

Shadows and subsystems of generalized probabilistic theories: when tomographic incompleteness is not a loophole for contextuality proofs

David Schmid^{1,2}, John H. Selby¹, Vinicius P. Rossi¹, Roberto D. Baldijão^{1,2}, and Ana Belén Sainz^{1,3}

¹International Centre for Theory of Quantum Technologies, University of Gdańsk, 80-309 Gdańsk, Poland

²Perimeter Institute for Theoretical Physics, N2L 2Y5 Waterloo, Canada

³Basic Research Community for Physics e.V., Germany

It is commonly believed that failures of tomographic completeness undermine assessments of nonclassicality in noncontextuality experiments. In this work, we study how such failures can indeed lead to mistaken assessments of nonclassicality. We then show that proofs of the failure of noncontextuality are robust to a very broad class of failures of tomographic completeness, including the kinds of failures that are likely to occur in real experiments. We do so by showing that such proofs actually rely on a much weaker assumption that we term *relative tomographic completeness*: namely, that one's experimental procedures are tomographic *for each other*. Thus, the failure of noncontextuality can be established even with coarse-grained, effective, emergent, or virtual degrees of freedom. This also implies that the existence of a deeper theory of nature (beyond that being probed in one's experiment) does not in and of itself pose any challenge to proofs of nonclassicality. To prove these results, we first introduce a number of useful new concepts within the framework of generalized probabilistic theories (GPTs). Most notably, we introduce the notion of a *GPT subsystem*, generalizing a range of preexisting notions of subsystems (including those arising from tensor products, direct sums, decoherence processes, virtual encodings, and more). We also introduce the notion of a *shadow* of a GPT fragment, which captures the information lost when one's states and effects are unknowingly not tomographic for one another.

Contents			
1 Introduction	2	7 Future directions	26
2 Generalized probabilistic theory preliminaries	3	Acknowledgements	27
3 GPT subsystems	10	References	27
4 Shadows of GPT fragments	13	A Proof of Lemma 16	31
5 The effect of shadow maps on simplex embeddability	17	B Proof of Theorem 17	32
6 Implications for noncontextuality	21	C Proof of Lemma 18	35
		D Proof of Theorem 20	35
		E Proof of Theorem 22	38

David Schmid: davidschmid10@gmail.com

John H. Selby: john.h.selby@gmail.com

1 Introduction

A gold standard universally applicable method for demonstrating that a given theory, experiment, or phenomenon resists classical explanation is to prove that it cannot be reproduced in any generalized noncontextual ontological model [1]. This approach has been motivated extensively in the literature [1–6], most notably by a version of Leibniz’s principle [2] and by its equivalence to the existence of a positive quasiprobabilistic representation [4, 5] and to the natural notion of classical-explainability in the framework of generalized probabilistic theories [5, 6]. Studies characterizing exactly which phenomena constitute proofs of nonclassicality—of the failure of generalized noncontextuality—have been carried out in the contexts of computation [7, 8], state discrimination [9–12], interference [13–15], compatibility [16–18], uncertainty relations [19], metrology [20], thermodynamics [20, 21], weak values [22, 23], coherence [24, 25], quantum Darwinism [26], information processing and communication [27–33], cloning [34], broadcasting [35], and Bell [36, 37] and Kochen-Specker scenarios [38–44]. In each of these contexts, contextuality constitutes a resource for generating novel kinds of correlations, which in many cases translates into advantages for information processing and other tasks. Henceforth, we will refer to generalized noncontextuality simply as noncontextuality, and throughout this work we will use the terms *contextuality* and *non-classicality* interchangeably.

It is often claimed that assessments of noncontextuality rest on the assumption that one’s laboratory procedures are *tomographically complete* [1, 45, 46]: that is, that they allow one to fully characterize the properties of the physical system in question (the causal mediary between one’s preparations and measurements). However, we will show here that assessments of noncontextuality rest on a weaker assumption, namely that one’s preparations and measurements are sufficient to fully characterize *each other*, regardless of whether they are tomographic for all properties of the true physical system. We term this the *relative tomographic completeness* assumption.

Consequently, we prove that assessments of nonclassicality will be valid even if

one’s laboratory operations are far from tomographically complete for any given systems in one’s experiment or theory, provided that the preparations and measurements used in one’s proof of nonclassicality probe *the same degree of freedom*.

To do this, we must first formalize what we mean by ‘degree of freedom’. Inspired by Refs. [6, 47], we do so within a framework for possible theories known as the framework of generalized probabilistic theories (GPTs) [48–50]. This framework captures quantum theory, classical theory, and a vast array of other theories as special cases, and allows for a unified study of all of these. While the notion of a system within a GPT is completely standard, we will here introduce the notion of a GPT *subsystem*, and then will demonstrate how this captures the relevant notion of a ‘degree of freedom’ for the study of noncontextuality. Our starting point will be a formalisation of GPTs which makes their representation more flexible, and that subsumes standard representations of GPTs (see Sec. 2.2).

Our notion of GPT subsystems is extremely general. It captures tensor product and direct sum subsystems in the usual senses defined in quantum theory (and in arbitrary GPTs), but it also captures such wide-ranging examples as virtual subsystems, stabilizer subsystems, rebit subsystems, decohered subsystems, and so on. Indeed, according to our definition, any fragment of a GPT which itself satisfies all the mathematical properties of a GPT constitutes a GPT subsystem of that GPT. This notion aims to capture what it means for a GPT system to “live inside” another, or to be “explainable by” another—in both cases, with all its GPT structure intact: if A is a GPT subsystem of B , then B can explain A through the notion of GPT explainability introduced by Ref. [51]. The main concepts that we introduce pertaining to GPT systems and subsystems, and how they relate to each other, are summarised in Fig. 4 in Sec. 4.5.

We then show that the only assumption needed for accurate assessments of nonclassicality is that the states and effects in one’s experiment are tomographic for a GPT subsystem of the true GPT. This is simply another way of stating the relative tomographic completeness assumption.

If the states and effects do not satisfy this property, and so are not relatively tomographic,

then one can indeed be led to incorrectly conclude that an experiment exhibits nonclassicality. Indeed, this is the sole nontrivial loophole in a state-of-the-art proof of the failure of generalized noncontextuality [45, 52]. (Note in particular that there are no loopholes due to detector inefficiencies or superdeterminism [16].) Yet this assumption, and what happens when it fails, has only been studied explicitly in a single prior work, Ref. [53] (which we comment on further in the conclusions).

In this work, we provide rigorous tools for studying this potential loophole, once again by introducing relevant concepts within the framework of generalized probabilistic theories.

Consider an experiment that accesses only some fragment of the preparations and measurements on one's physical system. If a given pair of distinct states (measurement effects) in one's experiment cannot be tomographically characterized by the effects (states) in the experiment, then those states (effects) will appear to be exactly identical within the context of the given experiment. This will then lead the experimenter to mistakenly treat distinct states (effects) as identical. Consequently, the GPT description one will give to the experiment will be incorrect—it will in fact be given by a particular informationally lossy function of the true GPT fragment describing the scenario. We will call this informationally lossy map the *shadow map*.

If (and only if) the assumption of relative tomography is not satisfied, the shadow map acts nontrivially, distorting the true GPT description of one's experiment. This distortion *can*, but does not necessarily, cause one to reach incorrect conclusions about classicality of the experiment. In the final portion of this work, we begin the project of characterizing under what conditions the failure of relative tomography does indeed lead to mistaken assessments of nonclassicality.

2 Generalized probabilistic theory preliminaries

2.1 State spaces, effect spaces, systems, and embeddings

Generalised probabilistic theories [48–50] (GPTs) provide a framework for studying operational theories at the level of the operational statistics they can generate. When defined in full generality,

GPTs have a rich compositional structure, but we will focus here on preparations and measurements on a single system within the theory. Our presentation is not entirely standard, but it has some advantages over prior presentations and introduces some useful new concepts. Like most presentations, we will here consider only finite-dimensional systems.

To begin, we will consider the structure of states in GPTs.

Definition 1 (GPT state spaces). *The states s for a given GPT system G are represented by vectors in a (finite dimensional) real vector space U_G . Among these state-vectors, there is a privileged compact convex subset $\bar{\Omega}_G$ of normalised states, where it must be the case that $\mathbf{0} \notin \text{AffSpan}[\bar{\Omega}_G]$, where $\mathbf{0}$ is the null state (the zero vector). The state space Ω_G of a GPT also includes subnormalized states, and so is the convex hull of the normalised states with the null state,*

$$\Omega_G := \text{Conv}[\mathbf{0} \cup \bar{\Omega}_G]. \quad (1)$$

Note that (unlike in many approaches to GPTs) we will not assume that $\text{Span}[\Omega_G] = U_G$, as this assumption is naturally violated when one starts to consider GPT subsystems, as we will do here.

We depict state s of GPT system G diagrammatically by

$$\begin{array}{c} |G \\ \hline s \end{array} \in \Omega_G. \quad (2)$$

We are here viewing GPT state spaces as objects in their own right, rather than merely as part of the specification of a GPT. That is, they are not merely a set of labels for computing probabilities, but more importantly contain a representation of the operational identities among states. An *operational identity* among GPT processes is simply a linear equality; for states on a single system¹, the general form of an operational identity is

$$\sum_{s \in \Omega_G} \alpha_s s = \mathbf{0} \quad (3)$$

for $\alpha_s \in \mathbb{R}$. The set of all such operational identities² encodes (and is encoded by) the convex

¹In general, operational identities can have much more structure; see Refs. [37, 46].

²This set can be easily computed [54] and can be described as a subspace of \mathbb{R}^{Ω_G} .

geometry of the state space, which in turn is a critical aspect of the structure of a given theory. As a simple example of the importance of this convex structure, note that mixed states in classical probability theory have a unique decomposition into pure states, while mixed quantum states admit of infinitely many decompositions into pure states.

We often want to consider more than one state space and maps between them. This may be because we are viewing them as state spaces belonging to two different systems within the same GPT and are considering physical transformations between them. Or, it could be that we are viewing them as state spaces belonging to two entirely different GPTs, in which case such maps would not be physical transformations, but rather representations of the states of one system within the state space of another.

Definition 2 (GPT state maps). *A state map between two state spaces Ω_G and Ω_H is a linear map*

$$\iota: \text{Span}[\Omega_G] \rightarrow \text{Span}[\Omega_H] \quad (4)$$

such that $\iota(\Omega_G) \subseteq \Omega_H$ ³.

A state map is said to be faithful if and only if it is injective, to be full (or surjective) if and only if it satisfies $\iota(\Omega_G) = \Omega_H$, and invertible if and only if it has a linear inverse which is itself a state map, i.e., such that $\iota^{-1}(\Omega_H) \subseteq \Omega_G$. Note that if a state map is both full and faithful then it is necessarily invertible.

Diagrammatically we denote state maps as

$$\begin{array}{c} H \\ | \\ \boxed{\iota} \\ | \\ G \end{array} \quad (5)$$

where we use the violet box to highlight the fact that we are not (necessarily) thinking of these as physical transformations. We represent their action on states via

$$\begin{array}{c} H \\ | \\ \boxed{\iota} \\ | \\ G \end{array} \quad :: \quad \begin{array}{c} G \\ \swarrow \searrow \\ s \end{array} \quad \mapsto \quad \begin{array}{c} H \\ | \\ \boxed{\iota} \\ | \\ G \end{array} \in \Omega_H, \quad (6)$$

where one reads this diagram from bottom to top as composition of linear maps.

³It is clear that GPT state maps are closed under sequential composition and that the identity linear map is a state map, so one can define the category **State** that has state spaces as objects and state maps as morphisms.

Because state maps are linear, they preserve operational identities. In addition, new operational identities are introduced if and only if one's state map is not faithful; that is, for nonfaithful ι , $\iota(\Omega_G)$ will exhibit equalities that are not present in Ω_G .

Note also that invertibility of a state map ι is a strictly stronger condition than invertibility of ι as a linear map, due to the condition that the inverse of ι must also be a state map. For instance, consider embedding a Bloch ball into another Bloch ball by a partially depolarizing map. This state map is invertible *as a linear map* (if it is not totally depolarizing), but its inverse is not a valid quantum channel (as it can map quantum states to outside the Bloch ball). Consequently the state map is not an invertible state map, since its linear inverse is not a state map. Similarly, note that surjectivity of a state map is a strictly stronger condition than surjectivity as a linear map, as the latter would merely demand that $\iota(\text{Span}[\Omega_G]) = \text{Span}[\Omega_H]$.

Two state spaces are isomorphic if there is an invertible state map between them. This is the physically meaningful notion of equivalence for state spaces of GPTs—one that demands that the convex structure (encoding the operational identities) is the same in the two state spaces. It is this notion of equivalence that ensures that it is sensible to define GPT state spaces in a way that is independent of the vector space dimension, as we have done. This is because embedding some state space $\Omega_G \subset U_G$ within a higher dimensional vector space W via some injective linear map $I_\Omega: U_G \rightarrow W$ gives a new state space $I_\Omega(\Omega_G) \subset W$ which is isomorphic to the original one. Up to this isomorphism, state spaces are insensitive to the vector space in which they are being represented. So, for example, the state space of a classical bit is the same whether or not one views it as living in a 2-dimensional vector space, or as embedded in (say) the 4-dimensional vector space of a qubit (we will see such an example later in Figure 2.)

Next we consider the structure of effects in GPTs.

Definition 3 (GPT effect spaces). *The effects e for a given GPT system G are represented by vectors in a (finite dimensional) real vector space V_G . The effect space, \mathcal{E}_G , is a set of such effect vectors that is necessarily convex and compact, that contains the null effect (the zero vector $\mathbf{0}$) and a*

We represent it as a dashed circle because it should not (necessarily) be interpreted as a state map or an effect map, but a map from state-effect pairs to the reals:

$$\langle p_G \rangle :: \left(\frac{\downarrow G}{\downarrow S}, \frac{\uparrow G}{\uparrow E} \right) \mapsto \langle p_G \rangle \in [0,1] \quad (15)$$

Putting the above together, we can define a system in a GPT as follows.

Definition 6 (GPT system). *A GPT system G is a triple of a state space Ω_G , an effect space \mathcal{E}_G , and a tomographic probability rule p_G :*

$$G := \left(\left\{ \frac{\downarrow G}{\downarrow S} \right\}_{s \in \Omega_G}, \left\{ \frac{\uparrow G}{\uparrow E} \right\}_{e \in \mathcal{E}_G}, \langle p_G \rangle \right). \quad (16)$$

We can then define the notion of a GPT system embedding as follows.

Definition 7 (GPT system embedding [47]). *A GPT system embedding⁶ from G to H is a pair of a state map $\iota: \Omega_G \rightarrow \Omega_H$ and an effect map $\kappa: \mathcal{E}_G \rightarrow \mathcal{E}_H$ which taken together preserve probabilities, so that $p_G(e, s) = p_H(\iota(s), \kappa(e))$ for all $s \in \Omega_G$ and $e \in \mathcal{E}_G$. Diagrammatically, it is a state and effect map satisfying*

$$\langle p_G \rangle = \langle p_H \rangle. \quad (17)$$

Note that both ι and κ must be faithful as a consequence of the fact that the GPT system has a tomographic probability rule. The embedding is said to be invertible if both ι and κ are moreover invertible.

In the language of Ref. [47], this is an exact unital embedding of one GPT into another.

If ι and κ are invertible state and effect maps from G to H , then (by definition) their inverses are

⁶It is straightforward to see that we can sequentially compose such embeddings by element-wise composition of the pair, and that any GPT system embeds into itself by the identity state and effect maps. We can therefore define the category **GPT-System** that has GPT systems as objects and GPT embeddings as morphisms.

state and effect maps from H to G . It follows that $(\iota^{-1}, \kappa^{-1})$ gives an embedding of H into G , since these maps necessarily also preserve probabilities:

$$\langle p_G \rangle = \langle p_H \rangle \implies \langle p_G \rangle = \langle p_H \rangle. \quad (18)$$

Consequently for any invertible GPT embedding from G into H with state and effect maps (ι, κ) , the inverse maps $(\iota^{-1}, \kappa^{-1})$ give a GPT embedding of H into G .

It is straightforward to check that GPT embeddings are transitive [47]: if we have an embedding of G into H by (ι, κ) and H into K by (δ, γ) , then the sequential composition of these, $(\delta \circ \iota, \gamma \circ \kappa)$, will be an embedding of G into K .

Two GPT systems are said to be isomorphic if and only if there is an invertible embedding between them.

Definition 8 (Equivalence of GPT systems). *Two GPT systems are equivalent (isomorphic) if and only if there is an invertible GPT system embedding (as in Definition 7) between them.*

This definition of equivalence subsumes the standard one found in the literature, as we show in the next section.

Note that if we simply embed the states and effects of a GPT system within some higher dimensional vector spaces via injective linear maps $I_\Omega: U_G \rightarrow W$ and $I_\mathcal{E}: V_G \rightarrow W'$, this defines an isomorphic GPT system, $(I_\Omega(\Omega_G), I_\mathcal{E}(\mathcal{E}_G), p_G(P_\Omega(\cdot), P_\mathcal{E}(\cdot)))$ where P_Ω is any left inverse of I_Ω and similarly for $P_\mathcal{E}$. This shows once again how our definitions are insensitive to the particular vector spaces in which one chooses to represent the states and effects.

2.2 Subsuming the standard notion of a GPT

Unlike in standard presentations of GPTs, our definition of a GPT system involves an explicit probability rule; moreover (as mentioned earlier), we do not assume that the states and effects on the system span the vector spaces in which they live. Our definition of a GPT system is not essentially distinct from that found in the literature, but at the representational level it is more flexible. For

instance, it subsumes two common but distinct kinds of representations of GPTs: those that represent effects as linear functionals on states, and those that represent effects as living in the same vector space as states. This is convenient, as it enables one to switch freely between these two descriptions, each of which is useful in certain contexts. We will see further advantages of our definitions when we come to the representation of GPT fragments (see the next section and the comment about accessible GPT fragments therein).

