example, the attitude adopted in Ref. [25].

or the commutativity paradigm. Both these ways of modelling spacelike separation lead to the same set of quantum correlations for any finite-dimensional Hilbert space $\mathcal{H}$ [79]. The question of whether the two paradigms lead to the same set of correlations in the case of infinite dimensional Hilbert spaces is the subject of Tsirelson's problem [79, 80]. Most studies of Bell-nonlocality are primarily concerned with finite dimensional Hilbert spaces; should one encounter infinite dimensional Hilbert spaces, the commutativity paradigm is the proper way to model spacelike separation.

⁵¹Note that the point of this demonstration is to show how, in the absence of spacelike separation justifying commutativity or a promise that the measurements are sharp, arbitrary correlations are achievable in quantum theory if unsharp measurements are allowed. All trivial POVMs are unsharp, but the converse is not true. That is, one can consider nontrivial POVMs that don't violate the CHSH-type inequality maximally, but which violate it (arbitrarily) more than is allowed by sharp measurements in quantum theory. One could construct them, for example, by just taking a convex combination of the PR-box joint POVM with some sharp (and thus product) joint POVM.

- However, if one is going beyond projective measurements, we know that commutativity is *only* a sufficient condition for joint measurability, not a necessary one [44].
- This brings us to our observation that the traditional notion of KS-noncontextuality is pathological once the most general situation in quantum theory is considered: arbitrary POVMs with the general notion of joint measurability (see, e.g., Ref. [44] for this notion and its relation to commutativity). In particular, in the absence of spacelike separation, there is no physical justification to restrict the notion of joint measurability to merely commutativity.
- A similar consideration applies at the level of a KS-noncontextual ontological model: there, factorizability is not justified in the absence of spacelike separation. So, on those grounds alone, one should go beyond KS-noncontextuality as one's notion of classicality; particularly, if one wants a notion of classicality that does not presume outcome determinism, just as *local causality* doesn't presume it. This was argued in Ref. [27]: imagine an adversarial setting where because of the absence of spacelike separation in a KS-contextuality experiment, two measurement settings on the same system can exhibit correlations that are independent of those induced by the system on which the measurements are being implemented, thus allowing them to exhibit stronger correlations than are possible in a KS-noncontextual model. We use trivial POVMs only to drive home that this can be done arbitrarily well (achieving PR-box type correlations, in fact) if there is no constraint on the strength of correlations the measurement settings can exhibit. The way such constraints on the correlations between the measurement settings show up in our analysis within the Spekkens framework is in terms of the quantity *Corr*: if *Corr* is really high, the measurements in a noncontextual ontological model cannot be arbitrarily strongly correlated, i.e., R cannot be arbitrarily high (cf. Eq. (72)).

D The KS-uncolourable hypergraph Γ_{18}

It is instructive to consider the KS-uncolourable hypergraph Γ_{18} , originally appearing in Ref. [51], and studied in the light of Spekkens contextuality in Ref. [12]. This hypergraph fails both criteria for the hypergraphs Γ considered in this paper, namely, $\mathcal{C}(\Gamma) \neq \emptyset$ (KS-colourability) and $\mathcal{CE}^1(\Gamma) = \mathcal{G}(\Gamma)$.

For probabilistic models on Γ_{18} , the following hold: $\mathcal{C}(\Gamma_{18}) = \emptyset \subsetneq \mathcal{CE}^1(\Gamma_{18}) \subsetneq \mathcal{G}(\Gamma_{18})$. This was con-

Figure 8: The hypergraph Γ_{27} and its subhypergraphs, i.e., Γ_{18} and Γ_3 , appearing in the three Bell-KS expressions of Eq. (143). The probabilistic model p considered in Eq. (143) is a probabilistic model on Γ_{27} , and *not* on the subhypergraphs. We have illustrated the subhypergraphs separately only for clarity regarding the subsets of vertices to which the Bell-KS expressions refer: the probabilities assigned to these vertices are obtained from probabilistic models on Γ_{27} .

sidered in Ref. [12], where $\mathcal{CE}^1(\Gamma_{18})$ excludes the extremal probabilistic model in $\mathcal{G}(\Gamma_{18})$ that corresponds to the upper bound on the noise-robust noncontextuality inequality of Ref. [12]. As argued in Ref. [12], this noise-robust noncontextuality inequality is the appropriate operational generalization (to possibly noisy measurements) of the Kochen-Specker contradiction first demonstrated in Ref. [51]; this generalization cannot be accommodated in our generalization of the CSW framework [22].

