Sheaf-Theoretic Preparation Contextuality via Stochastic Extension

Tom Williams   Mina Doosti
School of Informatics
Quantum Software Lab, University of Edinburgh
United Kingdom
T.A.H.Williams@sms.ed.ac.uk
mdoosti@ed.ac.uk

Farid Shahandeh
Department of Computer Science
Royal Holloway, University of London
United Kingdom
farid.shahandeh@rhul.ac.uk

1 Introduction

Contextuality, originating in the Bell–Kochen–Specker theorem [2, 4], is the obstruction to modelling operational data by a single classical description. Abramsky–Brandenburger [1] recast contextuality as an extension obstruction in a sheaf-theoretic framework, while Spekkens gave an operational account that extends contextuality to preparations and transformations beyond measurement incompatibility and outcome determinism [5]. This suggests a natural question: can preparation contextuality itself be formulated as an extension obstruction within a sheaf-theoretic framework?

In this work, we provide an affirmative answer. We give an account of preparation contextuality directly within the sheaf-theoretic framework, treating preparation data as the primary empirical object and characterising contextuality as an obstruction to stochastic extension.

In measurement scenarios, restriction from global to local data is canonical and given by marginalisation. In preparation scenarios the dual operation is stochastic completion of partially specified preparation instances. Unlike marginalisation, no canonical completion rule exists. We introduce minimal structural requirements on admissible completion maps—extension-respecting, input independence, and compositionality—and show that these assumptions force a rigid product form for admissible stochastic extension maps. Preparation contextuality arises when an empirical model admits a preparation-compatible extension structure but none admits a global response matrix reproducing the empirical data.

We illustrate the framework with a Pusey–Barrett–Rudolph (PBR) type scenario [3], where the combinatorial zero-structure of the empirical data yields an explicit obstruction to stochastic extension.

2 Measurement contextuality as an extension problem

Let X be a finite set of measurements with outcome set O, and let $\mathscr{C}$ be a cover of X by jointly measurable contexts [1]. For $C \in \mathscr{C}$, empirical data for a preparation p specify distributions $E_{C|p} \in \Delta(O^C)$.

For $V \subseteq U \subseteq X$, marginalisation is encoded by

$$M_{V|U}(s_V \mid s_U) = \mathbf{1}\left[(s_U)|_V = s_V\right].$$

Measurement noncontextuality asserts the existence of a single global distribution $D_{X|p} \in \Delta(O^X)$ such that

$$E_{C|p} = M_{C|X}\, D_{X|p} \qquad \forall C \in \mathscr{C}.$$

Thus measurement contextuality is a failure of linear extension with fixed restriction maps.

To appear in EPTCS.

3 Preparation scenarios and empirical data

Fix a measurement m with finite outcome set O. Let Y be a finite set of preparation sources, each with instance set I. Joint implementability is encoded by a cover $\mathscr{S}$ of Y, whose elements $\Gamma \subseteq Y$ are source contexts.

For $\Gamma \in \mathscr{S}$, a joint preparation instance is $\sigma_\Gamma \in I^\Gamma$. Performing m yields conditional probabilities $E_{m|\Gamma}(o \mid \sigma_\Gamma)$, equivalently a column-stochastic matrix $E_{m|\Gamma}$ of size $|O| \times |I^\Gamma|$. The empirical model is the family $\{E_{m|\Gamma}\}_{\Gamma\in\mathscr{S}}$.

4 Global response matrices and stochastic extension

A global explanation assigns outcome statistics to every fully specified instance $\sigma_Y \in I^Y$, via a stochastic matrix $D_{m|Y}(o \mid \sigma_Y)$.

To relate $D_{m|Y}$ to empirical data on Γ, one specifies stochastic extension maps $S_{Y|\Gamma}(\sigma_Y \mid \sigma_\Gamma)$.

The compatibility condition is

$$E_{m|\Gamma}(o \mid \sigma_\Gamma) = \sum_{\sigma_Y} D_{m|Y}(o \mid \sigma_Y)\, S_{Y|\Gamma}(\sigma_Y \mid \sigma_\Gamma). \tag{1}$$

In matrix form, $E_{m|\Gamma} = D_{m|Y}\, S_{Y|\Gamma}$. Unlike marginalisation, $S_{Y|\Gamma}$ is not canonical.

5 Admissible extension families

For $W \subseteq V \subseteq U \subseteq Y$, consider stochastic maps $S_{U|V} : I^V \rightsquigarrow I^U$.

Definition 1. *An extension family $\{S_{U|V}\}$ is*

(i) extension-respecting *if*

$$S_{U|V}(\sigma_U \mid \sigma_V) > 0 \;\Rightarrow\; (\sigma_U)|_V = \sigma_V;$$

(ii) input independent *if for each $V \subseteq U$ there exists $s_{U\setminus V} \in \Delta(I^{U\setminus V})$ such that*

$$S_{U|V}(\sigma_U \mid \sigma_V) = \mathbf{1}[(\sigma_U)|_V = \sigma_V]\, s_{U\setminus V}((\sigma_U)|_{U\setminus V});$$

(iii) compositional *if*

$$S_{U|W} = S_{U|V}\, S_{V|W} \qquad \forall W \subseteq V \subseteq U.$$

Extension-respecting expresses that completion does not alter already specified preparation choices. Input independence captures the idea that the statistical rule for completing unspecified sources depends only on which sources are absent, not on the values of those already present. Compositionality formalises consistency of staged completion: extending in one step or in successive refinements yields the same distribution. These conditions are minimal structural requirements on any operationally well-defined completion procedure.

Theorem 1 (Product-form extension). *If $\{S_{U|V}\}$ is extension-respecting, input independent, and compositional, then there exist single-source distributions $\{s_p \in \Delta(I)\}_{p\in Y}$ such that*

$$S_{U|V}(\sigma_U \mid \sigma_V) = \mathbf{1}[(\sigma_U)|_V = \sigma_V] \prod_{p\in U\setminus V} s_p(\sigma_p).$$

Proof sketch. Input independence reduces each $S_{U|V}$ to a distribution $s_{U\setminus V}$. Substituting into compositionality forces multiplicativity constraints for disjoint extensions. Iterating along single-source extensions yields factorisation into independent single-source terms. □

6 Preparation contextuality

Definition 2 (Preparation compatibility)**.** *An empirical model* $\{E_{m|\Gamma}\}_{\Gamma\in\mathscr{S}}$ *and an extension family* $\{S_{U|V}\}$ *are* preparation compatible *if for all overlapping contexts* $\Gamma,\Gamma'\in\mathscr{S}$*,*

$$E_{m|\Gamma}\, S_{\Gamma|\Gamma\cap\Gamma'} = E_{m|\Gamma'}\, S_{\Gamma'|\Gamma\cap\Gamma'}.$$

Preparation compatibility expresses that statistics obtained by completing preparation control to a common overlap are independent of the larger source context in which they arise.