Theorem 9. *Every GPT system in our framework is isomorphic to one in each of the standard forms found in the literature. Conversely, every GPT in one of the standard forms found in the literature is an instance of a GPT system in our framework.*

Proof. First, let us show how to convert any given GPT system in our framework to a GPT in the standard forms found in the literature [48,49,55,56]. If one is given a GPT system whose states do not span the vector space U_G in which they live and whose effects do not span the vector space V_G in which they live, one first transforms this into an isomorphic GPT system in which the states live in a smaller vector space which they *do* span, and in which the effects live in a smaller vector space which they *do* span. One can do this, for instance, using the projection maps P_Ω and $P_{\mathcal{E}}$, where P_Ω is a projection $U_G \rightarrow \text{Span}[\Omega_G]$ and where $P_{\mathcal{E}}$ is a projection $V_G \rightarrow \text{Span}[\mathcal{E}_G]$; these give a GPT isomorphism between the two GPT systems.

One can then lump the probability rule together with the effects to transform them into linear functionals on $\text{Span}[\Omega_G]$, i.e., for each $e \in \mathcal{E}_G$ we can define a linear functional

$$\begin{array}{c} \triangle \\ \textcolor{teal}{e} \\ \textcolor{teal}{\text{I}}_G \\ \textcolor{teal}{\langle} \textcolor{teal}{p}_G \textcolor{teal}{\rangle} \\ \textcolor{teal}{\text{I}}_G \\ \triangle \end{array} : \Omega_G \rightarrow [0,1] \quad :: \quad \begin{array}{c} \textcolor{teal}{\text{I}}_G \\ \textcolor{teal}{\langle} \textcolor{teal}{s} \textcolor{teal}{\rangle} \\ \triangle \end{array} \quad \mapsto \quad \begin{array}{c} \triangle \\ \textcolor{teal}{e} \\ \textcolor{teal}{\text{I}}_G \\ \textcolor{teal}{\langle} \textcolor{teal}{p}_G \textcolor{teal}{\rangle} \\ \textcolor{teal}{\text{I}}_G \\ \triangle \end{array}. \quad (19)$$

This recovers the common view [55,56] of effects as living in the vector space dual to that for states—i.e., the view wherein effects are *defined* as linear functionals on states.

One can then transform this representation into the other common representation [48,49] of GPTs, where states and effects are vectors in the same vector space (as opposed to dual vector spaces). The equivalence between these two kinds

of representations is well known, and relies on choosing an (arbitrary) inner product and using the Riesz representation theorem to get an inner product representation.

Conversely, a given GPT in either of these common forms is already an instance of a GPT system as we have defined it. A GPT wherein effects are viewed as linear functionals is a special case of a GPT system as in our Definition 6, where the probability rule is given by the evaluation map (i.e., by directly applying the linear functional to the state: $p(s,e) = e(s)$). A GPT wherein states and effects live in the same vector space is also a special case of a GPT system as in our Definition 6, where the probability rule is given by the inner product: $p(s,e) = \langle e, s \rangle$. \square

Our notion of equivalence for GPT systems (Definition 8) also subsumes the standard one in the literature: that two GPTs are equivalent if one is generated from the other by acting an invertible linear map on its state space and the inverse of that map on the effect space, while keeping the probability rule fixed.

2.3 Fragments and fragment embeddings

A GPT, viewed as a theory, is understood as describing every physically possible process in some (possibly hypothetical) world. However, sometimes one is interested in giving a GPT description of some particular experiment or particular processes within the world. In such cases, one needs a more general notion which allows, for example, for the possibility that the set of states and set of effects that are of interest might not be tomographic for each other. This leads to the notion of a *GPT fragment*. A fragment that is not itself a GPT system is typically defined relative to some background GPT system.

Definition 10 (Fragment of a GPT). *A fragment f of a GPT system $G := (\Omega_G, \mathcal{E}_G, p_G)$ has a GPT state space defined by a subset $\bar{\Omega}_f \subseteq \bar{\Omega}_G$ of the normalized states in G . As with all GPT state spaces, the full state space is $\Omega_f := \text{Conv}[\mathbf{0} \cup \bar{\Omega}_f]$. It has a GPT effect space defined by a subset $\mathcal{E}_f \subseteq \mathcal{E}_G$ of effects from the full GPT. As with all effect spaces, \mathcal{E}_f is convex and compact, must include the null effect $\mathbf{0}$ and the unit effect $u_f := u_G$, and satisfy $\forall e \in \mathcal{E}_f \exists e^\perp \in \mathcal{E}_f$ s.t. $e + e^\perp = u_f$. Probabilities in the fragment are computed using the probability rule p_G , but with its domain restricted to (the span of)*

embedding map for states that fails to be injective, then one will have at least one pair of distinct states, $s \neq s'$, for which $\iota(s) = \iota(s')$.

We emphasize this point because these facts will be crucial for understanding shadow maps and their implications for assessing nonclassicality. As an extreme example, every fragment with a trivial probability rule (i.e., one which is constant on normalized states) is embeddable into a trivial one with only one state (the embedding can simply map every normalized state in the original fragment to this unique state).

Finally, the notion of equivalence for GPT fragments is the following (analogous to that for GPT systems).

Definition 12 (Equivalence of GPT fragments). *Two GPT fragments are equivalent (isomorphic) if and only if there is an invertible GPT fragment embedding (as defined in Definition 11) between them.*

Finally, note that a GPT fragment embedding necessarily takes normalized states to normalized states, as a consequence of probability preservation together with the preservation of the unit effect.

In earlier works [17, 57], two related concepts were introduced: *fragments* of a GPT and *accessible GPT fragments*. The distinction arises because the states and effects in a fragment of a given GPT need not span the vector space of the GPT, in which case one has a choice: one can either represent the processes in the fragment as living within the space they span, or one can represent them in the (larger) vector space of the original GPT. The former view was introduced in Ref. [17] under the name accessible GPT fragments, whereas the latter is simply called a fragment.

Because we have represented GPT state spaces and effects spaces in a way that is insensitive to the vector space in which they are represented, it follows that our approach (e.g., in Definition 10) subsumes both representations and treats them as equivalent (in the sense of Definition 12).

2.4 Nonclassicality as simplex-embeddability

Assessing the noncontextuality [1] of an operational prepare-measure scenario is equivalent to assessing whether or not the GPT

representation of that scenario can be embedded into a simplicial GPT, as was shown in Ref. [6]. Consequently, we can take simplex-embeddability as our notion of classical-explainability for a GPT. In doing so, we inherit all of the prior foundational motivations that have been given for generalized noncontextuality [1–3].

To formalize this notion of classical explainability, we must first formalize the notion of a simplicial GPT system.

Definition 13 (Simplicial GPT systems). *A simplicial GPT system is defined by the tuple $\Lambda := (\Delta, \Delta^*, p_\Lambda)$, where Δ is a state space which forms a simplex living in a finite-dimensional vector space \mathbb{R}^d ; Δ^* is the effect space which forms the logical dual of the simplex (that is, the set of vectors whose inner product with vectors in the simplex is between 0 and 1); and p_Λ is the Euclidean inner product in \mathbb{R}^d .*

Simplicial GPT systems represent strictly classical systems—those for which every state has a unique interpretation as a state of knowledge about some fundamental ontic states (which can be associated with the extremal states—the vertices of the simplex), and for which every logically possible measurement is allowed (including a single maximally informative measurement). They also embody the notion of classicality introduced by Leggett and Garg under the name macroscopic realism [58, 59].

A given GPT system or GPT fragment is classically explainable if and only if it can be embedded into a simplicial GPT system.

Definition 14 (Simplex embedding for a GPT fragment). *A GPT fragment f admits of a simplex embedding if there exists a simplicial GPT system Λ and a GPT fragment embedding $(\iota, \kappa): f \rightarrow \Lambda$.*

It follows from this definition that every fragment of a simplex-embeddable GPT system will also admit of a simplex embedding. On the other hand, impossibility of simplex embedding for a fragment necessarily implies impossibility of simplex embedding for the full GPT system of which the fragment is part.

This is the natural notion of classical explainability within the framework of GPTs.

3 GPT subsystems

The notion of GPT system embeddings leads to a natural and very general notion of *subsystems* within the framework of GPTs. This aims to capture what it means for a GPT system to “live inside” another, or to be “explainable by” another—in both cases, with all its GPT structure intact.

Definition 15 (GPT subsystems). *A subsystem of a GPT system G is a GPT system F , denoted $F \subseteq G$, such that there exists a GPT embedding (in the sense of Definition 11) of F into G .*

By definition, every GPT subsystem is itself a GPT system, and so constitutes a relatively tomographic fragment.

Two obvious but useful facts about GPT subsystems are the following, the first of which follows from transitivity of GPT embeddings:

Proposition 1. *The notion of a GPT subsystem is transitive, so that $F \subseteq G$ and $G \subseteq H$ implies $F \subseteq H$.*

Second, from the definition of simplex-embedding, it follows that:

Proposition 2. *A GPT system F is simplex embeddable if and only if it is a GPT subsystem of a simplicial system; that is, if and only if $F \subseteq \Lambda$ for some simplicial GPT system Λ .*

This definition of a GPT subsystem is general enough to subsume a variety of interesting cases:

1. a component F of a composite system $F \otimes G$ is a GPT subsystem of $F \otimes G$,
2. a component F of a direct sum system $F \oplus G$ is a GPT subsystem of $F \oplus G$,
3. virtual quantum systems (so that, e.g., a logical qubit is a GPT subsystem of the physical qubits it is encoded in),
4. stabilizer quantum systems [60, 61] are GPT subsystems of quantum systems of the same dimension (so that, e.g., a stabilizer qubit is a GPT subsystem of a qubit),
5. restricting the states and effects on a GPT system to those living in a linear subspace defines a GPT subsystem (so that, e.g., a rebit [62, 63] system is a GPT subsystem of a qubit),

6. the set of all processes that are the image of a decoherence process on a GPT system (so that, e.g., a classical bit is a GPT subsystem of a quantum bit).
7. any fragment of a GPT system G which is itself a valid GPT system constitutes a GPT subsystem of G .

We expand on these examples and prove that they are special cases of GPT subsystems below. Many of these examples are pictured in Figure 1.

One might wonder if the notion we have defined is *too* general, so much so that the terminology of ‘subsystem’ might be inappropriate. For instance, only examples 1, 2, 3, and (arguably) 6 above are conventionally considered subsystems. In the context of quantum theory, for example, the notion of a subsystem typically comes with algebraic structure. However, in the context of GPTs, a system generally has no structure beyond convex structure, and so it is reasonable to require that a “GPT subsystem” require only linear embeddability but nothing specific beyond that (such as preservation of any particular algebraic structure that may be present in special cases, like in quantum theory). As an example (that we return to later in this section), the rebit is typically not considered to be a *quantum* subsystem of a qubit, but it makes sense to consider it a *GPT* subsystem—and according to our definition, it is.

This notion of GPT subsystem is also motivated by its usefulness in the study of noncontextuality, as we will see in later sections.

Example 1. To see that a component GPT system F of a given composite GPT system $F \otimes G$ is always a GPT subsystem of $F \otimes G$, it suffices to note that the composition operation in any GPT, \otimes , is a bilinear map on states $\otimes : \Omega_F \times \Omega_G \rightarrow \Omega_{F \otimes G}$ and on effects $\otimes : \mathcal{E}_F \times \mathcal{E}_G \rightarrow \mathcal{E}_{F \otimes G}$, where $\otimes :: (u_F, u_G) \mapsto u_F \otimes u_G = u_{F \otimes G}$, and where $(e_F \otimes e_G)[(s_F \otimes s_b)] = e_F[s_F]e_G[s_G]$. (In this and the other examples in this section, we view effects as linear functionals on states.) These properties hold for the composition rule in arbitrary GPTs, including quantum theory, the minimal and maximal tensor products [64], composition rules between the minimal and maximal tensor products, and even for composition rules in tomographically nonlocal GPTs. So as we now prove, it follows that F is always a GPT subsystem of $F \otimes G$. One

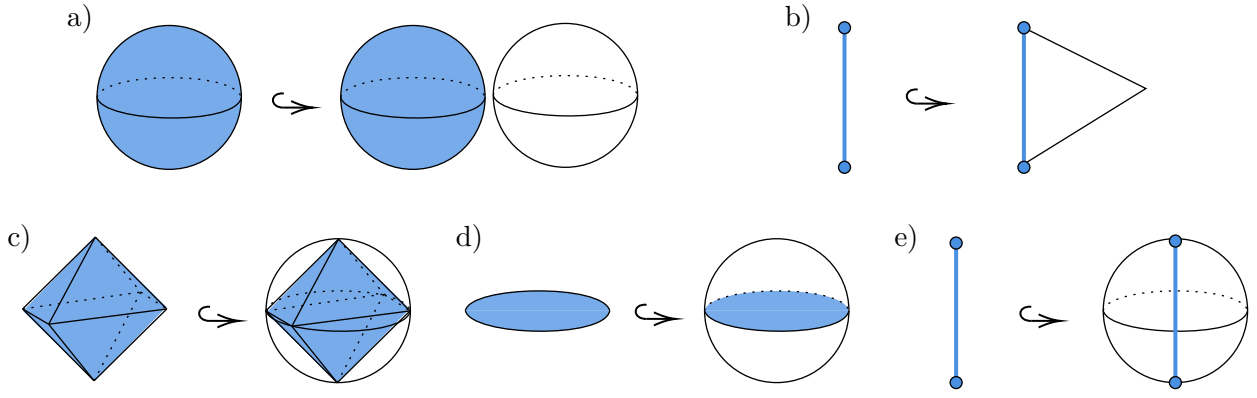


Figure 1: Five examples of GPT subsystems (in blue) of larger systems (in black). a) One qubit is a GPT subsystem of two qubits. b) A two-level classical system is a GPT subsystem of a three-level classical system. c) The convex hull of the stabilizer states and effects forms a GPT subsystem of a qubit. d) The rebit is a GPT subsystem of a qubit. e) A classical bit is a GPT subsystem of a qubit.

can define a GPT embedding of F into $F \otimes G$ given by the linear map $\iota : \Omega_F \rightarrow \Omega_{FG} :: s \mapsto s \otimes s'$ for any fixed normalized $s' \in \bar{\Omega}_G$, and the linear map $\kappa : \mathcal{E}_F \rightarrow \mathcal{E}_{FG} :: e \mapsto e \otimes u_G$. These maps preserve the probabilistic predictions, since they satisfy $e[s] = \kappa(e)[\iota(s)]$ as $\kappa(e)[\iota(s)] = e \otimes u[s \otimes s'] = e[s]u[s'] = e[s]$. They also act on the unit effect appropriately, since they satisfy $\kappa(u_F) = u_{FG}$, as $\kappa(u_F) = u_F \otimes u_G = u_{FG}$.

Example 2. Similarly, we now show that a direct sum factor F of a given GPT system $F \oplus G$ is indeed a GPT subsystem of $F \oplus G$. The properties of the \oplus operation that we need are merely that $(e_F \oplus e_G)[(s_F \oplus s_G)] = e_F(s_F) + e_G(s_G)$ and $u_F \oplus u_G = u_{F \oplus G}$ for effects. Hence, we can define a GPT embedding of F into $F \oplus G$ given by the linear map $\iota : \Omega_F \rightarrow \Omega_{F \oplus G} :: s \mapsto s \oplus 0$ and the linear map $\kappa : \mathcal{E}_F \rightarrow \mathcal{E}_{F \oplus G} :: e \mapsto e \oplus u_G$. This is a valid embedding, since these properties of the direct sum give $e[s] = \kappa(e)[\iota(s)]$ as $\kappa(e)[\iota(s)] = e \oplus u[s \oplus 0] = e[s] + u[0] = e[s]$ and $\kappa(u_F) = u_F \oplus u_G = u_{F \oplus G}$.

That a tensor factor of some larger system is a subsystem is completely standard. That components of a direct sum system should be considered subsystems is also standard. The direct sum structure $F \oplus G$ entails that the system is either described by a state in F or by a state in G . As F forms a space of states, each of which is a possible description of the system $F \oplus G$, it makes sense to consider F as a subsystem of $F \oplus G$. Consider for example a classical d -level system and a classical D -level system with $D > d$. A d -level system can be viewed as living inside the D -level system simply by ensuring that only

d of the possible states are ever accessed in one's experiment. Since classical systems can always be decomposed as direct sums, $\Delta_D = \Delta_d \oplus \Delta_{D-d}$, this is an instance of F (the d -level system) being a subsystem of $F \oplus G$ (the D -level system). Note that the direct sum structure also arises in superselected quantum systems [65], where the components are not necessarily classical.

Example 3. That our notion of a GPT subsystem also deems virtual quantum systems [66] to be subsystems is quite immediate, as these are isomorphic to any other quantum systems of the same dimension, and so trivially admit a linear embedding in our sense. So, our notion subsumes coarse-grained degrees of freedom encoded in some (potentially highly delocalized) larger set of fundamental systems. This includes quantum codes [67], decoherence-free subsystems [68], and time-delocalized subsystems [69], as these are all virtual subsystems [66]. Consider for example the simple quantum code that encodes a logical qubit into three physical qubits via the linear map taking $|0\rangle \rightarrow |0\rangle \otimes |0\rangle \otimes |0\rangle$ and $|1\rangle \rightarrow |1\rangle \otimes |1\rangle \otimes |1\rangle$. This defines a qubit on the Hilbert space defined by $\text{Span}[|0\rangle \otimes |0\rangle \otimes |0\rangle, |1\rangle \otimes |1\rangle \otimes |1\rangle]$. This qubit is a GPT subsystem of the three-qubit Hilbert space \mathbb{C}^8 . Indeed, the state map for the GPT subsystem embedding can simply be taken to be the encoding map itself, and the effect map can be taken to be the decoding map (acting contravariantly on effects).