If one extends the KS-uncolourable Γ_{18} to a KS-colourable hypergraph Γ_{27} with 9 “no-detection” events, one for each hyperedge, then we have $\mathcal{C}(\Gamma_{27}) \neq \emptyset$, but it's still the case that $\mathcal{C}(\Gamma_{27}) \subsetneq \mathcal{CE}^1(\Gamma_{27}) \subsetneq \mathcal{G}(\Gamma_{27})$ for this hypergraph.⁵² Hence, Γ_{27} cannot be understood in our generalization of the CSW framework either.⁵³

Indeed, if one “blindly” writes down a CSW classical bound for some Bell-KS expression defined on

⁵²This follows from noting that extremal probabilistic models on Γ_{18} are still extremal probabilistic models on Γ_{27} : ones where the no-detection events are assigned zero probabilities. See Theorem 2.5.3 of Ref. [23].

⁵³Note that adding these no-detection events is equivalent to allowing subnormalized probabilities (i.e., sum of probabilities assigned to measurement events in a hyperedge can be less than 1) on Γ_{18} . Hence, even allowing for subnormalization on Γ_{18} , which means that one is looking at probabilistic models on the hypergraph Γ_{27} , does not eliminate the gap between $\mathcal{CE}^1$ probabilistic models and general probabilistic models, so that any upper bound on a Bell-KS expression given by probabilistic models in $\mathcal{CE}^1(\Gamma_{27})$ is not always the same as the general probabilistic upper bound from probabilistic models in $\mathcal{G}(\Gamma_{27})$. The CSW framework only considers the upper bound given by $\mathcal{CE}^1(\Gamma_{27})$ probabilistic models.

Figure 9: Going from the orthogonality graph, G , of Γ_{18} to the hypergraph Γ_G (on the right) to which our noise-robust noncontextuality inequality pertains.

$O(\Gamma_{18})$, then such a bound is equivalently a bound for the same Bell-KS expression defined on Γ_{27} (where normalization is restored). Further, the E1 bound on Γ_{18} is a CE^1 bound on Γ_{27} . The GPT bound *happens* to agree with the CE^1 bound for a particular Bell-KS expression (sum of all probabilities) but *differs* for some other Bell-KS expressions defined on this hypergraph. Consider, for example, the following three expressions (see Fig. 8):

$$\begin{aligned} \text{Expr}_1 &\equiv \sum_{v \in V(\Gamma_{18})} p(v), \\ \text{Expr}_2 &\equiv \sum_{v \in V(\Gamma_3)} p(v), \\ \text{Expr}_3 &\equiv \sum_{v \in V(\Gamma_{18})} p(v) + \sum_{v \in V(\Gamma_3)} p(v). \end{aligned} \quad (143)$$

We have:

$$\begin{aligned} \text{Expr}_1 &\stackrel{C(\Gamma_{27})}{\leq} 8 \stackrel{C\mathcal{E}^1(\Gamma_{27})}{<} 9 \stackrel{\mathcal{G}(\Gamma_{27})}{=} 9, \\ \text{Expr}_2 &\stackrel{C(\Gamma_{27})}{\leq} 1 \stackrel{C\mathcal{E}^1(\Gamma_{27})}{=} 1 \stackrel{\mathcal{G}(\Gamma_{27})}{<} \frac{3}{2}, \\ \text{Expr}_3 &\stackrel{C(\Gamma_{27})}{\leq} 9 \stackrel{C\mathcal{E}^1(\Gamma_{27})}{<} 10 \stackrel{\mathcal{G}(\Gamma_{27})}{<} 10.5. \end{aligned} \quad (144)$$

Thus, Expr_3 is a Bell-KS expression that discriminates between probabilistic models at all three levels of the hierarchy. Indeed, the upper bound on Expr_3 for $C\mathcal{E}^1(\Gamma_{27})$ models can be saturated by projective quantum realizations of the hypergraph, in particular the standard realization with 18 rays, with the zero operator for the no-detection events [51]. The fact that there exists such a Bell-KS expression as Expr_3 means that the CE^1 upper bounds from the CSW approach can be violated by a general probabilistic model, i.e., the upper bounds for CE^1 models and

general probabilistic models don't agree, and we cannot take the graph-theoretic upper bounds of CSW for granted in our noise-robust noncontextuality inequalities. Indeed, the general probabilistic upper bound for any Bell-KS expression defined on a contextuality scenario is a hypergraph invariant — in the sense that it is a property that is shared by all hypergraphs isomorphic to each other — that may or may not be expressible as a graph invariant à la CSW.