Definition 3. *An empirical model* $\{E_{m|\Gamma}\}$ *is* preparation noncontextual *if there exists an admissible extension family* $\{S_{U|V}\}$ *such that*

1. (E,S) *satisfies preparation compatibility, and*
2. *there exists a global response matrix* $D_{m|Y}$ *such that*

$$E_{m|\Gamma} = D_{m|Y}\, S_{Y|\Gamma} \qquad \forall\Gamma\in\mathscr{S}.$$

If there exists at least one admissible extension family satisfying preparation compatibility but no such compatible extension family admits a global response matrix, the empirical model is preparation contextual.

Let $\overline{E}_{m|\Gamma}$ denote the possibilistic reduction of $E_{m|\Gamma}$.

Theorem 2. *Possibilistic preparation contextuality implies probabilistic preparation contextuality.*

Proof sketch. A probabilistic solution induces a possibilistic one by taking supports. Zeros are preserved under stochastic multiplication. □

7 Quantum illustration (PBR-type scenario)

We illustrate preparation contextuality using a construction in the spirit of PBR [3]. Consider four sources $Y=\{a,b,a',b'\}$ with binary instances $I=\{0,1\}$ and contexts $\mathscr{S}=\{\{a,b\},\{a,b'\},\{a',b\},\{a',b'\}\}$. Sources a,b prepare computational-basis states $|0\rangle,|1\rangle$, while a',b' prepare Hadamard-basis states $|+\rangle,|-\rangle$. Each context $\Gamma\in\mathscr{S}$, therefore, prepares one of four product states ρ^{Γ}_{ij}, $i,j\in\{0,1\}$.

A fixed four-outcome POVM $\{M_k\}_{k=1}^{4}$, as introduced in [3], has the property that in each context and for each pair (i,j) exactly one outcome k is assigned zero probability:

$$E_{m|\Gamma}(k\mid i,j) = \mathrm{Tr}[M_k\rho^{\Gamma}_{ij}], \qquad \text{with exactly one zero per column.}$$

Assume there exist a global response matrix $D_{m|Y}$ and an admissible extension family reproducing $\{E_{m|\Gamma}\}$ via (1). By Theorem 1, admissible completion maps factor across sources. Hence every global refinement of a locally specified instance receives a nonzero weight. Each empirical zero therefore forces the corresponding global response entry to vanish for all compatible global assignments.

Combining the constraints from the four contexts yields a parity-type contradiction: there exists a global assignment $\sigma_Y\in I^Y$ for which all four outcomes are forced to vanish simultaneously, contradicting column-stochasticity of $D_{m|Y}$. Since the empirical model admits a preparation-compatible extension structure (for instance the uniform choice of single-site distributions $\{s_p\}$), and the above argument is independent of the choice of $\{s_p\}$, no admissible extension map admits a global response matrix. The empirical model is therefore (possibilistically) preparation contextual.

References

[1] Samson Abramsky & Adam Brandenburger (2011): *The sheaf-theoretic structure of non-locality and contextuality*. New Journal of Physics 13(11), p. 113036, doi:10.1088/1367-2630/13/11/113036. Available at `https://dx.doi.org/10.1088/1367-2630/13/11/113036`.

[2] JOHN S. BELL (1966): *On the Problem of Hidden Variables in Quantum Mechanics*. Rev. Mod. Phys. 38, pp. 447–452, doi:10.1103/RevModPhys.38.447. Available at `https://link.aps.org/doi/10.1103/RevModPhys.38.447`.

[3] Matthew F. Pusey, Jonathan Barrett & Terry Rudolph (2012): *On the reality of the quantum state*. Nature Physics 8(6), pp. 475–478, doi:10.1038/nphys2309.

[4] E. Specker Simon Kochen (1968): *The Problem of Hidden Variables in Quantum Mechanics*. Indiana Univ. Math. J. 17, pp. 59–87.

[5] R. W. Spekkens (2005): *Contextuality for preparations, transformations, and unsharp measurements*. Phys. Rev. A 71, p. 052108, doi:10.1103/PhysRevA.71.052108. Available at `https://link.aps.org/doi/10.1103/PhysRevA.71.052108`.

Preparation Contextuality as an Obstruction to Stochastic Extension

Tom Williams,[1] Mina Doosti,[1] and Farid Shahandeh[2]

[1]*School of Informatics, Quantum Software Lab, University of Edinburgh, United Kingdom*
[2]*Department of Computer Science, Royal Holloway, University of London, United Kingdom*

We introduce a preparation-dual notion of contextuality, formulated as an obstruction to stochastic extension. In direct analogy with the sheaf-theoretic formulation of measurement contextuality, preparation contextuality arises when locally specified preparation statistics cannot be extended to a single global response model compatible with all source contexts. Whereas measurement contextuality concerns the incompatibility of restriction maps (marginalisation), the preparation setting requires stochastic extension of partial conditioning data, which is inherently non-unique. We identify minimal structural and preparation compatibility conditions on admissible extension maps and show that they enforce a rigid product form. This leads to a notion of preparation contextuality: whenever an empirical model admits at least one admissible preparation-compatible extension structure, it is contextual precisely when none of these compatible extension structures admits a global response matrix. The framework is formulated explicitly in matrix form and illustrated by a quantum-mechanical example exhibiting preparation contextuality.

Contextuality has emerged as a central nonclassical feature of quantum theory, capturing the impossibility of explaining certain experimental statistics by a single underlying classical description [1]. It plays a foundational role in quantum information, where it underlies nonlocality and Bell inequalities [2], and admits operational [3], logical [4], and sheaf-theoretic [5] formulations.

In the standard treatment, contextuality is associated with measurements. Empirical data are obtained from different measurement contexts, and noncontextuality corresponds to the existence of a single global assignment of outcomes whose restrictions reproduce the observed statistics [5]. Equivalently, measurement contextuality can be characterised as the failure to extend a family of compatible local distributions to a global one, closely related to Fine's joint-distribution criterion in Bell scenarios [6]. In the sheaf-theoretic formulation, contextuality arises as an obstruction to global section existence under canonical restriction by marginalisation.