Example 4. Another example of a GPT subsystem arises when considering the stabilizer

subtheory of quantum theory. Consider a stabilizer bit [60, 61], defined as the convex hull of the six eigenstates of Pauli operators on a qubit, together with the convex hull of the six effects that are projectors onto these (together with null and unit effects). According to our definition, the stabilizer qubit is a GPT subsystem of a standard quantum bit. Similarly, every stabilizer system of dimension d is a GPT subsystem of every quantum system of dimension d .

Example 5. An example of a GPT subsystem arises when considering real quantum theory. Consider a real bit, or rebit [62, 63], defined by the set of density matrices and POVM elements which can be written as linear combinations of the Pauli I , X , and Z operators (with no component of the Pauli Y operator). The rebit is a GPT subsystem of a qubit. To see this, consider the natural representation of a given state (or effect) O as a 3-dimensional real-valued vector $(1, \text{Tr}[OX], \text{Tr}[OZ])$; then, one can embed the states and effects of the rebit back into the full qubit simply by associating to each the Bloch vector $(1, \text{Tr}[OX], 0, \text{Tr}[OZ])$. This embedding clearly satisfies all the necessary properties. Similarly, any d -dimensional system in real quantum theory is a GPT subtheory of any d -dimensional system in quantum theory.

Example 6. Another example of a GPT subsystem is that of a decohered system, which can naturally be taken to live within the fundamental (pre-decohered) system, and consequently viewed as a subsystem thereof. The notion of decoherence (relative to some fixed basis) in quantum theory is well established, and corresponds to destroying all off-diagonal terms (relative to that basis) in the matrix representation of a state or effect. Within the context of general GPTs, one can generalize this idea to define hyperdecoherence processes [70–73], which correspond to idempotent maps taking GPT processes in the fundamental theory to GPT processes in the decohered theory. The fact that hyperdecoherence processes always leads to GPT subsystems of the fundamental theory follows from Ref. [47], which proved that there is always a linear embedding from a hyperdecohered system into the fundamental system that preserves the probabilities. We will expand on all of this in Section 6.6, where we also draw out the important consequences of this fact for the study

of noncontextuality.

Example 7. Finally, any fragment of a GPT G which satisfies all the mathematical properties of a GPT constitutes a GPT subsystem of G , since the one can always take the identity embedding of a relatively tomographic GPT fragment into itself as a valid GPT embedding.

3.1 Two GPT systems which are subsystems of each other are equivalent

Any good notion of subsystem should satisfy the condition that if two systems are subsystems of each other, then they are the same system (at least up to isomorphism). Here we show that our notion of subsystem does indeed satisfy this property.

For the proof of the main proposition in this Section, we will need the following Lemma:

Lemma 16. *Any embedding (ι, κ) of a GPT system (Ω, \mathcal{E}, p) into itself is a GPT isomorphism.*

The proof of this is in Appendix A

We can now prove the main result of this section.

Proposition 3. *Two GPT systems are equivalent if and only if they are GPT subsystems of each other.*

Proof. Suppose that the GPT systems $F := (\Omega_F, \mathcal{E}_F, p_F)$ and $G := (\Omega_G, \mathcal{E}_G, p_G)$ are equivalent. Then, by definition, we have a GPT isomorphism which embeds F into G and G into F . Therefore F is a GPT subsystem of G and G is a GPT subsystem of F .

Now, let us prove the converse. That F is a GPT subsystem of G means there is a (faithful) embedding $(\iota_1, \kappa_1) : F \rightarrow G$, and that G is a subsystem of F means that there is a (faithful) embedding $(\iota_2, \kappa_2) : G \rightarrow F$. We want to show that these facts give us an isomorphism between the two GPT systems.

We know that the ι_i and κ_i are injective, since GPT embeddings are necessarily faithful. What remains to be shown is that ι_i and κ_i are surjective state and effect maps, respectively.

To see this, consider the composite embedding of F into itself given by $(\iota_2 \circ \iota_1, \kappa_2 \circ \kappa_1) : F \rightarrow F$. By Lemma 16, this self-embedding is in fact an isomorphism. In particular, this tells us that $\iota_2 \circ \iota_1$ and $\kappa_2 \circ \kappa_1$ are surjective, which implies that ι_2 and κ_2 are surjective. Similarly, we can consider the composite embedding of G into itself to show that ι_1 and κ_1 are surjective.

It immediately follows that (ι_1, κ_1) and (ι_2, κ_2) are (possibly distinct) isomorphisms between F and G . Hence, F and G are equivalent GPT systems. \square

4 Shadows of GPT fragments

For any particular experiment on a given system, there is some true GPT describing the system, whether or not one knows what this GPT is. The laboratory operations one implements in the experiment will then be represented by a fragment of this true GPT. The states (effects) in this fragment will generally not be tomographically complete for the true GPT system.

If a given pair of distinct states (effects) in one's experiment cannot be discriminated by any of the effects (states) in the experiment, then those states (effects) will appear to be exactly identical within the context of the given experiment. In this case, the states (effects) will appear to exhibit operational identities that are not genuine—i.e., that do not hold in the true GPT. If one blindly trusts these apparent operational equivalences, then the GPT description one would give to the experiment will be incorrect—it will in fact be given by a particular informationally lossy function of the true GPT fragment describing the scenario. We will call this informationally lossy map the *shadow map*. The image of this map on the fragment in one's experiment captures how the GPT states and effects *appear* relative to the effects and states (respectively) one has implemented in one's experiment—although, as just described, this apparent characterization will necessarily be incorrect if the states and/or effects are not tomographic for each other.

If the states and effects are tomographic for each other, then this kind of mischaracterization cannot occur. We say that such states and effects are *relatively tomographic*, or simply that the fragment containing them is relatively tomographic.

Before formalizing these ideas fully, we present an illustrative example.

4.1 Example 1: a failure of relative tomographic completeness

To get some intuition for what happens when the states and effects in one's experiment are not relatively tomographic, consider a first simple

example.

Imagine that one's experimental apparatus can prepare all quantum states, but that one's measurement devices can only measure the Pauli Z observable. The set of states one can access in such an experiment is the full Bloch ball, while the set of effects contains only $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$, as well as the effects in the convex hull of these two together with the null effect and unit effect (as these can be obtained by post-processing). This fragment of states and effects are represented in the Bloch sphere in Figure 2.




Figure 2: If one restricts attention to only measurements of the Z component of a quantum state, then any two quantum states with the same Z component are informationally equivalent, regardless of their X and Y component. Consequently, the entire Bloch ball of states appears informationally equivalent to a single classical bit of states (whose extremal states correspond to eigenstates of the Pauli Z operator).

The statistics one can observe in any such experiment will necessarily be identical for any two states whose Z component are the same, regardless of their X and Y components, since no measurements in the experiment are sensitive to the latter. Consequently, this Z component of the state is the only feature of a given state that is relevant for the empirical data one can observe. One can then define the equivalence class of states which are identical relative to these measurements—one such class for every possible Z value. Identifying all states with the same value of Z results in a new ‘apparent’ state space which is simply a line segment containing the convex hull of the states $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$. Geometrically, these equivalence classes of states form disks at a fixed height, and one can choose a canonical representative of each equivalence class by orthogonally projecting the Bloch ball onto the line segment defined by the effects, as shown in Figure 2. This is an example of a shadow map, as defined formally in the next section.

In short, the shadow map applied to the *true* GPT states and effects in one's experiment

describes how those states and effects will appear, provided that one does not have access to any information aside from the data generated by combinations of those particular states and effects.

4.2 From GPT fragments to their shadows

In this subsection, we will define the notion of a shadow map and a GPT fragment shadow more formally. We begin by formalizing the notion of equivalence exemplified in the previous subsection and defining quotienting maps relative to it. This provides a first definition of GPT shadows. We then give a simpler definition that highlights the key abstract features of a shadow, and then show that the two definitions are essentially equivalent.

As motivated in the previous subsection, the idea of a shadow is to construct a proper GPT that reproduces the predictions of the fragment by taking the original fragment and identifying states (and effects) which cannot be distinguished by any of the effects (resp. states) in the theory. Formally, this means that we define an equivalence relation on the states (effects) and quotient with respect to this equivalence relation. That is, for a fragment $f = (\Omega_f, \mathcal{E}_f, p_f)$ we define an equivalence relation on states as

$$s_1 \sim s_2 \iff p_f(s_1, e) = p_f(s_2, e) \quad \forall e \in \mathcal{E}_f. \quad (24)$$

Since the right domain of p_f is $\text{Span}[\mathcal{E}_f]$, this condition is equivalent to the condition that $p_f(s_1 - s_2, \cdot) = 0$, or in other words, the condition that $s_1 - s_2 \in \text{Ker}_L[p_f]$, where this denotes the left kernel of p_f . This mathematical reframing is useful, because it is now clear that what we are doing is in fact a linear quotient with respect to the subspace $\text{Ker}_L[p_f]$. We denote the linear map defined by quotienting in this manner by $\sim: \Omega_f \rightarrow \Omega_f/\sim$. (A more explicit notation would be $\sim_{\text{Ker}_L[p_f]}$, but as we only have a single notion of quotienting for states in this work, we leave the subscript implicit. Also, the use of the upside down symbol \sim is to distinguish from the usual notion of quotienting \sim that takes unquotiented operational theories to GPTs [5, 74].)

Similarly, we define an equivalence relation (and an equivalent linear algebraic reframing of it) for the effects in the fragment as

$$e_1 \sim e_2 \iff p_f(s, e_1) = p_f(s, e_2) \quad \forall s \in \Omega_f \quad (25)$$

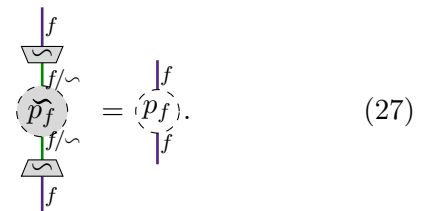
$$\iff e_1 - e_2 \in \text{Ker}_R[p_f], \quad (26)$$

where we have introduced the obvious notation for the right Kernel of p_f . Identifying effects which

cannot be distinguished by states in the fragment is done by quotienting relative to this equivalence relation; this quotienting defines a linear map that we also denote $\sim: \mathcal{E}_f \rightarrow \mathcal{E}_f/\sim$. (A more explicit notation one could also use would be $\sim_{\text{Ker}_R[p_f]}$.)

It is critical to note that this quotienting is distinct from the one that takes unquotiented operational theories to quotiented operational theories [5, 74] and that is also critically important for the study of noncontextuality [1, 6, 37]. Here, the equivalence relation one quotients relative to is defined by the specific experimental scenario, *not* by the set of all possible preparation and measurement procedures one could in principle have done on the system in question. Moreover, the quotienting we consider in this work maps one between two quotiented operational theories: from the true quotiented operational theory (the true GPT) to an *apparent* quotiented operational theory (the apparent GPT).

The two quotients defined above map the states (effects) of the fragment to the states (effects) in its shadow, but it remains to specify the probability rule for the shadow. We define the probability rule simply by demanding that it assigns the same probability for the equivalence classes as the original rule assigned for the elements of those classes. That is, if we explicitly denote the equivalence classes of states and effects as \tilde{s} and \tilde{e} , respectively, then the probability rule \tilde{p}_f for the shadow is defined by $\tilde{p}_f(\tilde{e}, \tilde{s}) := p_f(e, s)$ for all $\tilde{s} \in \Omega_f/\sim$ and all $\tilde{e} \in \mathcal{E}_f/\sim$. That $\tilde{p}_f(\tilde{e}, \tilde{s})$ depends only on the equivalence classes follows from the fact that $p_f(e, s)$ is the same for all $e \in \tilde{e}$ and is the same for all $s \in \tilde{s}$; that it actually constitutes a valid map—one which is linear and unique—follows from the universal property of quotients [75, Sec. 13]. That is, one can define \tilde{p}_f as the unique map satisfying





$$\begin{array}{c} f \\ \downarrow \\ \text{---} \circ \tilde{p}_f \text{---} \\ \downarrow \\ f \end{array} = \begin{array}{c} f \\ \downarrow \\ \text{---} \circ \tilde{p}_f \text{---} \\ \downarrow \\ f \end{array} \quad (27)$$

Definition 17 (GPT fragment shadow). *Given a GPT fragment $(\Omega_f, \mathcal{E}_f, p_f)$, the GPT shadow f/\sim*

is the GPT system given by

$$f/\sim := \left(\left\{ \begin{array}{c} f/\sim \\ \text{trapezoid} \\ f \\ \text{trapezoid} \\ s \end{array} \right\}_{s \in \Omega_f}, \left\{ \begin{array}{c} \text{trapezoid} \\ e \\ f \\ \text{trapezoid} \\ f/\sim \end{array} \right\}_{e \in \mathcal{E}_f}, \left(\begin{array}{c} f/\sim \\ \text{trapezoid} \\ \tilde{p} \\ \text{trapezoid} \\ f/\sim \end{array} \right), \right) \quad (28)$$

where  and  are defined as quotienting relative to the equivalence relations we defined for states (Eq. (24)) and for effects (Eq. (25)), respectively, and the probability rule \widetilde{p}_f is defined as in Eq. (27).⁹

Note that a GPT shadow of any fragment is always a valid GPT system (even if the fragment is not a valid GPT system—i.e., is not relatively tomographic).¹⁰

Thus, quotienting defines a map from an arbitrary fragment to its shadow. This *shadow map* is a particular GPT fragment embedding of f into f/\sim . This follows from the fact that the shadow reproduces the empirical probabilities of the original fragment, since

$$\begin{array}{c} \triangle_e \\ |f| \\ \sim \\ \text{---} \\ |f| \\ \tilde{p}_f \\ \text{---} \\ |f| \\ \sim \\ \text{---} \\ |f| \\ \triangle_s \end{array} = \left(p_f \right), \quad \forall s \in \Omega_f, e \in \mathcal{E}_f, \quad (29)$$

together with the facts that the quotienting map on states is a state map (from Ω_f to $\Omega_{f/\sim}$) and that the quotienting map on effects is an effect map (from \mathcal{E}_f to $\mathcal{E}_{f/\sim}$).

At its core, the GPT shadow of a fragment is simply *a proper GPT system that reproduces the probabilities of the fragment*.¹¹ That is, it is a *relatively tomographic* fragment that satisfies Eq. (29).

⁹Note that the shadow can also be defined (up to unique isomorphism) by the universal property of quotients in the category **GPT-Fragment**.

¹⁰We can therefore think of the shadow as a map from GPT fragments to GPT systems, that is, some $S : |\text{GPT-Fragment}| \rightarrow |\text{GPT-System}|$.

¹¹If one defines GPT fragments in a slightly more general way that allows for subnormalized states whose normalized counterparts are not in the fragment [17], then the shadow of such a fragment is not strictly a GPT system, but rather is a relatively tomographic fragment.

That the shadow is a relatively tomographic fragment is clear from the way that quotienting is defined. Explicitly: for any $\mathfrak{F} \in \Omega_f / \sim$ one can define a linear functional $\tilde{p}_{\mathfrak{F}}(\mathfrak{F}; \cdot)$ taking \mathcal{E}^A / \sim to \mathbb{R} . By the definition of the equivalence relation, it follows that the set of all such linear functionals are separating for \mathcal{E}_f / \sim . A similar argument holds for effects.

These two facts are indeed the *only* essential features of a GPT shadow.

Consequently, we can give an alternative but equivalent definition of GPT shadows as follows.

Definition 17' (GPT fragment shadow). *Given a GPT fragment $(\Omega_f, \mathcal{E}_f, p_f)$, a GPT shadow $\mathcal{S}(f) := (\sigma(\Omega_f), \tau(\mathcal{E}_f), p_S)$ of f is a GPT system defined by two maps $\sigma : \text{Span}(\Omega_f) \rightarrow U$ and $\tau : \text{Span}(\mathcal{E}_f) \rightarrow V$ and a tomographic probability rule $p_S : \text{Span}[\sigma(\Omega_f)] \times \text{Span}[\tau(\mathcal{E}_f)] \rightarrow \mathbb{R}$ such that*

$$\begin{array}{c} \textcolor{violet}{|f\rangle} \\ \textcolor{violet}{(pf)} \\ \textcolor{violet}{|f\rangle} \end{array} = \begin{array}{c} \textcolor{violet}{|f\rangle} \\ \boxed{\tau} \\ \textcolor{green}{|S(f)\rangle} \\ \textcolor{green}{|S(f)\rangle} \\ \boxed{\sigma} \\ \textcolor{violet}{|f\rangle} \end{array} . \quad (30)$$

Any such map from f to $\mathcal{S}(f)$ is a GPT fragment embedding that we term a shadow map.

In the strictest sense, this is more general than Definition 17; for example, the states and effects of the shadow in this definition are not assumed to span the space they live in, while the states and effects in a quotiented fragment as in Definition 17 will always do so. However, as we have already established that GPT state (effect) spaces are independent (up to isomorphism) of the vector spaces they are taken to live in, this does not constitute a meaningful distinction between the two definitions. Indeed, the two definitions are equivalent up to this representational distinction, as is demonstrated by the following theorem.

Theorem 17. *Any two shadows of a given GPT fragment are equivalent to each other (no matter how they are constructed). Moreover, any GPT system equivalent to a shadow is itself a shadow.*

The full proof is given in Appendix B, but we summarize the proof here.

Proof sketch. The only variability in the shadow is the choice of the σ and τ , which are necessarily linear maps (Lemma 26 in Appendix B); once these

are chosen, there is a unique probability rule which reproduces the correct probabilities. Consequently, we sometimes refer to the shadow map as a tuple (σ, τ) . Moreover, it is clear that the properties that σ and τ must satisfy in order that there does exist a probability rule which reproduces the probabilities is that the kernels of σ and τ need to be contained in the left and right kernels of p_f respectively. Moreover, as we demand that the probability rule be tomographic then this means that the kernels of σ and τ are actually equal to the left and right kernels of p_f respectively. The universal property of quotients then implies that σ and τ factor through the quotienting maps as $\sigma = \tilde{\sigma} \circ \smile$ and $\tau = \tilde{\tau} \circ \smile$ for some injective $\tilde{\sigma}$ and $\tilde{\tau}$. It is then easy to see that $(\tilde{\sigma}, \tilde{\tau})$ defines a GPT isomorphism between f/\smile and $\mathcal{S}(f)$, and we can then compose these isomorphisms to give an isomorphism between any two shadows $\mathcal{S}(f)$ and $\mathcal{S}'(f)$. \square

In other words, all methods for constructing a shadow lead to the same end result¹². Thus, we sometimes refer to ‘the’ shadow of a fragment, since all shadows are equivalent.