What, then, do the bounds given by graph invariants of CSW for $O(\Gamma_{18})$ mean in our generalization of the CSW framework? Following our approach, outlined in Sec. III.B, we can go from $G = O(\Gamma_{18})$ to the hypergraph $\Gamma_G = \Gamma_{O(\Gamma_{18})}$ (see Fig. 9) for which we have (by construction) $\mathcal{C}(\Gamma_{O(\Gamma_{18})}) \neq \emptyset$ (so that the underlying hypergraph is no longer KS-uncolourable) and $\mathcal{C}\mathcal{E}^1(\Gamma_{O(\Gamma_{18})}) = \mathcal{G}(\Gamma_{O(\Gamma_{18})})$ (so that, for any Bell-KS expression, the upper bound given by the fractional packing number $\alpha^*(G, w)$ in the CSW framework agrees with the general probabilistic upper bound). Since this construction proceeds by converting all maximal cliques in Γ_{18} to hyperedges in $\Gamma_{O(\Gamma_{18})}$ and adding a new vertex to each such hyperedge, it achieves both purposes: firstly, adding a (no-detection) vertex to every maximal clique that is a hyperedge in Γ_{18} ensures the KS-colourability of $\Gamma_{O(\Gamma_{18})}$, i.e., $\mathcal{C}(\Gamma_{O(\Gamma_{18})}) \neq \emptyset$, and secondly, adding a vertex to every maximal clique that is not a hyperedge in Γ_{18} ensures that $\mathcal{C}\mathcal{E}^1(\Gamma_{O(\Gamma_{18})}) = \mathcal{G}(\Gamma_{O(\Gamma_{18})})$. Once these two properties are satisfied, the graph invariants of CSW [22] become applicable to any Bell-KS expression defined for any set of vertices in the sub-hypergraph Γ_{18} of $\Gamma_{O(\Gamma_{18})}$.

Our noise-robust noncontextuality inequality then applies to the KS-colourable hypergraph $\Gamma_{O(\Gamma_{18})}$, where the graph invariants of CSW make sense, rather than the KS-uncolourable hypergraph Γ_{18} . On the other hand, an appropriate noise-robust noncontextuality inequality for the KS-uncolourable hypergraph Γ_{18} is, then, the one reported in Ref. [12].⁵⁴

References

- [1] L. Hardy, “Quantum Theory From Five Reasonable Axioms”, [arXiv:quant-ph/0101012](#) (2001).
- [2] L. Masanes and M. P. Mueller, “A derivation of quantum theory from physical requirements”, [New J. Phys.](#) **13**, 063001 (2011).
- [3] G. Chiribella, G. M. D’Ariano, and P. Perinotti, “Probabilistic theories with purification”, [Phys. Rev. A](#) **81**, 062348 (2010).
- [4] J. S. Bell, “On the Einstein-Podolsky-Rosen paradox”, [Physics](#) **1**, 195 (1964). Reprinted in Ref. [6], Chapter 2.

⁵⁴The approach for KS-uncolourable hypergraphs will be further developed in hypergraph-theoretic terms in forthcoming work [34].