In this work we develop a complementary, preparation-dual notion of contextuality within the sheaf-theoretic framework. Existing approaches to preparation contextuality, most notably Spekkens' operational formulation [3], characterise noncontextuality in terms of ontological models constrained by operational equivalences. While powerful, these approaches do not formulate preparation contextuality as an obstruction to extending compatible local preparation data to a single global object, in the manner that the sheaf-theoretic treatment does for measurement contextuality [5]. Our aim is to identify and formalise precisely this extension structure for preparation scenarios.

The key structural distinction is that, although both measurement and preparation contextuality can be formulated as extension problems, the passage from global to local data is fundamentally different. In the measurement case, local statistics are obtained from global data by canonical marginalisation. In the preparation case, they arise through stochastic averaging over unobserved preparation degrees of freedom. The corresponding completion maps are therefore not fixed: there are generally many inequivalent ways to extend a partial preparation specification to a global one. This non-uniqueness is not merely a technical complication. In the absence of a canonical completion rule, one must impose structural coherence conditions on admissible extension maps. No analogous step is required in the measurement framework, where restriction maps are determined a priori.

We formulate preparation contextuality as an obstruction to stochastic extension. An empirical model is preparation noncontextual if there exists an admissible preparation-compatible extension map together with a single global response matrix reproducing the observed local preparation statistics. If the empirical model admits at least one preparation-compatible extension structure but none admits a global response matrix, it is preparation contextual. Admissible extension maps are defined by three minimal structural requirements: extension-respecting, input independence, and compositionality. These conditions rigidly constrain the form of completion, forcing all admissible extension maps to factor across sources. Preparation contextuality arises when this global response matrix fails to exist.

Throughout, we adopt an explicit matrix formulation that renders the underlying structure transparent and makes the extension problem a question of stochastic matrix feasibility. This perspective clarifies the duality between restriction and extension and facilitates comparison with operational and factorisation-based approaches to contextuality. The general framework is illustrated by a quantum-mechanical example inspired by Pusey–Barrett–Rudolph-type constructions [7], which exhibits preparation contextuality.

Before turning to preparations, we briefly recall how the sheaf-theoretic formulation of *measurement* contextuality can be written as an explicit linear-algebraic extension problem [5].

Let X be a finite set of measurements, each with outcome set O, and let $\mathcal{C}$ be a cover of X whose elements represent jointly measurable contexts. For a context $C \in \mathcal{C}$, the joint outcome space is O^C.[1] Operationally, a preparation p together with a choice of context C produces a distribution over these

[1] Here $O^C := \{ s : C \to O \}$ denotes the set of all functions assigning an outcome in O to each measurement in the context C.

joint outcomes; we denote it by

$$E_{C|p} \in \Delta(O^C), \tag{1}$$

where $\Delta(Z)$ is the simplex of probability distributions on a finite set Z. The sheaf perspective emphasizes that whenever two contexts overlap, the statistics on the overlap are fixed by restricting (marginalizing) the larger joint distribution. In the matrix language, these restrictions are captured by fixed 0–1 maps that depend only on the inclusion of contexts. In particular, if $C, C' \in \mathcal{C}$ are contexts, then marginal consistency (no-signalling) requires that the two induced distributions on $O^{C\cap C'}$ coincide:

$$M_{(C\cap C')|C}\, E_{C|p} \;=\; M_{(C\cap C')|C'}\, E_{C'|p}. \tag{2}$$

Here $M_{V|U}$ is the fixed 0–1 marginalization matrix from O^U to O^V, defined entrywise by

$$M_{V|U}(s_V \mid s_U) := \mathbf{1}\big[(s_U)|_V = s_V\big], \qquad (V \subseteq U), \tag{3}$$

where $s_U \in O^U$ denotes a joint assignment of outcomes to the measurements in U, and $(s_U)|_V \in O^V$ denotes the restriction of the outcomes s_U to the subset $V \subseteq U$. In Bell-type scenarios, these consistency conditions coincide with the standard no-signalling constraints [2].

Measurement noncontextuality is the claim that there exists a single global distribution $D_{X|p} \in \Delta(O^X)$ which reproduces all observed context-wise distributions by marginalization. Concretely, for each context C this requires

$$E_{C|p} = M_{C|X} D_{X|p}, \tag{4}$$

where $M_{C|X}$ is the corresponding marginalization matrix.

$$M_{C|X}(s_C \mid s_X) = \mathbf{1}[(s_X)|_C = s_C]. \tag{5}$$

Thus $M_{C|X}$ simply "forgets" the outcomes of measurements in $X \setminus C$ by summing over all global assignments s_X compatible with a given local assignment s_C. Contextuality is precisely the failure of such a global distribution to exist, despite the local data being operationally well-defined.

Example (Abramsky–Brandenburger Bell model [5]). Let $X = \{a, a', b, b'\}$, $O = \{0, 1\}$, and

$$\mathcal{C} = \big\{\{a, b\}, \{a', b\}, \{a, b'\}, \{a', b'\}\big\}. \tag{6}$$

Prepare the Bell state

$$p = |\Phi_+\rangle = \frac{1}{\sqrt{2}}\big(|00\rangle + |11\rangle\big). \tag{7}$$

Let $\theta := \pi/6$, $c := \cos\theta$, $s := \sin\theta$, and define projective measurements by the following orthonormal bases:

$$\begin{aligned} &a = b = \{|0\rangle, |1\rangle\}, \\ &a' = \{|a'_0\rangle, |a'_1\rangle\}, \qquad |a'_0\rangle = c\,|0\rangle + s\,|1\rangle, |a'_1\rangle = -s\,|0\rangle + c\,|1\rangle, \\ &b' = \{|b'_0\rangle, |b'_1\rangle\}, \qquad |b'_0\rangle = c\,|0\rangle - s\,|1\rangle, |b'_1\rangle = s\,|0\rangle + c\,|1\rangle. \end{aligned} \tag{8}$$

For each context $C = \{x, y\} \in \mathcal{C}$ and $(i, j) \in O^C$, define the empirical model by the Born rule

$$E_{C|p}(i, j) = \langle\Phi_+|\left(\Pi_i^{(x)} \otimes \Pi_j^{(y)}\right)|\Phi_+\rangle, \tag{9}$$

where $\Pi_i^{(x)} = |x_i\rangle\langle x_i|$.