Yet another equivalent characterization of GPT shadows is the following: a GPT shadow of a fragment f is any GPT system G such that there is a full GPT embedding of f into G . This is very useful, as it means that one can verify that a given embedding defines a shadow map simply by checking that 1) its codomain is a valid GPT system, and 2) it is a full embedding.

In Appendix C, we also prove the following:

Lemma 18. *The shadow map for a given fragment is faithful if and only if the fragment is relatively tomographic.*

For example, the shadow map in the example of Subsection 4.1 is a (nontrivial) projection on the nontomographic fragment in question, and so fails to be faithful.

The idea of GPT fragment shadows has appeared (at least indirectly) a number of times in the literature [47, 76, 77]. Most notably, the notion of an “alternative reduced space” representation introduced in Ref. [76] is an instance of Definition 17’ when one takes the probability rule to be an inner product (and specializes to quantum

theory). Indeed, Ref. [76] presents two different ways of computing such a shadow via projections. These two methods are illustrated in Example 2 in Subsection 4.4 below.

4.3 The shadow map acts trivially on GPT subsystems

We now note a simple but critical fact: that shadow maps act essentially trivially on GPT fragments that are themselves GPT systems—that is, on GPT subsystems (fragments that are relatively tomographic).

Theorem 19. *Every GPT subsystem is equivalent to its own shadow.*

Proof. This follows immediately from the definition of shadows in terms of quotienting (Definition 17). As the states and effects in a GPT subsystem are tomographic for each other, it follows that the equivalence relation is trivial, and the maps \smile (for both states and effects) are the identity. \square

In other words, a GPT fragment whose states and effects are relatively tomographic is isomorphic to its shadow, with the isomorphism given by the shadow map itself.

We illustrate this by an example in the next section.

4.4 Example 2: shadows of a GPT subsystem

We can get some more intuition for the shadow map by considering a simple example within quantum theory. In this example, the fragment in question is relatively tomographic, and hence the shadow map acts trivially in the sense that its image and preimage are equivalent as GPT systems (by Theorem 19). This example was suggested by Markus Müller.

Consider a fragment of a 2-dimensional quantum system comprised of an equatorial disk of states and an equatorial disk of effects, but where these disks are neither identical nor orthogonal, as shown in Figure 3.

That is, the set of states in the fragment corresponds to the set of density matrices associated with points in the disk labeled by Ω , and the set of effects in the fragment corresponds to the set of POVM elements associated with the points in the convex hull of the disk labeled by

¹²The results of Ref. [76] prove that all methods for computing shadow maps are equivalent for assessing simplex-embeddability; our result is stronger as it refers to full equivalence as GPTs.




Figure 3: Consider the fragment of quantum theory whose state space is the disk of states shown here as Ω and whose effect space is the disk of effects shown here as \mathcal{E} (strictly speaking, the effect space is the convex hull of this disc with the zero and unit effects). Two natural GPT shadows are constructed by projecting one of these onto the other, as shown in (a) and (b) respectively. These two shadows are equivalent as GPT systems, in accordance with Theorem 17. They are also equivalent to the original GPT fragment, in accordance with Theorem 19.

\mathcal{E} , the zero effect and the unit effect. Denoting a generic density operator in Ω by s and a generic POVM element in \mathcal{E} by E , we assign probabilities to state-effect pairs following the Born rule, as $p(E, s) = \text{Tr}[Es]$.

We describe two different methods of constructing shadows of this fragment. These are depicted in Fig. 3(a) and (b), respectively, and are essentially the same as the two methods for computing a ‘reduced space representation’ introduced in Ref. [76].

The first shadow map we consider is simply the orthogonal projection of the effect space onto the state space. This is depicted in Fig. 3(a), and results in a new effect space which constitutes a smaller disk lying in the same plane as the state space (which is unchanged). Note that the probabilities assigned by the Born rule to any given state-effect pair in the original fragment are unchanged by this projection, since the projection only changes the y-component of the effects and since the y-component of every state in the fragment is already zero.

This is an instance of the shadow map in Definition 17', where the map σ is taken to be trivial and τ is taken to be the projection map just described; since these maps together preserve the probabilities while mapping the states and effects to a valid GPT system, they define a valid shadow map.

The second method is the orthogonal projection of the state space onto the effect space. This is depicted in Fig. 3(b), and results in a new state

space which constitutes a smaller disk lying in the same plane as the effect space (which is unchanged). (Here, it is τ that is trivial and σ that is defined by a projection map.)

The two resulting GPT systems are equivalent GPT systems, in accordance with Theorem 17. (This is evident visually by the fact that they are identical subsets of the Bloch sphere up to a rotation.) Moreover, they are also equivalent as GPT systems to the original fragment itself, as a consequence of the fact that the latter constitutes a GPT system together with Theorem 19. This is the sense in which the shadow map is said to act trivially (despite the fact that it does change the precise representation of the fragment in question).

This fragment of the qubit has states and effects which are tomographic for one another, but which do not span the same vector space. Not all papers in the literature would consider such an example to be a valid GPT, since many papers demand that the states and effects of a GPT span the same vector space. But such more general cases are no more or less expressive (in terms of the probabilities they generate as well as the operational identities they encode) as those where the states and effects span the same space, so we consider them to be valid GPTs as well. Indeed, as just shown, the original fragment (whose states and effects are tomographic for each other but not spanning the same space) is equivalent to the two shadows of it we constructed (whose states and effects are both tomographic *and* span the same space). This again demonstrates how our definitions are more flexible than standard definitions.

4.5 Summary of concepts

We finish this section by presenting a pictorial depiction of the main concepts introduced from Sections 2 to 4, and how they relate to each other (see Fig. 4). In the remaining of the paper we discuss their impact on assessments of nonclassicality.

5 The effect of shadow maps on simplex embeddability

In this section, we consider the effect of shadow maps on simplex-embeddability of a given fragment.

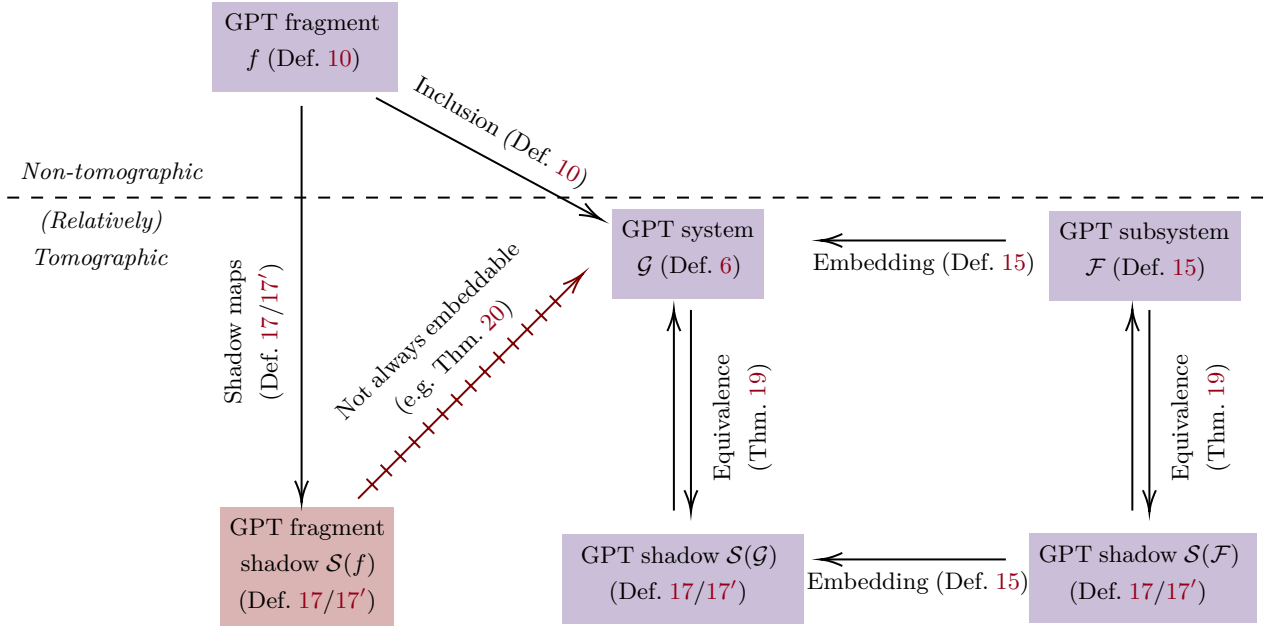


Figure 4: Summary of concepts introduced from Sections 2 to 4, and how they relate to each other. **Concepts:** GPT system \mathcal{G} , GPT subsystem \mathcal{F} , fragment f of a GPT system \mathcal{G} , shadow $\mathcal{S}(f)$ of a GPT fragment f , shadow $\mathcal{S}(\mathcal{G})$ of a GPT system \mathcal{G} , and shadow $\mathcal{S}(\mathcal{F})$ of a GPT subsystem \mathcal{F} . **Relations:** \mathcal{G} and $\mathcal{S}(\mathcal{G})$ are equivalent, just as \mathcal{F} and $\mathcal{S}(\mathcal{F})$ are also equivalent. Subsystem \mathcal{F} of \mathcal{G} can be embedded into \mathcal{G} , just as $\mathcal{S}(\mathcal{F})$ can be embedded into $\mathcal{S}(\mathcal{G})$. The GPT fragment f can be included into \mathcal{G} via an inclusion map, and mapped into the shadow $\mathcal{S}(f)$ via the shadow maps. The shadow $\mathcal{S}(f)$ cannot always be embedded into \mathcal{G} . As we demonstrate in the coming sections, only GPT shadows that do not constitute subsystems (red box) impose challenges to assessments of nonclassicality, while all other objects (purple boxes) yield trustworthy assessments.

5.1 Shadow maps can break simplex embeddability

We now prove that the shadow of a GPT fragment may not be simplex-embeddable, even if the fragment itself is simplex-embeddable. This follows from a general fact: that *every possible* polytopic GPT system can be recovered as the shadow of some fragment of a simplicial GPT system.

Theorem 20. *Every polytopic GPT system is the shadow of some fragment of a strictly classical (i.e., simplicial) GPT system of some dimension.*

We prove this in Appendix D. This theorem dates back to Holevo [77] and Beltrametti-Bugajski [78], and an analogous construction works for arbitrarily close approximations of nonpolytopic GPTs [47]. Our proof aims to give as clear and explicit a construction of the fragment as possible.

The idea of the construction is to simply imagine that all extremal preparations in one’s experiment are perfectly discriminable—the vertices of a simplicial theory, and that one merely has not implemented sufficient measurements to see this

fact. Under this assumption, it is easy to construct a simplicial fragment that reproduces one’s data table. All that remains is to figure out which subset of the logically possible measurements are the ones that reproduce the observed data. This is done simply by associating to each effect e the functional on vertices of the simplex defined by $\xi_e(\cdot) := p(e, \cdot)$.

We illustrate this theorem by a simple example: by demonstrating how the gbit is the shadow of a fragment of a simplicial theory. This example has also been discussed previously in the literature [47, 57, 77]. (Here, we give some extra details relative to those earlier presentations.)

5.2 Example 3: the gbit as the shadow of a fragment of a simplicial theory

We now repeat a well-known example that was originally due to Holevo [77, Sec. 1.5, p. 18].

The gbit \mathcal{G} has four pure states which form a regular square in a plane of \mathbb{R}^3 that does not contain the zero vector. We can write these as

$$s_1 = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, s_2 = \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}, s_3 = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, s_4 = \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix}. \quad (31)$$

The normalized state space $\bar{\Omega}_{\mathcal{G}}$ of the gbit is the convex hull of those extremal states. The effects in the gbit include all of the logically possible effects (so that the GPT is said to satisfy the no-restriction hypothesis [74]). These effects constitute the convex hull $\mathcal{E}_{\mathcal{G}} = \text{Conv}(\{0, u, e_{13}, e_{14}, e_{23}, e_{24}\})$, where e_{ij} are the effects that assign probability one to all states on the edge of the square that contains states s_i and s_j . By defining $e(s) = e \circ s$, where \circ is the usual matrix product¹³, these vectors take the explicit form

$$e_{13}^T = \frac{1}{2} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, e_{14}^T = \frac{1}{2} \begin{bmatrix} 1 \\ 1 \\ -1 \end{bmatrix}, e_{23}^T = \frac{1}{2} \begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix}, e_{24}^T = \frac{1}{2} \begin{bmatrix} 1 \\ -1 \\ -1 \end{bmatrix}. \quad (32)$$

The states and effects in Eq. (31) and Eq. (32) are plotted in purple in Figure 5. Defining the unit effect as

$$u^T = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad (33)$$

one can check that $u(\omega) = 1$ for all $\omega \in \bar{\Omega}_{\mathcal{G}}$ and that $e_{13} + e_{24} = e_{14} + e_{23} = u$. Thus, the gbit is taken to have two binary outcome measurements, defined by $\{e_{13}, e_{24}\}$ and $\{e_{14}, e_{23}\}$.

Now we will construct a fragment of a simplicial GPT system whose shadow reproduces the gbit, which we hereby denote as the *Holevo fragment* \mathcal{H} . This construction is an example of the general construction in Appendix D that is used to prove Theorem 20.

There are $n = 4$ extremal states of the gbit, so the classical system in question has a 4-vertex simplex as its normalized state space—i.e., a tetrahedron Δ_3 in three dimensions. We can take each of the extremal states in \mathbb{R}^4 to be unit basis vectors

$$\mu_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \mu_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \mu_3 = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \mu_4 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}, \quad (34)$$

and we take the set of normalized states in the fragment to be $\bar{\Omega}_{\mathcal{H}} := \text{Conv}(\{\mu_1, \mu_2, \mu_3, \mu_4\})$.

To construct the effects in the Holevo fragment, we first define the matrix L in which each column

¹³Note that the effects are naturally described as row vectors, as they are elements of the dual. Here we write their transposes (as column vectors) to spare space.




Figure 5: Here we depict the Holevo fragment: states and effects from a simplicial theory whose shadow map (depicted as downward arrows) gives the gbit state and effect space, $\sigma(\Omega_{\mathcal{H}}) = \Omega_{\mathcal{G}}$ and $\tau(\mathcal{E}_{\mathcal{H}}) = \mathcal{E}_{\mathcal{G}}$, respectively. Note that we have only shown the (convex hull of the) four extremal effects that are not the unit and null effect, since depicting the full effect space would require four dimensions.

is a vector s_i of the gbit, namely

$$L = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & -1 \end{bmatrix}. \quad (35)$$

We then act this matrix on the right of each gbit effect $e \in \mathcal{E}_{\mathcal{G}}$ to get the effects of the Holevo fragment, defining $\mathcal{E}_{\mathcal{H}} := \text{Conv}(\{\xi_0, \xi_{13}, \xi_{14}, \xi_{23}, \xi_{24}, \xi_u\})$ where $\xi_e = e \circ L$. Explicitly, these are the effects $\xi_0 = \vec{0}$, $\xi_u = \begin{bmatrix} 1 & 1 & 1 & 1 \end{bmatrix}$, and

$$\xi_{e_{13}}^T = \begin{bmatrix} 1 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \xi_{e_{14}}^T = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 1 \end{bmatrix}, \xi_{e_{23}}^T = \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix}, \xi_{e_{24}}^T = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 1 \end{bmatrix}. \quad (36)$$

In summary, the Holevo fragment \mathcal{H} is defined by $\Omega_{\mathcal{H}} := \text{Conv}(0 \cup \bar{\Omega}_{\mathcal{H}})$, $\mathcal{E}_{\mathcal{H}}$, and the probability rule given by the evaluation map (i.e. by simply composing the effect with the state).

The states in this fragment span \mathbb{R}^4 . However, the effects do not span the whole dual space $(\mathbb{R}^4)^*$, since the linear dependence $\xi_{13} + \xi_{24} = \xi_{14} + \xi_{23}$ between them implies that only three of them are linearly independent. Consequently, $\dim[\text{Span}(\mathcal{E}_{\mathcal{H}})] < \dim[\text{Span}(\Omega_{\mathcal{H}})]$, and so this Holevo fragment is not relatively tomographic. More specifically, the effects cannot separate the states.

Finally, let us show that the shadow $\mathcal{S}(\mathcal{H})$ of this Holevo fragment is the gbit \mathcal{G} . This fragment and its shadow are depicted in Figure 5.

As we prove in Appendix D, the shadow of a fragment of a simplicial system constructed in this manner is given by taking the shadow maps to be $\tau = L$ and $\sigma = L_r^{-1}$, where L_r^{-1} is a right inverse of L . That this defines a shadow map will be justified below (by verifying that it leads to valid GPT system and that $\tau = L$ and $\sigma = L_r^{-1}$ are full).