- [5] J. S. Bell, “On the problem of hidden variables in quantum mechanics”, *Rev. Mod. Phys.* **38**, 447 (1966). Reprinted in Ref. [6], Chapter 1.
- [6] J. S. Bell, “Speakable and Unspeakable in Quantum Mechanics”, 2nd Edition, *Cambridge University Press*, 2004.
- [7] 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).
- [8] N. Brunner, D. Cavalcanti, S. Pironio, V. Scarani, and S. Wehner, “Bell nonlocality”, *Rev. Mod. Phys.* **86**, 419 (2014).
- [9] B. Hensen *et al.*, “Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres”, *Nature* **526**, 682 - 686 (2015).
- [10] Lynden K. Shalm *et al.*, “Strong Loophole-Free Test of Local Realism”, *Phys. Rev. Lett.* **115**, 250402 (2015).
- [11] M. Giustina *et al.*, “Significant-Loophole-Free Test of Bell’s Theorem with Entangled Photons”, *Phys. Rev. Lett.* **115**, 250401 (2015).
- [12] R. Kunjwal and R. W. Spekkens, “From the Kochen-Specker Theorem to Noncontextuality Inequalities without Assuming Determinism”, *Phys. Rev. Lett.* **115**, 110403 (2015).
- [13] M. D. Mazurek, M. F. Pusey, R. Kunjwal, K. J. Resch, R. W. Spekkens, “An experimental test of noncontextuality without unphysical idealizations”, *Nat. Commun.* **7**, 11780 (2016).
- [14] A. Krishna, R. W. Spekkens, and E. Wolfe, “Deriving robust noncontextuality inequalities from algebraic proofs of the Kochen-Specker theorem: the Peres-Mermin square”, *New J. Phys.* **19**, 123031 (2017).
- [15] D. Schmid and R. W. Spekkens, “Contextual Advantage for State Discrimination”, *Phys. Rev. X* **8**, 011015 (2018).
- [16] R. Kunjwal and R. W. Spekkens, “From statistical proofs of the Kochen-Specker theorem to noise-robust noncontextuality inequalities”, *Phys. Rev. A* **97**, 052110 (2018).
- [17] D. Schmid, R. W. Spekkens, and E. Wolfe, “All the noncontextuality inequalities for arbitrary prepare-and-measure experiments with respect to any fixed set of operational equivalences”, *Phys. Rev. A* **97**, 062103 (2018).
- [18] R. W. Spekkens, “Contextuality for preparations, transformations, and unsharp measurements”, *Phys. Rev. A* **71**, 052108 (2005).
- [19] S. Kochen and E. P. Specker, “The Problem of Hidden Variables in Quantum Mechanics”, *J. Math. Mech.* **17**, 59 (1967). Also available at JS-TOR.
- [20] N. Harrigan and R. W. Spekkens, “Einstein, Incompleteness, and the Epistemic View of Quantum States,” *Found. Phys.* **40**, 125 (2010).
- [21] A. Cabello, S. Severini, and A. Winter, “(Non-)Contextuality of Physical Theories as an Axiom”, [arXiv:1010.2163 \[quant-ph\]](https://arxiv.org/abs/1010.2163) (2010).
- [22] A. Cabello, S. Severini, and A. Winter, “Graph-Theoretic Approach to Quantum Correlations”, *Phys. Rev. Lett.* **112**, 040401 (2014).
- [23] A. Acín, T. Fritz, A. Leverrier, and A. B. Sainz, “A Combinatorial Approach to Nonlocality and Contextuality”, *Comm. Math. Phys.* **334**(2), 533-628 (2015).
- [24] J. Barrett, “Information processing in generalized probabilistic theories”, *Phys. Rev. A* **75**, 032304 (2007).
- [25] S. Abramsky and A. Brandenburger, “The sheaf-theoretic structure of non-locality and contextuality”, *New J. Phys.* **13**, 113036 (2011).
- [26] A. Fine, “Hidden Variables, Joint Probability, and the Bell Inequalities”, *Phys. Rev. Lett.* **48**, 291 (1982).
- [27] R. Kunjwal, “Fine’s theorem, noncontextuality, and correlations in Specker’s scenario”, *Phys. Rev. A* **91**, 022108 (2015).
- [28] R. W. Spekkens, “The Status of Determinism in Proofs of the Impossibility of a Noncontextual Model of Quantum Theory”, *Found. Phys.* **44**, 1125-1155 (2014).
- [29] A. Cabello, “What do we learn about quantum theory from Kochen-Specker quantum contextuality?”, [PIRSA:17070034](https://arxiv.org/abs/1707.0034) (2017).
- [30] G. Chiribella and X. Yuan, “Measurement sharpness cuts nonlocality and contextuality in every physical theory”, [arXiv:1404.3348 \[quant-ph\]](https://arxiv.org/abs/1404.3348) (2014).
- [31] G. Chiribella and X. Yuan, “Bridging the gap between general probabilistic theories and the device-independent framework for nonlocality and contextuality”, *Information and Computation*, **250**, 15-49 (2016).
- [32] R. Chaves and T. Fritz, “Entropic approach to local realism and noncontextuality”, *Phys. Rev. A* **85**, 032113 (2012).
- [33] Tobias Fritz and Rafael Chaves, “Entropic Inequalities and Marginal Problems”, *IEEE Trans. on Information Theory*, vol. 59, pages 803 - 817 (2013).
- [34] R. Kunjwal, “Hypergraph framework for irreducible noncontextuality inequalities from logical proofs of the Kochen-Specker theorem”, [arXiv:1805.02083 \[quant-ph\]](https://arxiv.org/abs/1805.02083) (2018).
- [35] A. Cabello, “Specker’s fundamental principle of quantum mechanics”, [arXiv:1212.1756 \[quant-ph\]](https://arxiv.org/abs/1212.1756) (2012).
- [36] R. W. Spekkens, “Noncontextuality: how we should define it, why it is natural, and what to do about its failure”, [PIRSA:17070035](https://arxiv.org/abs/1707.0035) (2017).
- [37] 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 \[quant-ph\]](#) (2017).
- [38] M. F. Pusey, L. del Rio, and B. Meyer, “Contextuality without access to a tomographically complete set”, [arXiv:1904.08699](#) (2019).