Explicit probability table. Ordering columns as $(0,0), (1,0), (0,1), (1,1)$, we obtain:

	$(0,0)$	$(1,0)$	$(0,1)$	$(1,1)$
(a,b)	$\frac{1}{2}$	0	0	$\frac{1}{2}$
(a',b)	$\frac{3}{8}$	$\frac{1}{8}$	$\frac{1}{8}$	$\frac{3}{8}$
(a,b')	$\frac{3}{8}$	$\frac{1}{8}$	$\frac{1}{8}$	$\frac{3}{8}$
(a',b')	$\frac{1}{8}$	$\frac{3}{8}$	$\frac{3}{8}$	$\frac{1}{8}$

This family is compatible (no-signalling), hence defines a candidate for gluing; nevertheless it admits *no global section* in $\mathcal{D}(O^X)$. Equivalently, there is no distribution $D_{X|p}$ whose restrictions to each $C \in \mathcal{C}$ via $M_{C|X} D_{X|p}$ coincide with $E_{C|p}$.

Sheaf-theoretic preparation framework: We now turn to the preparation-dual setting, in which contextuality is probed by varying preparation procedures rather than measurement contexts. To isolate the structure associated with preparations, we fix a single measurement m with outcome set O, and regard all nonclassical features as arising from the manner in which preparation devices can be controlled and combined.

By a *source* we mean a preparation device equipped with a finite classical control interface. Choosing a value of this classical control specifies a particular *preparation instance* of the device. Preparation instances thus label operationally distinct procedures that can be selected prior to measurement, without making any assumptions about the underlying physical theory describing the prepared system. Operationally, a preparation instance plays a role dual to that of a measurement outcome: whereas a measurement outcome is the realised value of a random variable revealed by a measurement, a preparation instance labels which preparation procedure is realised, either by deliberate choice, random selection, or post-selection.

We write Y for the finite set of sources available in the experiment, and I for the finite set of preparation instances associated with each source. In general, not all families of sources can be implemented together within a single experimental arrangement. Instead, only certain collections of sources admit a joint implementation as alternative preparation procedures of a single device. This is directly analogous to the measurement scenario, where a measurement context specifies which sets of observables admit a joint measurement.

This implementability structure is encoded by a cover $\mathcal{S}$ of Y, whose elements $\Gamma \subseteq Y$ are referred to as *source contexts*. A source context represents a family of sources that can be realized together as coarse-grainings of a single underlying preparation device with a refined classical control. Choosing a source context Γ therefore fixes which joint refinements of preparation instances are operationally available. This is directly analogous to the measurement scenario, where a measurement context specifies which sets of observables admit a joint measurement.

For a given source context $\Gamma \in \mathcal{S}$, the experimenter selects a tuple of preparation instances

$$\sigma_\Gamma \in I^\Gamma, \tag{10}$$

corresponding to a choice of classical control for each source in the context. Upon subsequently performing the measurement m, the experimenter observes conditional outcome statistics

$$E_{m|\Gamma}(o \mid \sigma_\Gamma), \qquad o \in O. \tag{11}$$

As in the measurement case, the crucial point is that source contexts specify which collections of preparation procedures admit a common refinement, while the empirical data assign outcome statistics to each such choice. No additional assumptions about convex structure, state spaces, or underlying dynamics are required at this stage. Contextuality will arise precisely when these locally specified statistics cannot be extended to a single global preparation model compatible with all source contexts. Before turning to our main structural results, we illustrate this framework with an example that makes these ideas concrete and exhibits a scenario in which preparation contextuality can arise.

Example (Quantum PBR-type preparation scenario). As a preparation–dual analogue of the Bell example, consider two preparation sites, Alice and Bob, each equipped with two preparation procedures, denoted a, a' for Alice and b, b' for Bob. Each procedure has two preparation instances labelled $i \in \{0, 1\}$. The joint preparation contexts are

$$\Gamma \in \big\{\{a, b\}, \{a, b'\}, \{a', b\}, \{a', b'\}\big\}. \tag{12}$$

On each run, one procedure is chosen at each site together with instance labels $(i, j) \in \{0, 1\}^2$, and a fixed two–qubit measurement m is performed.

In the quantum realisation, each local preparation instance corresponds to a single–qubit pure state. For Alice, we take

$$\begin{aligned} \rho_0^a &= |0\rangle\langle 0|, & \rho_1^a &= |1\rangle\langle 1|, \\ \rho_0^{a'} &= |+\rangle\langle +|, & \rho_1^{a'} &= |-\rangle\langle -|. \end{aligned} \tag{13}$$

and similarly for Bob,

$$\begin{aligned} \rho_0^b &= |0\rangle\langle 0|, & \rho_1^b &= |1\rangle\langle 1|, \\ \rho_0^{b'} &= |+\rangle\langle +|, & \rho_1^{b'} &= |-\rangle\langle -|. \end{aligned} \tag{14}$$

where $|\pm\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}$. Thus the joint preparations are product states,

$$\rho_{ij}^{\{x,y\}} = \rho_i^x \otimes \rho_j^y, \quad \text{for all } x \in \{a, a'\} \text{ and } y \in \{b, b'\}. \tag{15}$$

For example, for $\Gamma = \{a, b\}$ the four joint states are

$$|00\rangle\langle 00|,\ |01\rangle\langle 01|,\ |10\rangle\langle 10|,\ |11\rangle\langle 11|, \tag{16}$$

while for $\Gamma = \{a', b'\}$ they are

$$|{+}{+}\rangle\langle{+}{+}|,\ |{+}{-}\rangle\langle{+}{-}|,\ |{-}{+}\rangle\langle{-}{+}|,\ |{-}{-}\rangle\langle{-}{-}|, \tag{17}$$

and similarly for the mixed contexts $\{a, b'\}$ and $\{a', b\}$. A fixed two–qubit POVM $\{M_k\}_k$ (the PBR measurement) is performed, defined by a four-outcome POVM whose effects are proportional to the projectors onto the states

$$\begin{aligned} |\xi_1\rangle &= |01\rangle + |10\rangle, \\ |\xi_2\rangle &= |0-\rangle + |1+\rangle, \\ |\xi_3\rangle &= |+1\rangle + |-0\rangle, \\ |\xi_4\rangle &= |+-\rangle + |-+\rangle. \end{aligned} \tag{18}$$

The empirical model consists of the conditional distributions

$$E_{m|\Gamma}(k \mid i, j) = \mathrm{Tr}\big[M_k\, \rho_{ij}^\Gamma\big]. \tag{19}$$

The empirical model E takes the explicit form

$$E = \left(\begin{array}{cccc|cccc|cccc|cccc} 0 & \frac{1}{2} & \frac{1}{2} & 0 & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & 0 & \frac{1}{2} & \frac{1}{2} & 0 & \frac{1}{2} & 0 & 0 & \frac{1}{2} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{2} & 0 & 0 & \frac{1}{2} & 0 & \frac{1}{2} & \frac{1}{2} & 0 & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \\ \frac{1}{2} & 0 & 0 & \frac{1}{2} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & 0 & \frac{1}{2} & \frac{1}{2} & 0 \end{array}\right), \tag{20}$$

where the columns are ordered in four consecutive blocks corresponding to the contexts $\{a, b\}$, $\{a, b'\}$, $\{a', b\}$, $\{a', b'\}$, with lexicographic order $(0, 0), (0, 1), (1, 0), (1, 1)$ within each block. Rows are ordered by outcomes $o \in \{1, 2, 3, 4\}$.$\square$

We now make precise what it would mean for observed preparation statistics to arise from a single globally defined preparation model.