The matrix representation of L_r^{-1} is

$$L_r^{-1} = \frac{1}{4} \begin{bmatrix} 1 & 2 & 0 \\ 1 & -2 & 0 \\ 1 & 0 & 2 \\ 1 & 0 & -2 \end{bmatrix}, \quad (37)$$

since $L \circ L_r^{-1} = \mathbb{1}_3$, as one can check. Now, we get the shadow states by applying L to the states of $\Omega_{\mathcal{H}}$, i.e., $\sigma(\mu_i) = L \circ \mu_i \in \Omega_{\mathcal{S}(\mathcal{H})}$, $i = 1, \dots, 4$ (recall that, given our choice of basis, each μ_i is a member of the standard basis of \mathbb{R}^4). Finally, we apply L_r^{-1} (via pre-composition) to the effects in the Holevo fragment to get the effects $\tau(\xi_j) \in \mathcal{E}_{\mathcal{S}(\mathcal{H})}$, $j \in \{0, e_{13}, e_{14}, e_{23}, e_{24}, u\}$. Working this multiplication out for the pure objects gives

$$\sigma(\mu_1) = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \sigma(\mu_2) = \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}, \quad (38)$$

$$\sigma(\mu_3) = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \sigma(\mu_4) = \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix}, \quad (39)$$

for states, and gives

$$\begin{aligned} [\tau(\xi_{e_{13}})]^T &= \frac{1}{2} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, [\tau(\xi_{e_{14}})]^T = \frac{1}{2} \begin{bmatrix} 1 \\ 1 \\ -1 \end{bmatrix}, \\ [\tau(\xi_{e_{23}})]^T &= \frac{1}{2} \begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix}, [\tau(\xi_{e_{24}})]^T = \frac{1}{2} \begin{bmatrix} 1 \\ -1 \\ -1 \end{bmatrix}, \end{aligned} \quad (40)$$

for effects, as well as $[\tau(\xi_0)]^T = \vec{0}$ and $[\tau(u)]^T = u$. Comparing these with Eq. (31) and Eq. (32), respectively, one sees that indeed $\Omega_{\mathcal{S}(\mathcal{H})} = \Omega_{\mathcal{G}}$ and $\mathcal{E}_{\mathcal{S}(\mathcal{H})} = \mathcal{E}_{\mathcal{G}}$, and the same probability rule is being used (namely the evaluation map, given by standard matrix multiplication), so the shadow $\mathcal{S}(\mathcal{H})$ is indeed \mathcal{G} .

One can now easily verify that taking $\tau = L$ and $\sigma = L_r^{-1}$ does indeed define a shadow map. It suffices to check that the image of these maps is a valid GPT system (which it is, as we just showed the image is \mathcal{G}), and that $\tau = L$ and $\sigma = L_r^{-1}$

define a full embedding. For σ to be full means that $\sigma(\Omega_{\mathcal{H}}) = \Omega_{\mathcal{G}}$, which is the case, as Eq. (38) shows that the image of the extremal states of $\Omega_{\mathcal{H}}$ coincides with the extremal points of $\Omega_{\mathcal{G}}$. Similarly, one can check that $\tau = L$ is full via Eq. (40).

Because the Holevo fragment is not relatively tomographic, its shadow is not faithful and so introduces new operational identities. We can see this explicitly from the fact that a generic state in the tetrahedron Δ_3 gets mapped under the shadow map as

$$\begin{bmatrix} p_1 \\ p_2 \\ p_3 \\ p_4 \end{bmatrix} \rightarrow \begin{bmatrix} 1 \\ p_1 - p_2 \\ p_3 - p_4 \end{bmatrix}. \quad (41)$$

So, two different states p and q in Δ_3 that happen to have components obeying $p_1 - p_2 = q_1 - q_2$ and $p_3 - p_4 = q_3 - q_4$ will get mapped to the same state in $\Omega_{\mathcal{S}(\mathcal{H})}$. This identification of different states occurs because the effects in $\mathcal{E}_{\mathcal{G}}$ that are able to distinguish these distributions are absent from $\mathcal{E}_{\mathcal{H}}$. Thus, experiments that fail to have access to these distinguishing effects lead one to the mistaken identification of fundamentally distinct states, which in turn leads one to believe that the GPT system in question is a gbit, even though it is by assumption a four-level classical system. This is an extreme example, in that gbts are known to be maximally contextual and compose to reach maximal violations of Bell-CHSH inequalities.

5.3 Shadow maps cannot introduce simplex-embeddability

One might also like to see an example of a fragment that is not simplex-embeddable, but whose shadow is simplex-embeddable. However, this is impossible.

Proposition 4. *Consider any fragment f of a GPT system G (even one whose states and effects are not tomographic for each other). If the shadow of f is simplex-embeddable, then f is simplex-embeddable.*

The proof follows immediately from transitivity of GPT embeddings: since every fragment embeds into its shadow, and since the shadow by assumption embeds into a simplicial GPT system, the fragment itself embeds into a simplicial theory.

We give an example of a shadow map that dramatically distorts the fragment

in question while nonetheless preserving simplex-embeddability in the next section.

Theorem 20 and Proposition 4 (respectively) can be summarized as follows: *Shadow maps can break simplex-embeddability, but cannot introduce it.*

5.4 Example 1 redux: a nontrivial shadow map that preserves simplex embeddability

Finally, reconsider the example from Section 4.1, where the shadow map acts in a highly nontrivial manner, so that (like in the Holevo example) the shadow is very different from the original GPT fragment. But despite this (and in contrast to the Holevo example), simplex-embeddability is nonetheless preserved by the shadow map in this example.

Consider again the fragment of a qubit containing all possible states in the Bloch ball and containing only those effects in the convex hull of the projectors $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$ (besides the zero and unit effects). This fragment is depicted in Figure 6 (which is simply a repeat of Figure 2).




Figure 6: Consider the fragment of quantum theory whose state space contains all states in the Bloch ball and whose effect space is the set of all effects in the line segment in red, besides the zero and unit effects. A GPT shadow of this fragment can be constructed by orthogonally projecting the set of states onto the line segment. This results in a shadow which is simply the GPT representing a classical bit.

One can compute a shadow of this fragment by orthogonally projecting the state space onto the line segment defined by the effects, as shown in the figure. Like in the projections in the previous section, this linearly maps the states and effects into a proper GPT system (here, the classical bit), while preserving the probabilities assigned (by the Born rule) to any state-effect pair. Consequently, it constitutes a valid shadow map.

The shadow constitutes all and only the states and effects of a classical bit. Therefore it is simplex-embeddable (and indeed is itself simplicial).

As one can check explicitly (e.g., via a linear program [57]), the original GPT fragment in this example is simplex-embeddable. So this constitutes an example where the shadow map takes a simplex-embeddable fragment to a simplex embeddable GPT system, despite dramatically altering the fragment.

6 Implications for noncontextuality

At this stage, readers familiar with the notion of simplex-embeddability may already see how the above results are relevant to the study of nonclassicality. However, we will emphasize and make more explicit these connections in this section.

In the literature, it is often claimed that assessments of noncontextuality rest on the assumption of tomographic completeness. However, our work implies that they rest on a weaker assumption, namely that the states and effects in the experimentally realized fragment are *relatively tomographic*, regardless of whether they are tomographic for the full set of effects and states of the true GPT system. We term this the *relative tomographic completeness* assumption.

Recall from Section 4.1 that when the states and effects in one's experiment are not relatively tomographic, the statistics one observes in any experiment with them will be insufficient to correctly characterize them. More specifically, some states and/or effects that are distinct in the true GPT governing the experiment will appear to be equivalent for all the data obtained in that experiment. Thus, in any experiment that makes use only of this data (as opposed to, for example, appealing to some theory-dependent arguments concerning the inner workings of one's laboratory apparatuses), any attempt to characterize the GPT state and effect vectors describing one's preparations and measurements in the experiment will in fact characterize a shadow of the fragment, rather than the fragment itself. We expand on this idea in the next subsection.

If the fragment in one's experiment is relatively tomographic, then this cannot lead to any problems, since characterizing such a fragment's shadow is equivalent to characterizing the true fragment. However, when relative tomography does fail, the mischaracterizations of one's GPT states and effects (and equivalently, of the

operational identities that they obey) described by the shadow map can lead to incorrect assessments of nonclassicality (simplex-embeddability). However, it does not *necessarily* imply that one will reach the wrong conclusion about nonclassicality; the specific details of the mischaracterization—the details of how the shadow map distorts the true fragment—will determine whether or not this happens. Later subsections will expand on the different possibilities, which are summarized in Table 1.

	Shadow is embeddable	Shadow is non-embeddable
Fragment is embeddable	- Any tomographic fragment of a simplicial theory; - Our example in Fig. 6.	- Holevo's example (Sec. 5.2).
Fragment is non-embeddable	IMPOSSIBLE (Proposition 4)	- A qubit; - Any tomographic non-embeddable fragment; - Our example in Fig 3.

Table 1: The different possible cases for how a shadow map affects the simplex-embeddability properties of a given GPT fragment, and examples of each. Mistaken assessments of nonclassicality occur *only* in the top right case (the cell with a black background).

6.1 Theory-agnostic tomography

We can exemplify and formalize the fact that theory-independent experiments characterize the shadow of one's true GPT fragment (rather than the true fragment itself) by considering theory-agnostic tomography [45, 79]. This is the state-of-the-art method for determining the characterization of GPT states and effects in a prepare-measure experiment.

We formalize theory-agnostic tomography using a novel diagrammatic representation. Note that in this section (and the associated Appendix E), we consider linear maps whose inputs (outputs) are not associated with a GPT system, and so are labeled by the domain (codomain) of the map (whereas in the rest of the paper, systems were labeled by the associated GPT system).

Imagine one has a fragment $f = (\Omega_f, \mathcal{E}_f, p_f)$ of some GPT system G representing some states and effects which are performed in an experiment. As discussed in Section 2.3, no real experiment will actually implement every state and effect in the fragment, as this set is infinite; however, an experiment involving the convexly extremal states

and effects of a fragment is sufficient to uniquely infer the existence of—and probabilities generated by—any states and effects in the convex hull of these. (Moreover, in a real, lossy experiment one will not have exactly normalized states, so a noise-robust analysis must be done to contend with this.)

Given a fragment f , we can construct a data table D_f in which we index columns by convexly extremal states and rows by convexly extremal effects, so that the entries in the data table are the probability of obtaining the given effect on the given state. We denote the entries of the data table

$$[D_f]_{i,j} := p_f(s_i, e_j), \quad (42)$$

where the s_i are the n convexly-extremal GPT states in the fragment, the e_j are the m convexly-extremal GPT effects in the fragment, and $[D_f]_{i,j}$ is a matrix element of $D_f \in M_{m \times n}(\mathbb{R})$, where $M_{m \times n}(\mathbb{R})$ is the set of real $m \times n$ matrices. We refer to this as the data table generated by the fragment.¹⁴

Definition 21 (Theory-agnostic tomography). *Given a data table $D \in M_{m \times n}(\mathbb{R})$, consider the following procedure:*

1. Find the smallest k such that¹⁵

$$\boxed{D}_{\mathbb{R}^n}^{\mathbb{R}^m} = \begin{array}{c} \mathbb{R}^m \\ \text{---} E \text{---} \\ \text{---} \mathbb{R}^k \\ \text{---} S \text{---} \\ \text{---} \mathbb{R}^n \end{array}, \quad (43)$$

where $E \in M_{m \times k}(\mathbb{R})$ and $S \in M_{k \times n}(\mathbb{R})$;

2. Consider the basis of unit vectors $\left\{ \begin{array}{c} \mathbb{R}^n \\ \text{---} i \text{---} \end{array} \right\}_i$ of \mathbb{R}^n . Then, for each $i \in \{1, \dots, n\}$, denote by \bar{s}_i the i th column of matrix S , namely

$$\begin{array}{c} \mathbb{R}^k \\ \text{---} \bar{s}_i \text{---} \end{array} := \begin{array}{c} \mathbb{R}^k \\ \text{---} S \text{---} \\ \text{---} \mathbb{R}^n \\ \text{---} i \text{---} \end{array}. \quad (44)$$

¹⁴This can be viewed as a map from fragments to data tables, which we denote by $\mathcal{D} : |\text{GPT-Fragment}| \rightarrow |\text{DataTable}| :: f \mapsto D_f$.

¹⁵For data tables generated by real (noisy) experiments, this matrix factorization will not be exact, and one must take a more sophisticated approach to find the best-fit factorization [45, 79].

Similarly, consider the basis of unit covectors $\left\{ \begin{array}{c} \triangle \\ \downarrow \\ \mathbb{R}^m \end{array} \right\}_j$ for \mathbb{R}^m . For each $j \in \{1, \dots, m\}$, denote by \bar{e}_j the j th row of matrix E , i.e.,

$$\begin{array}{c} \triangle \\ \downarrow \\ \mathbb{R}^k \end{array} \bar{e}_j := \begin{array}{c} \triangle \\ \downarrow \\ \mathbb{R}^m \\ \downarrow \\ E \\ \downarrow \\ \mathbb{R}^k \end{array} ; \quad (45)$$

3. Define a GPT $G_D := (\Omega_{G_D}, \mathcal{E}_{G_D}, p_{G_D})$ by

$$\bar{\Omega}_{G_D} := \text{Conv} \left[\left\{ \begin{array}{c} \triangle \\ \downarrow \\ \mathbb{R}^k \\ \downarrow \\ \bar{s}_i \end{array} \right\}_{i=1}^n \right]; \quad (46)$$

$$\mathcal{E}_{G_D} := \text{Conv} \left[\left\{ \begin{array}{c} \triangle \\ \downarrow \\ \mathbb{R}^k \\ \downarrow \\ \bar{e}_j \end{array} \right\}_{j=1}^m \right], \quad (47)$$

where the probability rule p_{G_D} is the evaluation map.

This procedure of theory-agnostic tomography takes as input a data table D and constructs the GPT representations of the states and effects in the GPT system G_D of minimal dimension compatible with D .¹⁶

Note that Ω_{G_D} will necessarily be a valid state space, since (by construction) the data table was generated from a collection of normalised states. Similarly \mathcal{E}_{G_D} will necessarily be a valid effect space, since (by construction) the data table was generated from a collection of effects containing the complement to every effects in the collection, the null effect, and the unit effect.¹⁷

We now prove that theory-agnostic tomography does not directly characterize the GPT fragment f that generated one's data table D_f , but rather the *shadow* of that fragment. That is, the output $\mathcal{S}^{\text{TA}}(f) := G_{D_f}$ of theory-agnostic tomography, applied to a data table generated by fragment f , is a shadow for every fragment f .¹⁸ This was not

¹⁶This defines a map from data tables to GPT systems which we denote by $\mathcal{T} : |\text{DataTable}| \rightarrow |\text{GPT-System}| :: D \mapsto G_D$.

¹⁷If these latter conditions were not satisfied, then in the last stage of theory-agnostic tomography, one would need to add the unit effect and null effect to \mathcal{E}_G and then close under coarse-graining in order to obtain a valid effect space.

¹⁸That is, the composite map $\mathcal{S}^{\text{TA}} := \mathcal{T} \circ \mathcal{D} : |\text{GPT-Fragment}| \rightarrow |\text{GPT-System}|$ acts as a shadow map for every fragment.

explicitly recognized previously, but is critical to understanding when theory-agnostic tomography does and does not work as expected. Prior works did recognize that theory-agnostic tomography is only guaranteed to be accurate under an assumption of tomographic completeness [45, 46, 52]; however, our work demonstrates that what is really required is *relative* tomographic completeness, and moreover formalizes what precisely goes wrong when this assumption fails. (Incidentally, the following theorem also constitutes a formal proof that the output of GPT tomography is always a valid GPT system.)

Theorem 22. Consider the data table generated by some GPT fragment. The output of theory-agnostic tomography (as in Definition 21) applied to this data table is the GPT shadow of the GPT fragment. That is, for all fragments f we have that $\mathcal{S}^{\text{TA}}(f) = G_{D_f}$ is a valid shadow of f .

We prove this theorem in Appendix E.

Consequently, correct assessments of noncontextuality occur if and only if the simplex-embeddability property of the GPT fragment describing one's experiment is identical to the simplex-embeddability property of its shadow.

6.2 Assessments of (non)contextuality for GPT subsystems of the true GPT are always accurate

A critical result is the following.

Theorem 23. Consider a fragment f of the true GPT system G . If f is a GPT subsystem of G , then the shadow of f is nonclassical if and only if f is nonclassical.

The proof is simple. Since f is a GPT subsystem, its states and effects must be relatively tomographic; consequently, f is equivalent to its shadow (by Lemma 19), and so either both are classical (simplex-embeddable), or neither is.

Together with the transitivity of GPT subsystems (Proposition 1), this implies that proving nonclassicality for a GPT subsystem is sufficient to establish nonclassicality of the full GPT system containing it. That such a result should hold for subsystems in the sense of the parallel composition operation \otimes is a basic consistency constraint that any sensible notion of classicality must satisfy, as it means that GPT tensor product subsystems of classically explainable (simplex-embeddable) GPTs are

classically explainable (simplex-embeddable).¹⁹ However, these results also apply to the wide class of other GPT subsystems defined and discussed in Section 3, including superselection sectors of a composite, virtual subsystems, stabilizer subsystems, and so on.

In summary: in noncontextuality experiments, what matters is not tomographic completeness of states and effects in some absolute sense, but tomographic completeness of states and effects *relative to each other*. They need not be tomographically complete for the fundamental system that causally mediates correlations between the preparations and measurements.

Thus, for example, one can genuinely prove nonclassicality using states and effects of a qubit even without ever accessing the Pauli- Y axis of the Bloch ball. Problems could only arise if either the states or the effects—but *not both*—had nontrivial expectation value on Y . One could similarly prove nonclassicality using some subset of the modes of an oscillator. In more extreme examples, one might even be probing some collective (coarse-grained) degree of freedom of a massive collection of more fundamental systems. This too poses no problem whatsoever to the goal of accurately assessing nonclassicality, as long as one has the ability to single out some set of states and effects on the *same* GPT subsystem—i.e., that are relatively tomographic.

Most importantly, this explains why we expect assessments of noncontextuality to hold up even if our current theories (e.g., quantum theory) are emergent rather than fundamental. We return to this point in Section 6.6.

6.3 Proofs of classical-explainability are robust to arbitrary failures of tomographic completeness

Recall from Proposition 4 that shadow maps cannot take simplex-embeddable fragments to non-simplex-embeddable fragments.