- [39] Y. C. Liang, R. W. Spekkens, H. M. Wiseman, “Specker’s parable of the overprotective seer: A road to contextuality, nonlocality and complementarity”, *Phys. Rep.* **506**, 1 (2011).
- [40] R. Kunjwal and S. Ghosh, “Minimal state-dependent proof of measurement contextuality for a qubit”, *Phys. Rev. A* **89**, 042118 (2014).
- [41] R. Kunjwal, C. Heunen, and T. Fritz, “Quantum realization of arbitrary joint measurability structures”, *Phys. Rev. A* **89**, 052126 (2014).
- [42] R. Kunjwal, “A note on the joint measurability of POVMs and its implications for contextuality”, [arXiv:1403.0470 \[quant-ph\]](#) (2014).
- [43] S. Popescu and D. Rohrlich, “Quantum nonlocality as an axiom”, *Found. Phys.* **24**, 379-385 (1994).
- [44] T. Heinosaari, D. Reitzner, and P. Stano, “Notes on Joint Measurability of Quantum Observables”, *Found. Phys.* **38**, 1133-1147 (2008).
- [45] R. Kunjwal, “How to go from the KS theorem to experimentally testable noncontextuality inequalities”, [PIRSA:17070059](#) (2017).
- [46] Konrad Engel, “Sperner theory: Encyclopedia of Mathematics and its Applications”, Vol. 65, Cambridge University Press, Cambridge (1997).
- [47] A. A. Klyachko, M. A. Can, S. Binicioğlu, and A. S. Shumovsky, “Simple Test for Hidden Variables in Spin-1 Systems”, *Phys. Rev. Lett.* **101**, 020403 (2008).
- [48] C. Held, “The Kochen-Specker Theorem”, *The Stanford Encyclopedia of Philosophy* (Spring 2018 Edition), Edward N. Zalta (ed.).
- [49] T. Gonda, R. Kunjwal, D. Schmid, E. Wolfe, and A. B. Sainz, “Almost Quantum Correlations are Inconsistent with Specker’s Principle”, *Quantum* **2**, 87 (2018).
- [50] M. Navascués, Y. Guryanova, M. J. Hoban, and A. Acín, “Almost quantum correlations”, *Nat. Commun.* **6**, 6288 (2015).
- [51] A. Cabello, Adán, J. Estebarez, and G. Garcia-Alcaine, “Bell-Kochen-Specker theorem: A proof with 18 vectors,” *Phys. Lett. A* **212**, 183 (1996).
- [52] E. G. Beltrametti and S. Bugajski, “A classical extension of quantum mechanics”, *J. Phys. A* **28**, 3329 (1995).
- [53] X. Zhan, E. G. Cavalcanti, J. Li, Z. Bian, Y. Zhang, H. M. Wiseman, and P. Xue, “Experimental generalized contextuality with single-photon qubits”, *Optica* **4**, 966-971 (2017).
- [54] R. Kunjwal, “Contextuality beyond the Kochen-Specker theorem”, [arXiv:1612.07250 \[quant-ph\]](#) (2016).
- [55] T. Fritz, A. B. Sainz, R. Augusiak, J. B. Brask, R. Chaves, A. Leverrier, and A. Acín, “Local orthogonality: a multipartite principle for correlations”, *Nat. Commun.* **4**, 2263 (2013).
- [56] R. W. Spekkens, “Nonclassicality as the failure of noncontextuality”, [PIRSA:15050081](#) (2015) (see the slide at 41:43 minutes).
- [57] R. W. Spekkens, “Quasi-Quantization: Classical Statistical Theories with an Epistemic Restriction”, In: Chiribella G., Spekkens R. (eds) *Quantum Theory: Informational Foundations and Foils. Fundamental Theories of Physics*, vol 181. Springer, Dordrecht.
- [58] T. Vidick and S. Wehner, “Does Ignorance of the Whole Imply Ignorance of the Parts? Large Violations of Noncontextuality in Quantum Theory”, *Phys. Rev. Lett.* **107**, 030402 (2011).
- [59] R. Raussendorf, “Contextuality in measurement-based quantum computation”, *Phys. Rev. A* **88**, 022322 (2013).
- [60] M. Howard, J. Wallman, V. Veitch, and J. Emerson, “Contextuality supplies the ‘magic’ for quantum computation”, *Nature* **510**, 351 (2014).
- [61] N. Delfosse, P. A. Guerin, J. Bian, and R. Raussendorf, “Wigner Function Negativity and Contextuality in Quantum Computation on Rebits”, *Phys. Rev. X* **5**, 021003 (2015).
- [62] J. Bermejo-Vega, N. Delfosse, D. E. Browne, C. Okay, R. Raussendorf, “Contextuality as a resource for qubit quantum computation”, *Phys. Rev. Lett.* **119**, 120505 (2017).
- [63] J. Singh, K. Bharti, and Arvind, “Quantum key distribution protocol based on contextuality monogamy”, *Phys. Rev. A* **95**, 062333 (2017).
- [64] A. Cabello, “Kochen-Specker Theorem for a Single Qubit using Positive Operator-Valued Measures”, *Phys. Rev. Lett.* **90**, 190401 (2003).
- [65] A. Grudka and P. Kurzyński, “Is There Contextuality for a Single Qubit?”, *Phys. Rev. Lett.* **100**, 160401 (2008).
- [66] P. Busch, “Quantum States and Generalized Observables: A Simple Proof of Gleason’s Theorem”, *Phys. Rev. Lett.* **91**, 120403 (2003).
- [67] C. M. Caves, C. A. Fuchs, K. Manne, and J. M. Renes, “Gleason-Type Derivations of the Quantum Probability Rule for Generalized Measurements”, *Found. Phys.* **34**, 193 (2004).
- [68] A. M. Gleason, “Measures on the closed subspaces of a Hilbert space”, *J. Math. Mech.* **6**, 885 (1957). Also available at JSTOR.
- [69] P. K. Aravind, “The generalized Kochen-Specker theorem”, *Phys. Rev. A* **68**, 052104 (2003).
- [70] A. A. Methot, “Minimal Bell-Kochen-Specker proofs with POVMs on qubits”, *Int. J. Quantum Inf.* **5**, 353 (2007).
- [71] Q. Zhang, H. Li, T. Yang, J. Yin, J. Du, J. W. Pan, “Experimental Test of the Kochen-Specker Theorem for Single Qubits using Positive Operator-Valued Measures”, [arXiv:quant-ph/0412049](#) (2004).