A putative global model assigns outcome statistics to every fully specified preparation instance. Such a model is described by a stochastic matrix

$$D_{m|Y}(o \mid \sigma_Y), \qquad \sigma_Y \in I^Y, \tag{21}$$

where σ_Y denotes a complete assignment of preparation instances to all sources in Y. At this stage, $D_{m|Y}$ is purely hypothetical: its role is to represent a candidate global explanation whose compatibility with the observed data is to be assessed.

Relating this global description to experimentally accessible statistics requires specifying how a partial assignment of preparation instances on a source context $\Gamma \subseteq Y$ is completed to a full assignment on Y. Operationally, this corresponds to specifying how preparation control for sources not present in Γ is filled in. We encode this refinement by a stochastic *extension map*

$$S_{Y|\Gamma}(\sigma_Y \mid \sigma_\Gamma), \tag{22}$$

which assigns a distribution over global instance assignments σ_Y compatible with the partial assignment σ_Γ.

Given a global response model $D_{m|Y}$ and an extension map $S_{Y|\Gamma}$, the predicted outcome statistics for the context Γ are obtained by averaging over the unobserved preparation instances,

$$E_{m|\Gamma}(o \mid \sigma_\Gamma) = \sum_{\sigma_Y} D_{m|Y}(o \mid \sigma_Y)\, S_{Y|\Gamma}(\sigma_Y \mid \sigma_\Gamma). \tag{23}$$

This may be viewed as a law of total probability, with the requirement that the global statistics depend only on the global preparation instance σ_Y, and not on the particular source context Γ through which it is accessed — i.e., $D_{m|Y,\Gamma}(o \mid \sigma_Y, \sigma_\Gamma) = D_{m|Y}(o \mid \sigma_Y)$. This equation plays an analogous formal role to the restriction maps used in the measurement scenario: local statistics are recovered from a global object by combining it with a fixed linear map.

In contrast to marginalisation, stochastic extension is not canonical. There may exist many inequivalent ways of completing a partial preparation assignment to a larger set of sources. Before asking whether a global response model exists, we must therefore specify what it means for such a completion procedure to be operationally well defined. We impose minimal structural requirements on admissible families of extension maps, expressing the idea that extension should represent a genuine refinement of preparation control rather than an arbitrary stochastic rewriting.

First, we require *extension-respecting*: completing a preparation assignment must not alter the instances that have already been specified. Second, we require *input independence*: the statistical rule used to fill in missing sources depends only on which sources are missing, and not on the particular values of those already specified. This ensures that the statistics assigned to unspecified sources do not depend on the values of those already fixed, and prevents contextual behaviour from being artificially encoded in the completion rule itself. Third, we require *compositionality*: extending in stages yields the same distribution as extending in a single step, ensuring that refinement is independent of arbitrary choices of intermediate implementation. Together, these conditions express that stochastic extension behaves as a coherent refinement procedure, independent of staging and of the values of previously fixed preparation instances.

To make these conditions precise, fix a finite set of sources Y and a finite instance set I. For all $W \subseteq V \subseteq U \subseteq Y$, let

$$S_{U|V} : I^V \rightsquigarrow I^U \tag{24}$$

be a stochastic *extension map*, i.e. a column-stochastic matrix with entries $S_{U|V}(\sigma_U \mid \sigma_V)$. Such a map specifies how a partial assignment of preparation instances on V is refined to an assignment on U.

Definition 1 (Extension-respecting)**.** *A family* $\{S_{U|V}\}$ *is* extension-respecting *if*

$$S_{U|V}(\sigma_U \mid \sigma_V) > 0 \implies (\sigma_U)|_V = \sigma_V. \tag{25}$$

Definition 2 (Input independence)**.** *A family* $\{S_{U|V}\}$ *is* input independent *if for every inclusion* $V \subseteq U$ *there exists a distribution*

$$s_{U\setminus V} \in \Delta(I^{U\setminus V}) \tag{26}$$

such that for all $\sigma_V \in I^V$ *and* $\sigma_U \in I^U$,

$$S_{U|V}(\sigma_U \mid \sigma_V) = \mathbf{1}[(\sigma_U)|_V = \sigma_V]\; s_{U\setminus V}((\sigma_U)|_{U\setminus V}). \tag{27}$$

Definition 3 (Compositionality)**.** *A family* $\{S_{U|V}\}$ *is* compositional *if for all* $W \subseteq V \subseteq U$,

$$S_{U|W} = S_{U|V}\, S_{V|W}, \tag{28}$$

with ordinary composition of stochastic matrices.

Note that input independence implies extension-respecting, but we state the latter separately to emphasise the interpretation of stochastic extension as a genuine refinement of specified preparation instances.

These structural requirements already rigidly constrain the form of admissible extension maps.

Theorem 1 (Product-form extension theorem)**.** *Assume* $\{S_{U|V}\}$ *is extension-respecting, input independent in the form* (27)*, and compositional in the sense of* (28)*. Then there exist single-site distributions* $\{s_p \in \Delta(I)\}_{p\in Y}$ *such that for all* $V \subseteq U \subseteq Y$,

$$S_{U|V}(\sigma_U \mid \sigma_V) = \mathbf{1}[(\sigma_U)|_V = \sigma_V] \prod_{p\in U\setminus V} s_p(\sigma_p), \tag{29}$$

in particular, for extensions from context $\Gamma \in \mathcal{S}$ *to* Y,

$$S_{Y|\Gamma}(\sigma_Y \mid \sigma_\Gamma) = \mathbf{1}[(\sigma_Y)|_\Gamma = \sigma_\Gamma] \prod_{p\in Y\setminus\Gamma} s_p(\sigma_p). \tag{30}$$