Thus, for *any* set of states and effects, a proof of classical-explainability is robust. So if one does theory-agnostic tomography and

¹⁹This also helps one understand the lab notebook argument of Ref. [46]: there, establishing nonclassicality for a physical system is sufficient for establishing nonclassicality for that system taken together with any other physical systems, including those keeping classical records of what experimental procedures were implemented in the experiment.

finds a classically-explainable GPT, one can be certain that the actual experiment is indeed classically-explainable. This is typically much less useful than, e.g., Theorem 23, since one typically wishes to certify *nonclassicality* rather than classicality. And it may still be the case that the *full* GPT system in question is nonclassical, but that this can only be witnessed by a more comprehensive experiment.

6.4 Necessary conditions for proofs of nonclassicality to be robust to failures of tomographic completeness?

Given this, the natural question becomes to characterize *exactly* when shadow maps do and do not cause simplex-embeddable fragments to become non-simplex-embeddable.²⁰ Can one give necessary and sufficient conditions (presumably geometric ones) for a GPT fragment to be simplex-embeddable if and only if its shadow is? This would constitute necessary and sufficient conditions for one’s assessments of nonclassicality to be correct.

By Theorem 23, a *sufficient* condition for simplex-embedding to be preserved by a shadow map is that the fragment in question is relatively tomographic—a GPT subsystem of the true GPT system. But this is not a necessary condition, as evidenced by the example in Section 5.4, where we exhibited a nontomographic fragment that was simplex-embeddable and whose shadow was also simplex-embeddable (despite being quite different from the fragment itself).

We suspect that the results of Ref. [53] (showing that one can certify proofs of nonclassicality even in the presence of spurious operational equivalences) are also closely connected to the question raised in this section. We leave this connection to future work.

6.5 Another useful perspective: operational identities for GPT subsystems are always genuine

In this section, we essentially repeat the conclusions of Section 6.2, but in a language that may be useful for readers who are more familiar with noncontextuality no-go theorems based on

²⁰A still more ambitious question would be to characterize how taking the shadow map *quantitatively* [57] affects simplex-embedding.

operational identities (or operational equivalences) than those based on simplex-embedding of GPTs.

Imagine one attempts to witness the failure of noncontextuality in an experiment on some particular physical system whose true GPT description is G . In any real experiment, one will only access a fragment of the states and effects of the system. From the statistics observed in the experiment, one must infer some operational identities. If one wishes to use the statistics of such an experiment to prove the failure of noncontextuality, then it is often said that one must implement a set of states and effects that are tomographically complete for the system.

But in fact, one can base such proofs on a strictly weaker assumption—that the operational identities used in one’s proof are valid identities in the true GPT (G). This is simply because Leibniz’s principle (in the methodological form introduced by Spekkens [2]) can be applied to any processes that are indistinguishable in principle. Showing the indistinguishability of two processes relative to a set of procedures that are tomographically complete for the physical system in one’s experiment is a *sufficient but not necessary* condition for in principle indistinguishability.

For instance, consider any proof of nonclassicality built on top of operational equivalences holding among states of a single quantum bit. Such proofs do not suddenly fail to be valid simply because one implements them experimentally using a physical qubit that is a tensor subsystem of a pair of qubits. This is an obvious prerequisite to proofs of nonclassicality being meaningful in the first place, since all systems are subsystems of some larger system (e.g., of the universe). But there are less trivial examples as well. One could equally well implement a proof of nonclassicality using a virtual qubit defined as a coarse-grained degree of freedom of some large collection of physical qubits. Or, one could use only the states and effects lying in a plane of the Bloch sphere (i.e., using a rebit).

Since the operational identities used in one’s proof are necessarily extracted from one’s experimental data, the question becomes, under what conditions are the apparent operational identities actually valid in the true GPT? Using the notion of a GPT subsystem, one can give a simple necessary and sufficient condition.

Lemma 24. *Consider a fragment f of some*

fundamental GPT system G . If f is a GPT subsystem, then all operational identities that hold for its states (effects) relative to its effects (states) are genuine. That is, they hold also relative to the full set of effects (states) in the true GPT system G .

This is because f being a GPT subsystem implies that its states and effects are relatively tomographic, and so can be discriminated with *or without* the effects and states in G that are not also in f .

For instance, consider again an experiment on a qubit that only accesses states and effects in the rebit. Although these states and effects would not be tomographically complete for the physical system (the qubit), this would not lead to any invalid operational equivalences, since the rebit states and rebit effects are relatively tomographic.

6.6 Theories that hyperdecohere to quantum theory are nonclassical

State-of-the-art proofs of the failure of noncontextuality are theory-independent, in the sense that they do not assume the validity of quantum theory. Thus, like Bell’s theorem, they provide a constraint on nature, independently of whether or not some future theory supersedes quantum theory. (See Section IV.B of Ref. [46] for a more detailed discussion of the similarities between these two kinds of arguments.)

However, one might be concerned that this theory-independence is undermined by the fact that proofs of noncontextuality (unlike proofs of Bell nonclassicality) rely on a characterization of the GPT states and effects governing one’s experiment—or equivalently, rely on particular operational identities that are known to hold among the GPT processes in one’s experiment. If quantum theory is superseded by some future theory, then might it be that one’s characterization of one’s laboratory procedures as GPT processes are mistaken (and consequently, that the operational identities one believed held among them were also mistaken)? If this did happen, then current assessments of nonclassicality would not necessarily hold up as our theories of physics develop over time.

A competing intuition, however, is that any empirically successful future theory must contain quantum theory in some sense, and so must also exhibit the same forms of nonclassicality (and

potentially more). In the following, we use the notion of GPT subsystems to show that this intuition can be made precise, and we show that as a consequence, current assessments of nonclassicality *will* remain unchanged even if our best theories of nature evolve in the future, if a well-motivated assumption holds.

The key argument was identified already by Müller and Garner in Lemma 11 of Ref. [47].

Any future successor to quantum theory will need to reproduce quantum theory in an appropriate limit, in order to explain why we have not seen beyond-quantum phenomena in any of our current experiments. Just as quantum phenomena are suppressed by decoherence in appropriate regimes, leading to the appearance of classical theory, so too we expect that such beyond-quantum phenomena would be suppressed (in appropriate regimes) by some generalized kind of decoherence process which leads to the appearance of quantum theory. A general notion of this kind, described within the framework of GPTs, was introduced in Refs. [70–73], and is known as hyperdecoherence.

Formally, a hyperdecoherence process on a GPT G is an idempotent, discard-preserving linear map H that takes states/effects in the GPT to states/effects in the GPT. Recall that idempotence of H simply means that $H \circ H = H$. All of these features of a hyperdecoherence map are of course satisfied by a standard decoherence process taking quantum theory to classical theory.

Given some initial GPT system G (where for simplicity we work with the standard representation where effects live in the dual vector space), the hyperdecohered GPT system can be defined as the set of all states/effects in the image of the hyperdecoherence map H , namely

$$\left\{ \begin{array}{c} \textcircled{G} \\ \boxed{H} \\ \textcircled{G} \\ \textcircled{S} \end{array} \right\}_{s \in \Omega_G} \quad \text{and} \quad \left\{ \begin{array}{c} \textcircled{E} \\ \textcircled{G} \\ \boxed{H} \\ \textcircled{G} \end{array} \right\}_{e \in \mathcal{E}_G}. \quad (48)$$

These hyperdecohered states and effects are always tomographic for each other, and so form a valid GPT system, as was proven, for example, in Appendix D.2 of Ref. [80].

Having defined the hyperdecohered theory, we now prove the key result of this section.

Theorem 25. *Consider a GPT system G and the hyperdecohered GPT system $H(G)$ defined by any given hyperdecoherence process on it. Then $H(G)$ is a GPT subsystem of the original GPT system G .*

The proof is simply that each process in the theory is already itself a process in G , by the assumption that H is a physical transformation in the theory²¹. Consequently, one can simply take the state map to be identity and the effect map to be identity in order to define the embedding of $H(G)$ into G .

As mentioned, the canonical example of hyperdecoherence is standard decoherence in quantum theory, e.g., taking a qubit to a classical bit. And as mentioned in Section 3, the classical bit is indeed a GPT subsystem of the qubit, in keeping with Theorem 25.

An immediate corollary of Theorem 25, first proved in Ref. [47], is the following:

Corollary 1. *Any theory which hyperdecoheres to quantum theory necessarily fails to admit of a generalised noncontextual ontological model.*

This is the sense in which every proof of the failure of noncontextuality that assumes quantum theory will remain valid even if quantum theory itself is superseded by a later successor theory.

If the manner by which quantum theory is recovered from some future successor theory is too exotic to fit the (quite general) mold of a hyperdecoherence process, then one would need to reassess whether such a beyond-quantum theory necessarily retains the same nonclassical features as quantum theory. This is analogous to how if one considers an exotic enough future theory (e.g., one which modified the nature of space and time), one would need to reassess the question of whether such a theory indeed was nonclassical in the sense seemingly implied by Bell’s theorem. But in both cases, the accuracy of the conclusions one draws about nonclassicality are quite theory-independent—backed up by general assumptions that plausibly must hold in any reasonable future theory.

7 Future directions

It would be interesting to see whether GPT subsystems subsume still other notions of systems that have been considered. Perhaps of special interest are attempts to define a notion of a system

²¹The fact that a hyperdecoherence process takes states/effects in the GPT system to states/effects in the GPT system is motivated by the idea that the process must itself be a physical transformation in the GPT.

(or subsystem) that is operationally defined rather than a primitive notion [81–83].

In forthcoming work, we generalize the idea of a GPT subsystem to the broader notion of a *GPT subtheory*, capturing the idea of what it means for a *theory* to live inside another. The definition is quite analogous to that given herein for GPT subsystems: one GPT is a GPT subtheory of another if there exists a structure-preserving embedding from the former into the latter. The novelty is that once one considers GPTs as full theories, there is compositional structure in addition to convex structure to be preserved.

Another important direction for future work is to consider how the shadow map *quantitatively* affects simplex-embeddability. In this work, we focused merely on the qualitative divide between simplex-embeddable or not. We also suspect that combining the current tools with those of Ref. [53] could lead to further progress on understanding when and how failures of relative tomography leads to mistaken assessments of nonclassicality. Note that our approach in this work was to investigate what happens to a true GPT description in the presence of limited experimental tests; Ref. [53], in contrast, can be viewed as characterizing what can be said about the scope of all possible true GPT descriptions of an experiment, given a constraint on how far off one’s set of processes is from being relatively tomographic.

Although we have given a simple sufficient condition (relative tomography of one’s preparations and measurements) for assessments of noncontextuality to be correct, this condition is not something one can experimentally guarantee (although one can certainly gather experimental evidence for it [45, 79]). It remains an open question whether one can find other sufficient conditions, and which of these conditions might be easiest to test experimentally. Or, might there even be physical conditions under which relative tomography is guaranteed? Very speculatively, one might obtain such a guarantee if one had a means to generate (all possible) physical symmetry transformations in some GPT.

Finally, we note that our conclusions regarding the relevant assumptions for witnessing the failure of generalized noncontextuality are also relevant to operational proofs of the failure of Kochen-Specker noncontextuality [84], since the Kochen-Specker assumptions are best motivated

by the same considerations, namely appeal to Leibniz’s principle [2].

Acknowledgements

DS thanks Rob Spekkens and Elie Wolfe for useful discussions regarding when proofs of nonclassicality may or may not be robust. Early inspiration for this work came from discussions with Rob Spekkens relating to proofs of nonclassicality on a rebit embedded in a qubit. We thank Markus P. Müller, Thomas Galley, Caroline L. Jones, and Howard Barnum for a delightful and insightful visit to IQOQI-Vienna where many ideas related to this manuscript were discussed. We thank Markus P. Müller especially for helping us recognize that relative tomography of states and effects cannot be captured by simply looking at the spans of their vector spaces, and in particular for coming up with the example in Section 4.4. We thank also Yilè Yīng, Victor Gitton, and Marco Erba for relevant discussions. Finally, we thank the referees for carefully reading this work and providing helpful feedback. JHS and DS were supported by the National Science Centre, Poland (Opus project, Categorical Foundations of the Non-Classicality of Nature, project no. 2021/41/B/ST2/03149). VR and RB acknowledge support by the Digital Horizon Europe project FoQaCiA, Foundations of quantum computational advantage, GA No. 101070558, funded by the European Union, NSERC (Canada), and UKRI (UK). This work is partially carried out under IRA Programme, project no. FENG.02.01-IP.05-0006/23, financed by the FENG program 2021-2027, Priority FENG.02, Measure FENG.02.01., with the support of the FNP. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development and by the Province of Ontario through the Ministry of Colleges and Universities. Some diagrams were prepared using TikZit and Mathcha.

References

- [1] R. W. Spekkens. “Contextuality for preparations, transformations, and unsharp

- measurements”. *Phys. Rev. A* **71**, 052108 (2005).
- [2] Robert W. Spekkens. “The ontological identity of empirical indiscernibles: Leibniz’s methodological principle and its significance in the work of Einstein” (2019). [arXiv:1909.04628](#).
 - [3] David Schmid. “Guiding our interpretation of quantum theory by principles of causation and inference”. *PhD thesis, University of Waterloo* (2021).
 - [4] Robert W. Spekkens. “Negativity and Contextuality are Equivalent Notions of Nonclassicality”. *Phys. Rev. Lett.* **101**, 020401 (2008).
 - [5] David Schmid, John H. Selby, Matthew F. Pusey, and Robert W. Spekkens. “A structure theorem for generalized-noncontextual ontological models”. *Quantum* **8**, 1283 (2024).
 - [6] David Schmid, John H. Selby, Elie Wolfe, Ravi Kunjwal, and Robert W. Spekkens. “Characterization of Noncontextuality in the Framework of Generalized Probabilistic Theories”. *PRX Quantum* **2**, 010331 (2021).
 - [7] David Schmid, Haoxing Du, John H. Selby, and Matthew F. Pusey. “Uniqueness of noncontextual models for stabilizer subtheories”. *Phys. Rev. Lett.* **129**, 120403 (2022).
 - [8] Farid Shahandeh. “Quantum computational advantage implies contextuality” (2021). [arXiv:2112.00024](#).
 - [9] David Schmid and Robert W. Spekkens. “Contextual advantage for state discrimination”. *Phys. Rev. X* **8**, 011015 (2018).
 - [10] Kieran Flatt, Hanwool Lee, Carles Roch I Carceller, Jonatan Bohr Brask, and Joonwoo Bae. “Contextual advantages and certification for maximum-confidence discrimination”. *PRX Quantum* **3**, 030337 (2022).
 - [11] Sumit Mukherjee, Shivam Naonit, and A. K. Pan. “Discriminating three mirror-symmetric states with a restricted contextual advantage”. *Phys. Rev. A* **106**, 012216 (2022).
 - [12] Jaehee Shin, Donghoon Ha, and Younghun Kwon. “Quantum contextual advantage depending on nonzero prior probabilities in state discrimination of mixed qubit states”. *Entropy* **23**, 1583 (2021).
 - [13] Lorenzo Catani, Matthew Leifer, David Schmid, and Robert W. Spekkens. “Why interference phenomena do not capture the essence of quantum theory”. *Quantum* **7**, 1119 (2023).
 - [14] Lorenzo Catani, Matthew Leifer, David Schmid, and Robert W. Spekkens. “Reply to ‘Comment on ‘Why interference phenomena do not capture the essence of quantum theory’”” (2022). [arXiv:2207.11791](#).
 - [15] Lorenzo Catani, Matthew Leifer, Giovanni Scala, David Schmid, and Robert W. Spekkens. “Aspects of the phenomenology of interference that are genuinely nonclassical”. *Phys. Rev. A* **108**, 022207 (2023).
 - [16] John H. Selby, David Schmid, Elie Wolfe, Ana Belén Sainz, Ravi Kunjwal, and Robert W. Spekkens. “Contextuality without incompatibility”. *Phys. Rev. Lett.* **130**, 230201 (2023).
 - [17] John H. Selby, David Schmid, Elie Wolfe, Ana Belén Sainz, Ravi Kunjwal, and Robert W. Spekkens. “Accessible fragments of generalized probabilistic theories, cone equivalence, and applications to witnessing nonclassicality”. *Phys. Rev. A* **107**, 062203 (2023).
 - [18] Armin Tavakoli and Roope Uola. “Measurement incompatibility and steering are necessary and sufficient for operational contextuality”. *Phys. Rev. Research* **2**, 013011 (2020).
 - [19] Lorenzo Catani, Matthew Leifer, Giovanni Scala, David Schmid, and Robert W. Spekkens. “What is nonclassical about uncertainty relations?”. *Phys. Rev. Lett.* **129**, 240401 (2022).
 - [20] Matteo Lostaglio. “Certifying Quantum Signatures in Thermodynamics and Metrology via Contextuality of Quantum Linear Response”. *Phys. Rev. Lett.* **125**, 230603 (2020).
 - [21] Naim Elias Comar, Danilo Cius, Luis Felipe Santos, Rafael Wagner, and Bárbara Amaral.