- [72] L. Mancinska, G. Scarpa, and S. Severini, “New Separations in Zero-Error Channel Capacity Through Projective Kochen Specker Sets and Quantum Coloring”, *IEEE Transactions on Information Theory* **59**, 4025 (2013).
- [73] J. Henson and A. B. Sainz, “Macroscopic non-contextuality as a principle for almost-quantum correlations”, *Phys. Rev. A* **91**, 042114 (2015).
- [74] D. A. Meyer, “Finite Precision Measurement Nullifies the Kochen-Specker Theorem”, *Phys. Rev. Lett.* **83**, 3751 (1999).
- [75] A. Kent, “Noncontextual Hidden Variables and Physical Measurements”, *Phys. Rev. Lett.* **83**, 3755 (1999).
- [76] R. Clifton and A. Kent, “Simulating quantum mechanics by non-contextual hidden variables”, *Proc. R. Soc. Lond. A: Vol. 456*, 2101-2114 (2000).
- [77] J. Barrett and A. Kent, “Non-contextuality, finite precision measurement and the Kochen-Specker theorem”, *Stud. Hist. Philos. Mod. Phys.* **35**, 151 (2004).
- [78] A. Winter, “What does an experimental test of quantum contextuality prove or disprove?”, *J. Phys. A: Math. Theor.* **47**, 424031 (2014).
- [79] V. B. Scholz and R. F. Werner, “Tsirelson’s Problem”, [arXiv:0812.4305 \[math-ph\]](https://arxiv.org/abs/0812.4305) (2008).
- [80] T. Fritz, “Tsirelson’s problem and Kirchberg’s conjecture”, *Rev. Math. Phys.* **24** (5), 1250012 (2012).