Proof. Fix $W \subseteq V \subseteq U$. Using (27) in the compositionality condition (28) gives, for all $\sigma_W \in I^W$ and $\sigma_U \in I^U$,

$$\begin{aligned}&\mathbf{1}[(\sigma_U)|_W = \sigma_W]\; s_{U\setminus W}(\sigma_U|_{U\setminus W})\\ &\quad= \sum_{\sigma_V\in I^V} \mathbf{1}[(\sigma_U)|_V = \sigma_V]\; \mathbf{1}[(\sigma_V)|_W = \sigma_W]\\ &\qquad\qquad \times s_{U\setminus V}(\sigma_U|_{U\setminus V})\; s_{V\setminus W}(\sigma_V|_{V\setminus W}).\end{aligned} \tag{31}$$

The indicator functions force $\sigma_V = \sigma_U|_V$, so the sum collapses to

$$\begin{aligned}&\mathbf{1}[(\sigma_U)|_W = \sigma_W]\; s_{U\setminus W}((\sigma_U)|_{U\setminus W})\\ &\quad= \mathbf{1}[(\sigma_U)|_W = \sigma_W]\; s_{U\setminus V}((\sigma_U)|_{U\setminus V})\; s_{V\setminus W}((\sigma_U)|_{V\setminus W}),\end{aligned} \tag{32}$$

and hence

$$s_{U\setminus W}(\sigma_{U\setminus V}, \sigma_{V\setminus W}) = s_{U\setminus V}(\sigma_{U\setminus V})\; s_{V\setminus W}(\sigma_{V\setminus W}), \tag{33}$$

for all $\sigma_{U\setminus V} \in I^{U\setminus V}$ and $\sigma_{V\setminus W} \in I^{V\setminus W}$.

Now fix $V \subseteq U$ and choose any chain

$$V = V_0 \subset V_1 \subset \cdots \subset V_n = U, \qquad V_{i+1} = V_i \cup \{p_{i+1}\}. \tag{34}$$

Repeated application of (33) yields

$$s_{U\setminus V}(\sigma_{U\setminus V}) = \prod_{i=1}^{n} s_{V_i\setminus V_{i-1}}(\sigma_{p_i}). \tag{35}$$

Defining $s_{p_i} := s_{V_i\setminus V_{i-1}} \in \Delta(I)$, this factorisation is independent of the chosen chain, since (33) implies $s_{A\cup B} = s_A\, s_B$ for disjoint A, B. Substituting this product form back into (27) establishes (29). ∎

The admissibility conditions, therefore, constrain extension maps to product form, but they do not fix them uniquely: different choices of single-site distributions $\{s_p\}_{p\in Y}$ yield inequivalent admissible families $\{S_{U|V}\}$.

Given an empirical model $\{E_{m|\Gamma}\}_{\Gamma\in\mathcal{S}}$ and an admissible family of extension maps $\{S_{U|V}\}$, we are interested in those extensions that induce consistent restrictions on overlaps of source contexts. Operationally, this expresses preparation compatibility: the effective description on $\Gamma\cap\Gamma'$ should not depend on whether it is accessed via Γ or via Γ'.

Definition 4 (Preparation compatibility)**.** *Let $\{E_{m|\Gamma}\}_{\Gamma\in\mathcal{S}}$ be an empirical model and let $\{S_{U|V}\}$ be an admissible family of extension maps. We say that the pair (E,S) satisfies* preparation compatibility *if for all $\Gamma,\Gamma'\in\mathcal{S}$,*

$$E_{m|\Gamma}\, S_{\Gamma|\Gamma\cap\Gamma'} = E_{m|\Gamma'}\, S_{\Gamma'|\Gamma\cap\Gamma'}. \tag{36}$$

Equation (36) states that the induced outcome statistics on the shared sources $\Gamma\cap\Gamma'$ are independent of which larger source context is used. In the sheaf-theoretic measurement framework, so-called generalized no-signalling expresses the compatibility of probability distributions on overlaps of measurement contexts [5]. In Bell scenarios this coincides with causal no-signalling, but in general it is a structural consistency condition on marginal distributions rather than a dynamical constraint. Eq. (36) plays the analogous mathematical role for preparation scenarios: it enforces compatibility of the induced statistics on overlapping source contexts, now mediated by extension maps rather than marginalisation. Unlike the measurement case, however, this condition does not correspond to a physical no-signalling requirement, but simply expresses coherence between the empirical model and the chosen extension structure.

With these ingredients in place, we can now state the preparation-dual notion of contextuality. The central question is whether the observed statistics on each source context can be reproduced by a single global response model for some admissible extension structure.

An empirical model specifies, for each source context $\Gamma\in\mathcal{S}$, a family of conditional outcome distributions $E_{m|\Gamma}$. Given such an empirical model, we consider admissible families of extension maps $\{S_{U|V}\}$ that satisfy preparation compatibility with E in the sense of Eq. (36). We say that the empirical model is preparation noncontextual if there exists such a compatible extension structure admitting a single global response matrix reproducing the observed context-wise statistics via Eq.(23). If at least one preparation-compatible extension structure exists but none admits such a global response matrix, the empirical model is preparation contextual.

Definition 5 (Preparation contextuality)**.** *An empirical model $\{E_{m|\Gamma}\}_{\Gamma\in\mathcal{S}}$ is* preparation noncontextual *if there exists an admissible family of stochastic extension maps $\{S_{U|V}\}$ such that*

1. *(E,S) satisfies preparation compatibility, and*
2. *there exists a global response matrix $D_{m|Y}$ satisfying*

$$E_{m|\Gamma} = D_{m|Y}\, S_{Y|\Gamma} \qquad \forall\,\Gamma\in\mathcal{S}.$$

If the empirical model admits at least one admissible extension family satisfying preparation compatibility but none admits such a global response matrix, then the empirical model is said to be preparation contextual*.*

In the remainder of this paper, preparation contextuality will always be understood in this sense.

It is often useful to consider a stronger, possibilistic notion of preparation contextuality, obtained by retaining only the support of the empirical statistics. Given an empirical model E, its possibilistic reduction $\overline{E}$ is obtained by mapping all nonzero probabilities to one. One defines possibilistic noncontextuality analogously, by replacing stochastic matrices with Boolean matrices and ordinary matrix multiplication with Boolean multiplication. The two notions are related as follows.