- “Contextuality in anomalous heat flow” (2024). [arXiv:2406.09715](#).
- [22] Matthew F. Pusey. “Anomalous Weak Values Are Proofs of Contextuality”. *Phys. Rev. Lett.* **113**, 200401 (2014).
- [23] Ravi Kunjwal, Matteo Lostaglio, and Matthew F. Pusey. “Anomalous weak values and contextuality: Robustness, tightness, and imaginary parts”. *Phys. Rev. A* **100**, 042116 (2019).
- [24] Vinicius P. Rossi, David Schmid, John H. Selby, and Ana Belén Sainz. “Contextuality with vanishing coherence and maximal robustness to dephasing”. *Phys. Rev. A* **108**, 032213 (2023).
- [25] Rafael Wagner, Anita Camillini, and Ernesto F. Galvão. “Coherence and contextuality in a Mach-Zehnder interferometer”. *Quantum* **8**, 1240 (2024).
- [26] Roberto D. Baldijão, Rafael Wagner, Cristhiano Duarte, Bárbara Amaral, and Marcelo Terra Cunha. “Emergence of Noncontextuality under Quantum Darwinism”. *PRX Quantum* **2**, 030351 (2021).
- [27] Robert W. Spekkens, D. H. Buzacott, A. J. Keehn, Ben Toner, and G. J. Pryde. “Preparation Contextuality Powers Parity-Oblivious Multiplexing”. *Phys. Rev. Lett.* **102**, 010401 (2009).
- [28] Andre Chailloux, Iordanis Kerenidis, Srijita Kundu, and Jamie Sikora. “Optimal bounds for parity-oblivious random access codes”. *New J. Phys.* **18**, 045003 (2016).
- [29] Andris Ambainis, Manik Banik, Anubhav Chaturvedi, Dmitry Kravchenko, and Ashutosh Rai. “Parity oblivious d-level random access codes and class of noncontextuality inequalities”. *Quantum Inf. Process.* **18**, 111 (2019).
- [30] Debashis Saha, Paweł Horodecki, and Marcin Pawłowski. “State independent contextuality advances one-way communication”. *New J. Phys.* **21**, 093057 (2019).
- [31] Shiv Akshar Yadavalli and Ravi Kunjwal. “Contextuality in entanglement-assisted one-shot classical communication”. *Quantum* **6**, 839 (2022).
- [32] Alley Hameedi, Armin Tavakoli, Breno Marques, and Mohamed Bourennane. “Communication games reveal preparation contextuality”. *Phys. Rev. Lett.* **119**, 220402 (2017).
- [33] Amanda M. Fonseca, Vinicius P. Rossi, Roberto D. Baldijão, John H. Selby, and Ana Belén Sainz. “Robustness of contextuality under different types of noise as quantifiers for parity-oblivious multiplexing tasks”. *Phys. Rev. A* **111**, 022217 (2025).
- [34] Matteo Lostaglio and Gabriel Senno. “Contextual advantage for state-dependent cloning”. *Quantum* **4**, 258 (2020).
- [35] Pauli Jokinen, Mirjam Weilenmann, Martin Plávala, Juha-Pekka Pellonpää, Jukka Kiukas, and Roope Uola. “No-broadcasting characterizes operational contextuality” (2024). [arXiv:2406.07305](#).
- [36] Victoria J. Wright and Máté Farkas. “Invertible map between bell nonlocal and contextuality scenarios”. *Phys. Rev. Lett.* **131**, 220202 (2023).
- [37] David Schmid, John H. Selby, and Robert W. Spekkens. “Unscrambling the omelette of causation and inference: The framework of causal-inferential theories” (2021). [arXiv:2009.03297](#).
- [38] Ravi Kunjwal and Robert W. Spekkens. “From the Kochen-Specker Theorem to Noncontextuality Inequalities without Assuming Determinism”. *Phys. Rev. Lett.* **115**, 110403 (2015).
- [39] Ravi Kunjwal and Robert W. Spekkens. “From statistical proofs of the kochen-specker theorem to noise-robust noncontextuality inequalities”. *Phys. Rev. A* **97**, 052110 (2018).
- [40] Ravi Kunjwal. “Contextuality beyond the Kochen-Specker theorem” (2016). [arXiv:1612.07250](#).
- [41] Ravi Kunjwal. “Beyond the Cabello-Severini-Winter framework: Making sense of contextuality without sharpness of measurements”. *Quantum* **3**, 184 (2019).
- [42] Ravi Kunjwal. “Hypergraph framework for irreducible noncontextuality inequalities from logical proofs of the Kochen-Specker theorem”. *Quantum* **4**, 219 (2020).

- [43] Ravi Kunjwal. “Fine’s theorem, noncontextuality, and correlations in specker’s scenario”. *Phys. Rev. A* **91**, 022108 (2015).
- [44] Tomáš Gonda, Ravi Kunjwal, David Schmid, Elie Wolfe, and Ana Belén Sainz. “Almost Quantum Correlations are Inconsistent with Specker’s Principle”. *Quantum* **2**, 87 (2018).
- [45] Michael D. Mazurek, Matthew F. Pusey, Kevin J. Resch, and Robert W. Spekkens. “Experimentally Bounding Deviations From Quantum Theory in the Landscape of Generalized Probabilistic Theories”. *PRX Quantum* **2**, 020302 (2021).
- [46] David Schmid, John H. Selby, and Robert W. Spekkens. “Addressing some common objections to generalized noncontextuality”. *Phys. Rev. A* **109**, 022228 (2024).
- [47] Markus P. Müller and Andrew J. P. Garner. “Testing quantum theory by generalizing noncontextuality”. *Phys. Rev. X* **13**, 041001 (2023).
- [48] L. Hardy. “Quantum Theory From Five Reasonable Axioms” (2001). [arXiv:quant-ph/0101012](https://arxiv.org/abs/quant-ph/0101012).
- [49] Jonathan Barrett. “Information processing in generalized probabilistic theories”. *Phys. Rev. A* **75**, 032304 (2007).
- [50] Peter Janotta and Raymond Lal. “Generalized probabilistic theories without the no-restriction hypothesis”. *Phys. Rev. A* **87**, 052131 (2013).
- [51] Andrew J. P. Garner and Markus P. Mueller. “Characterization of the probabilistic models that can be embedded in quantum theory” (2020). [arXiv:2004.06136](https://arxiv.org/abs/2004.06136).
- [52] Michael D Mazurek, Matthew F Pusey, Ravi Kunjwal, Kevin J Resch, and Robert W Spekkens. “An experimental test of noncontextuality without unphysical idealizations”. *Nature communications* **7**, 1–7 (2016).
- [53] Matthew F. Pusey, Lídia del Rio, and Bettina Meyer. “Contextuality without access to a tomographically complete set” (2019). [arXiv:1904.08699](https://arxiv.org/abs/1904.08699).
- [54] David Schmid, Roberto D Baldijão, John H Selby, Ana Belén Sainz, and Robert W Spekkens. “Noncontextuality inequalities for prepare-transform-measure scenarios” (2024). [arXiv:2407.09624](https://arxiv.org/abs/2407.09624).
- [55] Lluís Masanes and Markus P Müller. “A derivation of quantum theory from physical requirements”. *New Journal of Physics* **13**, 063001 (2011).
- [56] Martin Plávala. “General probabilistic theories: An introduction” (2021). [arXiv:2103.07469](https://arxiv.org/abs/2103.07469).
- [57] John H. Selby, Elie Wolfe, David Schmid, Ana Belén Sainz, and Vinicius P. Rossi. “Linear program for testing nonclassicality and an open-source implementation”. *Phys. Rev. Lett.* **132**, 050202 (2024).
- [58] A. J. Leggett and Anupam Garg. “Quantum mechanics versus macroscopic realism: Is the flux there when nobody looks?”. *Phys. Rev. Lett.* **54**, 857–860 (1985).
- [59] David Schmid. “A review and reformulation of macroscopic realism: resolving its deficiencies using the framework of generalized probabilistic theories”. *Quantum* **8**, 1217 (2024).
- [60] Daniel Gottesman. “Stabilizer codes and quantum error correction” (1997). [arXiv:quant-ph/9705052](https://arxiv.org/abs/quant-ph/9705052).
- [61] Daniel Gottesman. “The Heisenberg representation of quantum computers” (1998). [arXiv:quant-ph/9807006](https://arxiv.org/abs/quant-ph/9807006).
- [62] Ernst CG Stueckelberg. “Quantum theory in real hilbert space”. *Helv. Phys. Acta* **33**, 458 (1960).
- [63] Carlton M Caves, Christopher A Fuchs, and Pranaw Rungta. “Entanglement of formation of an arbitrary state of two rebits”. *Foundations of Physics Letters* **14**, 199–212 (2001).
- [64] Howard Barnum and Alexander Wilce. “Information Processing in Convex Operational Theories”. *Elec. Notes Theo. Comp. Sci.* **270**, 3–15 (2011).
- [65] Stephen D Bartlett, Terry Rudolph, and Robert W Spekkens. “Reference frames, superselection rules, and quantum information”. *Reviews of Modern Physics* **79**, 555 (2007).

- [66] Paolo Zanardi. “Virtual quantum subsystems”. *Physical Review Letters* **87**, 077901 (2001).
- [67] Benjamin Schumacher. “Quantum coding”. *Phys. Rev. A* **51**, 2738–2747 (1995).
- [68] Daniel A. Lidar and K. Birgitta Whaley. “Decoherence-free subspaces and subsystems”. Pages 83–120. Springer Berlin Heidelberg. Berlin, Heidelberg (2003).
- [69] Ognyan Oreshkov. “Time-delocalized quantum subsystems and operations: on the existence of processes with indefinite causal structure in quantum mechanics”. *Quantum* **3**, 206 (2019).
- [70] Karol Życzkowski. “Quartic quantum theory: an extension of the standard quantum mechanics”. *Journal of Physics A: Mathematical and Theoretical* **41**, 355302 (2008).
- [71] Jonathan G. Richens, John H. Selby, and Sabri W. Al-Safi. “Entanglement is necessary for emergent classicality in all physical theories”. *Phys. Rev. Lett.* **119**, 080503 (2017).
- [72] Ciarán M Lee and John H Selby. “A no-go theorem for theories that decohere to quantum mechanics”. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **474**, 20170732 (2018).
- [73] James Hefford and Stefano Gogioso. “Hyper-decoherence in density hypercubes”. *Electronic Proceedings in Theoretical Computer Science* **340**, 141–159 (2021).
- [74] Giulio Chiribella, Giacomo Mauro D’Ariano, and Paolo Perinotti. “Probabilistic theories with purification”. *Phys. Rev. A* **81**, 062348 (2010).
- [75] Theodore W. Gamelin and Robert Everist Greene. “Introduction to topology”. *Dover Publications*. Mineola, N.Y. (1999, 2nd ed.).
- [76] Victor Gitton and Mischa P. Woods. “Solvable criterion for the contextuality of any prepare-and-measure scenario”. *Quantum* **6**, 732 (2022).
- [77] Alexander Holevo. “Probabilistic and statistical aspects of quantum theory”. *Edizioni della Normale*. (2011).
- [78] E G Beltrametti and S Bugajski. “A classical extension of quantum mechanics”. *J. Phys. A* **28**, 3329–3343 (1995).
- [79] Michael J. Grabowecky, Christopher A. J. Pollack, Andrew R. Cameron, Robert W. Spekkens, and Kevin J. Resch. “Experimentally bounding deviations from quantum theory for a photonic three-level system using theory-agnostic tomography”. *Phys. Rev. A* **105**, 032204 (2022).
- [80] Daniel Centeno, Marco Erba, David Schmid, John H. Selby, Robert W. Spekkens, Sina Soltani, Jacopo Surace, Alex Wilce, and Yìlè Yīng. “Twirled worlds: symmetry-induced failures of tomographic locality” (2024). [arXiv:2407.21688](https://arxiv.org/abs/2407.21688).
- [81] Giulio Chiribella. “Agents, subsystems, and the conservation of information”. *Entropy* **20**, 358 (2018).
- [82] Lea Krämer and Lídia Del Rio. “Operational locality in global theories”. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **376**, 20170321 (2018).
- [83] Victor Gitton and Mischa P. Woods. “On the system loophole of generalized noncontextuality” (2022). [arXiv:2209.04469](https://arxiv.org/abs/2209.04469).
- [84] Simon Kochen and Ernst Specker. “The problem of hidden variables in quantum mechanics”. *J. Math. & Mech.* **17**, 59–87 (1967).
- [85] Charalambos D Aliprantis. “Principles of real analysis”. *Academic Press*. San Diego, CA (1998). 3 edition.
- [86] Robert Lang. “A note on the measurability of convex sets”. *Archiv der Mathematik* **47**, 90–92 (1986).
- [87] Keith Conrad. “Lecture Notes on Bilinear Forms”.

A Proof of Lemma 16

Lemma 16. *Any embedding (ι, κ) of a GPT system (Ω, \mathcal{E}, p) into itself is a GPT isomorphism.*

Proof. To begin, recall that every GPT embedding is faithful (Def. 7), which means that ι and κ are injective state and effect maps. Moreover, they are injective as linear maps $\iota: \text{Span}[\Omega] \rightarrow \text{Span}[\Omega]$ and

$\kappa : \text{Span}[\mathcal{E}] \rightarrow \text{Span}[\mathcal{E}]$, so they automatically have a linear left inverse. All that remains to be shown is that they are also surjective (and hence invertible) as state and effect maps. In other words, $\iota(\Omega) = \Omega$ and $\kappa(\mathcal{E}) = \mathcal{E}$.

First, we show that ι and κ cannot change the volume of Ω and \mathcal{E} . Fix a Lebesgue measure on $\text{Span}[\Omega]$ and write $\text{vol}()$ for the corresponding volume (and analogously for $\text{Span}[\mathcal{E}]$). Note that

$$\text{vol}[\iota(\Omega)] = \det[\iota] \text{vol}[\Omega], \quad (49)$$

$$\text{vol}[\kappa(\mathcal{E})] = \det[\kappa] \text{vol}[\mathcal{E}], \quad (50)$$

since the determinant of a linear map specifies the scaling factor of volume for any measurable full-dimensional convex set (see [85, Lemma 40.4, p. 389]).²² The conditions $\iota(\Omega) \subseteq \Omega$ and $\kappa(\mathcal{E}) \subseteq \mathcal{E}$ imply that the volumes can only decrease, so $\det[\iota] \leq 1$ and $\det[\kappa] \leq 1$. Now, recall that p is tomographic, which means that it constitutes a nondegenerate bilinear pairing, $p : \text{Span}[\Omega] \times \text{Span}[\mathcal{E}] \rightarrow \mathbb{R}$. That is,

$$p(s, e) = 0 \quad \forall e \iff s = 0, \quad (51)$$

and

$$p(s, e) = 0 \quad \forall s \iff e = 0. \quad (52)$$

We can therefore define a transpose relative to p , $*_p$, such that for any linear map $f : \text{Span}[\Omega] \rightarrow \text{Span}[\Omega]$, there exists another linear map $f^{*p} : \text{Span}[\mathcal{E}] \rightarrow \text{Span}[\mathcal{E}]$ such that

$$p(f(s), e) = p(s, f^{*p}(e)) \quad \forall s, e. \quad (53)$$

Tomography of p and probability preservation of embeddings lead to

$$p(s, e) = p(\iota(s), \kappa(e)) = p(s, \iota^{*p}(\kappa(e))) \quad \forall s, e, \quad (54)$$

which implies that

$$\iota^{*p}(\kappa(e)) = e, \quad \forall e. \quad (55)$$

From this it follows that

$$\iota^{*p} \circ \kappa = \text{Id}_{\text{Span}(\mathcal{E})}, \quad \text{and} \quad \kappa = (\iota^{*p})^{-1}. \quad (56)$$

Therefore, $\det[\kappa] = \det[(\iota^{*p})^{-1}] = \det[\iota^{*p}]^{-1}$, which together with the fact that $\det[\iota^{*p}] = \det[\iota]$ (see

²²Here, our state and effect spaces are convex, and therefore measurable with respect to Lebesgue measure [86, Thm. 2].

Ref. [87, Theorem 3.21, p. 21]) implies that $\det[\kappa] = \det[\iota]^{-1}$. Together with $\det[\iota] \leq 1$ and $\det[\kappa] \leq 1$, we get $\det[\iota] = \det[\kappa] = 1$. So indeed $\text{vol}[\iota(\Omega)] = \text{vol}[\Omega]$ and $\text{vol}[\kappa(\mathcal{E})] = \text{vol}[\mathcal{E}]$.

The second step is to note that $\iota(\Omega) \supseteq \text{int}(\Omega)$, where the interior is taken within $\text{Span}[\Omega]$ (i.e., the relative interior). Indeed, if any relative interior point $x \in \Omega$ were not in $\iota(\Omega)$, there would exist a sufficiently small full-dimensional ball around x contained in Ω but disjoint from $\iota(\Omega)$. This would imply a strictly smaller volume for $\iota(\Omega)$, contradicting the equality above. Thus every interior point of Ω lies in $\iota(\Omega)$. Moreover, since both $\iota(\Omega)$ and Ω are closed we have that $\iota(\Omega) = \overline{\iota(\Omega)} \supseteq \overline{\text{int}(\Omega)} = \Omega$, where $\overline{\iota(\Omega)}$ and $\overline{\text{int}(\Omega)}$ denote the closures of these sets. Finally, recalling that $\iota(\Omega) \subseteq \Omega$ (since ι is a state map from Ω to itself), we conclude $\iota(\Omega) = \Omega$. A similar argument shows that $\kappa(\mathcal{E}) = \mathcal{E}$.

Therefore, ι and κ are surjective state and effect maps, and the pair (ι, κ) constitutes a GPT isomorphism.

B Proof of Theorem 17

First, we state and prove three useful lemmas.

Lemma 26. *Shadow maps are necessarily linear.*

Proof. Consider first a state $s \in \Omega_f$ of a GPT fragment f , such that

$$\downarrow \frac{f}{s} = \sum_i \alpha_i \downarrow \frac{f}{s_i}, \quad (57)$$

where $\{\alpha_i\}$ are real numbers and s_i are arbitrary states in Ω_f . Then for any effect $e \in \mathcal{E}_f$,

$$\downarrow \frac{e}{p_f} = \downarrow \frac{e}{p_s} \downarrow \frac{f}{s} = \downarrow \frac{e}{p_s} \downarrow \frac{f}{s_i} \alpha_i. \quad (58)$$

Now, we appeal to the bilinearity of the map p_f . It will follow that

$$\downarrow \frac{e}{p_f} = \downarrow \frac{e}{p_s} \downarrow \frac{f}{s} = \downarrow \frac{e}{p_s} \downarrow \frac{f}{s_i} \alpha_i. \quad (59)$$

$$= \begin{array}{c} \triangleleft e \\ \downarrow f \\ \boxed{p_f} \\ \downarrow f \\ \boxed{\sum_i \alpha_i} \\ \downarrow f \\ \triangleleft s_i \end{array} \quad (60)$$

$$= \sum_i \alpha_i \begin{array}{c} \triangleleft e \\ \downarrow f \\ \boxed{p_f} \\ \downarrow f \\ \triangleleft s_i \end{array} \quad (61)$$

$$= \sum_i \alpha_i \begin{array}{c} \triangleleft e \\ \downarrow f \\ \boxed{\tau} \\ \downarrow S(f) \\ \boxed{p_S} \\ \downarrow S(f) \\ \boxed{\sigma} \\ \downarrow f \\ \triangleleft s_i \end{array}, \quad (62)$$

and since this holds for any $\tau(e) \in \mathcal{S}(\mathcal{E}_f)$, it follows from the fact that p_S is tomographic that

$$\begin{array}{c} \boxed{\sigma} \\ \downarrow f \\ \boxed{\sum_i \alpha_i} \\ \downarrow f \\ \triangleleft s_i \end{array}^{S(f)} = \sum_i \alpha_i \begin{array}{c} \boxed{\sigma} \\ \downarrow f \\ \triangleleft s_i \end{array}^{S(f)}. \quad (63)$$

This must be true for any $s \in \Omega_f$ and any decomposition of it in terms of some s_i , so σ is a linear map.