Theorem 2 (Possibilistic implies probabilistic contextuality)**.** *If an empirical model is possibilistically preparation contextual, then it is also preparation contextual in the probabilistic sense.*

Proof. Any probabilistic global response model induces a possibilistic one by taking supports: replacing each positive entry by 1 and each zero by 0. Hence the existence of a probabilistic noncontextual model would imply the existence of a possibilistic one. The contrapositive yields the claim. ■

Preparation-contextuality of Quantum PBR-type preparation scenario. We illustrate preparation contextuality by demonstrating that the PBR-type preparation scenario constructed earlier is preparation contextual in the sense of our definition. The full proof is given in Appendix B; here we present the essential structure of the argument.

Recall that this empirical model E takes the explicit form (see Eq. (20) for index convention)

$$E = \left(\begin{array}{cccc|cccc|cccc|cccc} 0 & \frac{1}{2} & \frac{1}{2} & 0 & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & 0 & \frac{1}{2} & \frac{1}{2} & 0 & \frac{1}{2} & 0 & 0 & \frac{1}{2} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{2} & 0 & 0 & \frac{1}{2} & 0 & \frac{1}{2} & \frac{1}{2} & 0 & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \\ \frac{1}{2} & 0 & 0 & \frac{1}{2} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & 0 & \frac{1}{2} & \frac{1}{2} & 0 \end{array}\right). \tag{37}$$

A key property of this empirical model is that, for each source context and each choice of preparation instances, exactly one measurement outcome is assigned zero probability. Equivalently, every empirical column contains a single forbidden outcome.

By Theorem 1, any admissible family of extension maps has product form,

$$S_{Y|\Gamma}(\sigma_Y \mid \sigma_\Gamma) = \mathbf{1}[(\sigma_Y)|_\Gamma = \sigma_\Gamma] \prod_{p\in Y\setminus\Gamma} s_p(\sigma_p), \tag{38}$$

for some choice of single-site distributions $\{s_p\}_{p\in Y}$. In particular, taking each s_p to be uniform yields an admissible extension structure that satisfies preparation compatibility with E. Thus, the scenario admits at least one preparation-compatible extension description.

To establish preparation contextuality, it therefore suffices to show that no global response matrix $D_{m|Y}$ can reproduce E via Eq. (23) for any admissible extension map.

Assume, for contradiction, that there exists a global response matrix $D_{m|Y}$ together with a product-form extension map reproducing the empirical statistics via Eq. (23). Because the extension map assigns a nonzero weight to every globally consistent refinement of a local preparation choice, each forbidden outcome at the empirical level propagates to a forbidden outcome at the level of the corresponding global preparation instances. Hence, the pattern of zeros in the empirical matrices imposes stringent support constraints on the columns of $D_{m|Y}$.

The crucial observation is that these constraints are incompatible with the requirement that every column of $D_{m|Y}$ represent a valid probability distribution. A parity-based combinatorial argument shows that, for any choice of single-site extension distributions $\{s_p\}$, at least one global preparation instance is forced to forbid all four measurement outcomes simultaneously, which is impossible. Hence, no global response model exists for any admissible preparation-compatible extension map, and the empirical model is preparation contextual.

Conceptually, the argument proceeds in three steps. First, the presence of a single forbidden outcome in each empirical column allows one to associate to every local preparation choice a unique outcome that cannot occur. Second, the product-form extension map ensures that whenever a global preparation instance is compatible with a given local choice, the corresponding forbidden outcome is also forbidden globally. This propagates the empirical zero constraints to the level of the global response matrix.

Finally, one examines the pattern of propagated constraints across the four source contexts. For global preparation instances of even parity, the forbidden outcomes arising from the four contexts are all distinct, forcing all four measurement outcomes to be absent simultaneously. For odd parity, fewer outcomes are forbidden, but the remaining support is still insufficient to satisfy consistency with all empirical columns. In both cases, one obtains a contradiction, showing that no global response matrix can exist. □

Discussion: We have introduced a preparation-dual notion of contextuality formulated as an obstruction to stochastic extension. In direct analogy with the sheaf-theoretic formulation of measurement contextuality [5], preparation contextuality arises when locally specified statistics cannot be extended to a single global response model compatible with all source contexts. In the measurement case, contextuality is the failure of marginal restriction to assemble into a global section. In the preparation-dual case developed here, contextuality is the failure of stochastic extension: although compatible extension structures exist, none admits a global response matrix $D_{m|Y}$ satisfying $E_{m|\Gamma} = D_{m|Y}\, S_{Y|\Gamma}$.

This formulation places preparation contextuality squarely within a linear-algebraic framework. The empirical model E, the extension maps S, and the global response matrix D are related by matrix equations, and preparation noncontextuality reduces to the existence of a stochastic matrix satisfying a finite family of linear constraints. In this respect, the present notion bears structural similarity to operational approaches to contextuality in which contextuality is characterised via constraints on stochastic maps.

A further direction is to clarify the relationship between this extension-based perspective and the COPE framework of Shahandeh, Doosti, and collaborators [8, 9], which formulates operational contextuality via constrained nonnegative matrix factorizations. Both approaches employ matrix descriptions, but the structural obstructions they identify are different. Determining how these viewpoints relate remains an open problem.

A deeper structural question concerns the existence of a preparation-analogue of Fine's theorem [6]. In the measurement case, Abramsky and Brandenburger showed that the existence of a global section is equivalent to the existence of a deterministic hidden-variable model [5]. This result provides a precise bridge between the sheaf-theoretic and ontological-model formulations of contextuality. An analogous theorem in the preparation setting would establish that stochastic extendability is equivalent to the existence of a preparation-noncontextual ontological model. Such a result would clarify the exact relationship between the sheaf-theoretic preparation-dual notion introduced here and operational preparation noncontextuality in the sense of Spekkens [3]. Establishing (or refuting) such an equivalence is an important direction for future work.

More broadly, the extension-based perspective suggests several avenues for further investigation. First, one may seek a cohomological characterisation of preparation contextuality paralleling developments for measurement contextuality [10]. Second, the linear-algebraic structure of the obstruction invites algorithmic analysis: preparation noncontextuality reduces to a feasibility problem for stochastic matrices subject to linear constraints. Third, it would be natural to investigate whether a unified categorical or sheaf-theoretic treatment of preparation and measurement contextuality could illuminate deeper dualities between restriction and extension.

Preparation contextuality thus isolates a particularly robust form of nonclassicality. It captures a structural incompatibility between local preparation descriptions and global response models. Understanding how this obstruction relates to ontological, operational, and rank-based formulations of contextuality remains an important problem for future research.

[1] E. S. Simon Kochen, Indiana Univ. Math. J. **17**, 59 (1968).

[2] J. S. Bell, Physics **1**, 195 (1964).

[3] R. W. Spekkens, Phys. Rev. A **71**, 052108 (2005).

[4] S. Abramsky and L. Hardy, Physical Review A **85**, 062114 (2012).

[5] S. Abramsky and A. Brandenburger, New Journal of Physics **13**, 113036 (2011).

[6] A. Fine, Physical Review Letters **48**, 291 (1982).

[7] M. F. Pusey, J. Barrett, and T. Rudolph, Nature Physics **8**, 475 (2012).

[8] F. Shahandeh and M. Doosti, Physical Review A **98**, 052117 (2018).
[9] F. Shahandeh and M. Doosti, Quantum **5**, 533 (2021).
[10] S. Abramsky, S. Mansfield, and R. S. Barbosa, Electronic Proceedings in Theoretical Computer Science **236**, 1 (2017).

Appendix A: PBR-type example: full proof

In this appendix we give a complete technical proof that the quantum model discussed in the main text is preparation contextual in the sense of our definition. By definition, preparation contextuality requires that the empirical model admit at least one admissible extension family satisfying preparation compatibility, but that no admissible extension family admit a global response matrix reproducing the empirical statistics. We first exhibit a preparation-compatible extension structure by taking the single-site distributions s_p to be uniform. We then show that no admissible extension map admits even a possibilistic global response matrix. Preparation contextuality in the probabilistic sense then follows from the general result established in the main text that possibilistic contextuality implies probabilistic contextuality for a fixed admissible extension map.

We work with the set of sources

$$Y = \{a, b, a', b'\},$$

each admitting binary instances $\{0, 1\} = \{0, 1\}$. The implementable source contexts are

$$\mathcal{S} = \big\{\{a, b\}, \{a, b'\}, \{a', b\}, \{a', b'\}\big\}.$$

A global instance assignment is a tuple

$$\sigma_Y = (a, b, a', b') \in \{0, 1\}^4.$$

We fix a concrete extension structure by taking s_p to be the uniform distribution on $\{0, 1\}$ for each $p \in Y$, and let S^{unif} denote the corresponding product-form extension map:

$$S^{\text{unif}}_{Y|\Gamma}(\sigma_Y \mid \sigma_\Gamma) = \mathbf{1}[(\sigma_Y)|_\Gamma = \sigma_\Gamma] \prod_{p \in Y \setminus \Gamma} \tfrac{1}{2}. \tag{A.1}$$

By Theorem 1, this map is admissible. A direct computation using the explicit form of the empirical model below shows that

$$E_{m|\Gamma} S^{\text{unif}}_{\Gamma|\Gamma\cap\Gamma'} = E_{m|\Gamma'} S^{\text{unif}}_{\Gamma'|\Gamma\cap\Gamma'}$$

for all overlapping contexts Γ, Γ'. Hence (E, S^{unif}) satisfies preparation compatibility. To establish preparation contextuality, it therefore suffices to show that no global response matrix exists for any admissible product-form extension map.

Columns of empirical matrices are ordered in four consecutive blocks corresponding to the contexts $\{a, b\}$, $\{a, b'\}$, $\{a', b\}$, $\{a', b'\}$, with lexicographic order $(0, 0), (0, 1), (1, 0), (1, 1)$ within each block. Rows are ordered by outcomes $o \in \{1, 2, 3, 4\}$, and rows of extension and global response matrices are ordered lexicographically by $\sigma_Y \in \{0, 1\}^4$.

With these conventions, the empirical model E takes the explicit form

$$E = \left(\begin{array}{cccc|cccc|cccc|cccc}
0 & \frac{1}{2} & \frac{1}{2} & 0 & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{2} & 0 & 0 & \frac{1}{2} \\
\frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & 0 & \frac{1}{2} & \frac{1}{2} & 0 & \frac{1}{2} & 0 & 0 & \frac{1}{2} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \\
\frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{2} & 0 & 0 & \frac{1}{2} & 0 & \frac{1}{2} & \frac{1}{2} & 0 & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \\
\frac{1}{2} & 0 & 0 & \frac{1}{2} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & 0 & \frac{1}{2} & \frac{1}{2} & 0
\end{array}\right).$$

Each column contains exactly one zero entry. Let $\overline{E}$ denote the possibilistic reduction of E, obtained by mapping $0 \mapsto 0$ and $(0, 1] \mapsto 1$. Let $\overline{S}^{\text{unif}}$ denote the possibilistic reduction of S^{unif}.

Since every column of $\overline{E}$ has a unique zero entry, there exists a unique function

$$\varphi : \{1, \ldots, 16\} \to \{1, 2, 3, 4\}$$

such that $\overline{E}(\varphi(j), j) = 0$ for each column j.

If $\overline{E} = \overline{D}\,\overline{S}^{\text{unif}}$, then Boolean matrix multiplication implies

$$\overline{S}^{\text{unif}}(k, j) = 1 \;\Rightarrow\; \overline{D}(\varphi(j), k) = 0.$$

For each global assignment index k, define

$$J(k) := \{j : \overline{S}^{\text{unif}}(k, j) = 1\}.$$

Then

$$\operatorname{supp}(\overline{D}(:, k)) \subseteq \{1, 2, 3, 4\} \setminus \varphi(J(k)), \qquad |\operatorname{supp}(\overline{D}(:, k))| \le 4 - |\varphi(J(k))|.$$

Under the uniform choice of $\{s_p\}$, every global assignment $\sigma_Y \in \{0,1\}^4$ has nonzero weight, so for each index k, the set $J(k)$ contains one column from each context block. A direct inspection of φ shows that if the tuple (a, b, a', b') corresponding to k has even parity,

$$a \oplus b \oplus a' \oplus b' = 0,$$

then $\varphi(J(k)) = \{1, 2, 3, 4\}$, whereas if the parity is odd,

$$a \oplus b \oplus a' \oplus b' = 1,$$

then $|\varphi(J(k))| = 2$.

In the even-parity case, the bound above forces $\overline{D}(:, k) = 0$, contradicting the requirement that columns of $\overline{D}$ be nonempty. In the odd-parity case, the bound implies $|\operatorname{supp}(\overline{D}(:, k))| \leq 2$, while Boolean consistency with the columns of $\overline{E}$ requires at least three supported outcomes, again yielding a contradiction.

Thus no possibilistic matrix $\overline{D}$ with nonempty columns satisfies

$$\overline{E} = \overline{D}\, \overline{S}^{\text{unif}}.$$

By the general implication established in the main text, this rules out probabilistic global response matrices satisfying

$$E = D\, S^{\text{unif}}.$$

Hence the empirical model is preparation contextual.