The proof of the linearity of τ follows a similar reasoning, relying now on the fact that $\mathcal{S}(\Omega_f)$ spans the dual of $\mathcal{S}(\mathcal{E}_f)$. \square

Lemma 27. Consider a fragment $f = (\Omega_f, \mathcal{E}_f, p_f)$. For any shadow $\mathcal{S}(f) = (\sigma(\Omega_f), \tau(\mathcal{E}_f), p_S)$ of the fragment, the maps σ and τ obey

$$\text{Ker}[\sigma] = \text{Ker}_L[p_f] \quad (64)$$

$$\text{Ker}[\tau] = \text{Ker}_R[p_f]. \quad (65)$$

Proof. The proof is as follows, where the first equivalence is from the definition of the kernel, the second is from the fact that the effects of $\mathcal{S}(f)$ are (by definition) all of the form $\tau(e)$ and these (by definition) span the right argument of p_S , the third is from the fact that the shadow map preserves the probabilistic predictions of the fragment, the fourth is from the fact that the effects e (by definition) span the left argument of p_f , and the final equivalence is the definition of the left kernel.

$$v \in \text{Ker}(\sigma) \iff \begin{array}{c} \boxed{\sigma} \\ \downarrow f \\ \triangleleft v \end{array}^{S(f)} = 0 \quad (66)$$

$$\iff \begin{array}{c} \triangleleft e \\ \downarrow f \\ \boxed{\tau} \\ \downarrow S(f) \\ \boxed{p_S} \\ \downarrow S(f) \\ \boxed{\sigma} \\ \downarrow f \\ \triangleleft v \end{array} = 0, \quad \forall e \in \mathcal{E}_f \quad (67)$$

$$\iff \begin{array}{c} \triangleleft e \\ \downarrow f \\ \boxed{p_f} \\ \downarrow f \\ \triangleleft v \end{array} = 0, \quad \forall e \in \mathcal{E}_f \quad (68)$$

$$\iff \begin{array}{c} \downarrow f \\ \boxed{p_f} \\ \downarrow f \\ \triangleleft v \end{array} = 0 \quad (69)$$

$$\iff v \in \text{Ker}_L[p_f]. \quad (70)$$

A similar argument shows that $\text{Ker}[\tau] = \text{Ker}_R[p_f]$. \square

In particular, if we consider some $v \in \text{Ker}[\sigma]$ such that $v = s - s'$ then this means that $\sigma(s) = \sigma(s') \iff s \sim s'$ and similarly, for $w \in \text{Ker}[\tau]$ such that $w = e - e'$, one has $\tau(e) = \tau(e') \iff e \sim e'$ where the relations \sim are the ones in Def. 17.

Lemma 28 (Shadow maps split). Consider a fragment $f = (\Omega_f, \mathcal{E}_f, p_f)$. Given any shadow $\mathcal{S}(f) = (\sigma(\Omega_f), \tau(\mathcal{E}_f), p_S)$ for f , the maps σ and τ split as

$$\sigma = \tilde{\sigma} \circ \smile, \quad (71)$$

$$\tau = \tilde{\tau} \circ \smile, \quad (72)$$

where \smile are the quotienting maps defining f/\smile (as in Def. 17) and where $\tilde{\sigma} : \Omega/\smile \rightarrow U$ and $\tilde{\tau} : \mathcal{E}/\smile \rightarrow V$ are unique injective linear maps. Diagrammatically,

$$\begin{array}{c} \boxed{\sigma} \\ \downarrow f \end{array}^{S(f)} = \begin{array}{c} \boxed{\tilde{\sigma}} \\ \downarrow f/\smile \end{array}^{S(f)}, \quad (73)$$

$$\begin{array}{c} \boxed{\tau} \\ \downarrow S(f) \end{array}^f = \begin{array}{c} \boxed{\tilde{\tau}} \\ \downarrow S(f)/\smile \end{array}^f. \quad (74)$$

Proof. This is a consequence of the universal property of quotients. Consider a linear map $\mu : T \rightarrow W$ where T and W are vector spaces, and consider a subspace $A \subseteq T$ such that $A \subseteq \text{Ker}[\mu]$. The universal property of the quotient (for vector spaces) ensures that there exists a unique

$\tilde{\mu} : T/A \rightarrow W$ such that $\mu = \tilde{\mu} \circ \sim_A$, where \sim_A is the quotienting map $\sim : T \rightarrow T/A$. Note that $\text{Ker}[\tilde{\mu}] = \text{Ker}[\mu]/A$. Therefore, if $A = \text{Ker}[\mu]$, then $\text{Ker}[\tilde{\mu}] = 0_{T/A}$, so $\tilde{\mu}$ is injective. In other words, the universal property tells us that there exists a unique map $\tilde{\mu}$ such that the following diagram commutes:

$$\begin{array}{ccc} T & \xrightarrow{\mu} & W \\ & \searrow \sim_A & \uparrow \tilde{\mu} \\ & & T/A \end{array} \quad (75)$$

Let us apply this theorem to our case, with the substitution $T = \text{Span}[\Omega_f]$, $W = U$, and $A = \text{Ker}[\sigma]$ for splitting σ . Since the quotienting map is the canonical projection onto $\text{Span}[\Omega_f]/\text{Ker}[\sigma]$, in our case this equals $\text{Span}[\Omega_f]/\text{Ker}_L[p_f]$, given Lemma 27. This is exactly the quotienting relation for states in Def. 17 (and therefore denoted by the backwards symbol \sim). Next, we apply the universal property with the substitution $T = \text{Span}[\mathcal{E}_f]$, $W = V$, and $A = \text{Ker}[\tau]$ in order to split τ . Here, we get that \sim is the quotienting of \mathcal{E}_f with respect to $\text{Ker}_R[p_f]$ (again using Lemma 27), therefore also denoted by the backwards symbol ' \sim '. In other words, the quotienting identifies any pair of states s, s' such that $p_f(s, e) = p_f(s', e)$ for all effects $e \in \mathcal{E}_f$ and similarly for effects.

The universal property of quotients then implies that there exists a unique map $\tilde{\sigma} : \text{Span}[\Omega_f/\sim] \rightarrow U$ such that $\sigma = \tilde{\sigma} \circ \sim$ and, similarly, there exists a unique map $\text{Span}[\mathcal{E}_f/\sim] \rightarrow V$ such that $\tau = \tilde{\tau} \circ \sim$. Since the quotienting maps are implemented with respect to the kernels of σ and τ , both maps $\tilde{\sigma}$ and $\tilde{\tau}$ are injective. \square

Lemmas 26 and 28 can be used to prove that any shadow with the form prescribed by Definition 17' is equivalent to the particular (canonical) shadow f/\sim constructed in Definition 17.

Lemma 29 (All shadows $\mathcal{S}(f)$ are equivalent to f/\sim). *Consider a fragment $f = (\Omega_f, \mathcal{E}_f, p_f)$ and an arbitrary shadow $\mathcal{S}(f) = (\sigma(\Omega_f), \tau(\mathcal{E}_f), p_{\mathcal{S}})$ of that fragment (as defined in Definition 17'). Then, the quotient shadow f/\sim (as defined in Definition 17) is equivalent to $\mathcal{S}(f)$. That is, there exists an invertible GPT fragment embedding of f/\sim into $\mathcal{S}(f)$.*

Proof. We will show that $(\tilde{\sigma}, \tilde{\tau})$ provides a faithful and full embedding of f/\sim into $\mathcal{S}(f)$. Recall that this implies that the embedding is invertible. First,

note that $\tilde{\sigma}$ inherits linearity from σ (which must be linear, by Lemma 26). Second, given Lemma 28, the domain of $\tilde{\sigma}$ is $\text{Span}[\Omega_f/\sim]$, and it is an injective state map, obeying $\tilde{\sigma}(\Omega_f/\sim) \subseteq \sigma(\Omega_f)$ and mapping any distinct pair of states $\mathfrak{s}, \mathfrak{s}' \in \Omega_f/\sim$ to distinct states in $\sigma(\Omega_f)$. We now show $\tilde{\sigma}$ is surjective (full). Any state in $\sigma(\Omega_f)$ can be written as $\sigma(s)$ with $s \in \Omega_f$, and using the splitting of σ , we get

$$\begin{array}{c} \text{S}(f) \\ \downarrow \sigma \\ \nabla \end{array} = \begin{array}{c} \text{S}(f) \\ \downarrow \tilde{\sigma} \\ \nabla \end{array} \circ \begin{array}{c} f \\ \downarrow \sim \\ \nabla \end{array} = \begin{array}{c} \text{S}(f) \\ \downarrow \tilde{\sigma} \\ \nabla \end{array} \quad \forall \sigma(s) \in \sigma(\Omega_f), \quad (76)$$

which implies that $\sigma(\Omega_f) \subseteq \tilde{\sigma}(\Omega_f/\sim)$. Therefore, $\tilde{\sigma}(\Omega_f/\sim) = \sigma(\Omega_f)$, and $\tilde{\sigma}$ is a full state map.

Analogously for effects, $\tilde{\tau}$ is linear (again as a consequence of Lemma 26), its domain is \mathcal{E}_f/\sim , and it acts as an injective effect map, as a consequence of Lemma 28. Finally, it is also a full effect map, since for any $\tau(e) \in \tau(\mathcal{E}_f)$ one has

$$\begin{array}{c} \triangle \\ \uparrow \tau \\ \square \end{array} \circ \begin{array}{c} e \\ \downarrow f \\ \text{S}(f) \end{array} = \begin{array}{c} \triangle \\ \uparrow \tilde{\tau} \\ \square \end{array} \circ \begin{array}{c} f \\ \downarrow \sim \\ \text{S}(f) \end{array} = \begin{array}{c} \triangle \\ \uparrow \tilde{\tau} \\ \square \end{array} \quad (77)$$

That $(\tilde{\sigma}, \tilde{\tau})$ preserve the probabilities is a consequence of the fact that both $\mathcal{S}(f)$ and f/\sim reproduce the probabilities of the fragment f . \square

Having established that any given shadow $\mathcal{S}(f)$ of a fragment f is equivalent to the quotient shadow f/\sim , we can now prove Theorem 17.

Theorem 17. *Any two shadows of a given GPT fragment are equivalent to each other (no matter how they are constructed). Moreover, any GPT system equivalent to a shadow is itself a shadow.*

Proof. If $\mathcal{S}_1(f)$ and $\mathcal{S}_2(f)$ are two shadows for f , Lemma 29 tell us that $\mathcal{S}_1(f)$ is equivalent to f/\sim , as well as $\mathcal{S}_2(f)$ is equivalent to f/\sim . Since equivalence is transitive and symmetric, we get that $\mathcal{S}_2(f)$ is equivalent to $\mathcal{S}_1(f)$. \square

The following Lemma applies to any fragment, and will be useful below.

Lemma 30. *(Probability rules are unique) Consider a fragment $f := (\Omega_f, \mathcal{E}_f, p_f)$, and suppose that $(\Omega_f, \mathcal{E}_f, p'_f)$ is such that p'_f also reproduces the probabilities of the fragment. Then, $p'_f = p_f$.*

Proof. Suppose $p'_f(s, e) = p_f(s, e)$ for all $s \in \Omega_f$ and $e \in \mathcal{E}_f$. Since probability rules are bilinear, then $p'_f(u, v) = p_f(u, v)$ for all $u \in \text{Span}[\Omega_f]$ and $v \in \text{Span}[\mathcal{E}_f]$. Since probability rules are defined with domains $\text{Span}[\Omega_f] \times \text{Span}[\mathcal{E}_f]$ this implies $p'_f = p_f$. \square

As an instance of Lemma 30, we get the following.

Corollary 2 (The only freedom in a shadow map is given by σ and τ). *Consider a fragment $f = (\Omega_f, \mathcal{E}_f, p_f)$ and two shadows for this fragment, $\mathcal{S}_1(f) = (\sigma_1(\Omega_f), \tau_1(\mathcal{E}_f), p_{s_1})$ and $\mathcal{S}_2(f) = (\sigma_2(\Omega_f), \tau_2(\mathcal{E}_f), p_{s_2})$. If $\sigma_1 = \sigma_2$ and $\tau_1 = \tau_2$, then $p_{s_1} = p_{s_2}$. That is, given σ and τ , the probability rule $p_{\mathcal{S}}$ of the shadow $\mathcal{S}(f)$ is fixed.*

C Proof of Lemma 18

We now prove Lemma 18.

Lemma 18. *The shadow map for a given fragment is faithful if and only if the fragment is relatively tomographic.*

Proof. To see that the shadow embedding (σ, τ) is injective if and only if the fragment is relatively tomographic, notice the following: if the fragment is relatively tomographic, then for $s_1, s_2 \in \Omega_f$ one has

$$\begin{aligned}
 \triangleleft_f^{s_1} &\neq \triangleleft_f^{s_2} \iff \begin{array}{c} \triangleleft_e^f \\ \downarrow f \\ (p_f) \\ \downarrow f \\ \triangleleft_f^{s_1} \end{array} \neq \begin{array}{c} \triangleleft_e^f \\ \downarrow f \\ (p_f) \\ \downarrow f \\ \triangleleft_f^{s_2} \end{array}, \quad \exists e \in \mathcal{E}_f \quad (78) \\
 &\iff \begin{array}{c} \triangleleft_e^f \\ \downarrow f \\ \tau \\ \downarrow \mathcal{S}(f) \\ (p_{\mathcal{S}}) \\ \downarrow \mathcal{S}(f) \\ \sigma \\ \downarrow f \\ \triangleleft_f^{s_1} \end{array} \neq \begin{array}{c} \triangleleft_e^f \\ \downarrow f \\ \tau \\ \downarrow \mathcal{S}(f) \\ (p_{\mathcal{S}}) \\ \downarrow \mathcal{S}(f) \\ \sigma \\ \downarrow f \\ \triangleleft_f^{s_2} \end{array}, \quad \exists e \in \mathcal{E}_f \quad (79) \\
 &\iff \begin{array}{c} \mathcal{S}(f) \\ \downarrow f \\ \sigma \\ \downarrow f \\ \triangleleft_f^{s_1} \end{array} \neq \begin{array}{c} \mathcal{S}(f) \\ \downarrow f \\ \sigma \\ \downarrow f \\ \triangleleft_f^{s_2} \end{array}. \quad (80)
 \end{aligned}$$

Here, the first equivalence follows from the fragment being relatively tomographic, the second from Eq. (30), and the third from the fact that the probability rule of a shadow is tomographic. Consequently, the map σ is injective. A similar

argument gives $e_1 \neq e_2 \implies \tau(e_1) \neq \tau(e_2)$ for any $e_1, e_2 \in \mathcal{E}_f$ of a tomographic fragment, so τ is injective. Since σ and τ are injective, the shadow map is injective.

Conversely, notice that for any $s_1, s_2 \in \Omega_f$ such that $\sigma(s_1) = \sigma(s_2)$, injectivity of σ gives us that

$$\begin{aligned}
 \triangleleft_f^{s_1} &= \triangleleft_f^{s_2} \iff \begin{array}{c} \mathcal{S}(f) \\ \downarrow f \\ \sigma \\ \downarrow f \\ \triangleleft_f^{s_1} \end{array} = \begin{array}{c} \mathcal{S}(f) \\ \downarrow f \\ \sigma \\ \downarrow f \\ \triangleleft_f^{s_2} \end{array} \quad (81) \\
 &\iff \begin{array}{c} \triangleleft_e^f \\ \downarrow f \\ \tau \\ \downarrow \mathcal{S}(f) \\ (p_{\mathcal{S}}) \\ \downarrow \mathcal{S}(f) \\ \sigma \\ \downarrow f \\ \triangleleft_f^{s_1} \end{array} = \begin{array}{c} \triangleleft_e^f \\ \downarrow f \\ \tau \\ \downarrow \mathcal{S}(f) \\ (p_{\mathcal{S}}) \\ \downarrow \mathcal{S}(f) \\ \sigma \\ \downarrow f \\ \triangleleft_f^{s_2} \end{array}, \quad \forall e \in \mathcal{E}_f \quad (82) \\
 &\iff \begin{array}{c} \triangleleft_e^f \\ \downarrow f \\ (p_f) \\ \downarrow f \\ \triangleleft_f^{s_1} \end{array} = \begin{array}{c} \triangleleft_e^f \\ \downarrow f \\ (p_f) \\ \downarrow f \\ \triangleleft_f^{s_2} \end{array}, \quad \forall e \in \mathcal{E}_f \quad (83)
 \end{aligned}$$

with a similar argument for effects $e_1, e_2 \in \mathcal{E}_f$ such that $\tau(e_1) = \tau(e_2)$. This is the definition of a relatively tomographic fragment. \square

D Proof of Theorem 20

We now prove Theorem 20.

Theorem 20. *Every polytopic GPT system is the shadow of some fragment of a strictly classical (i.e., simplicial) GPT system of some dimension.*

Proof. Consider an arbitrary polytopic GPT system G —that is, any GPT system in which the state space Ω_G is a polytope and the effect space \mathcal{E}_G is a polytope.²³ As G is a proper GPT system rather than a generic fragment, its states and effects are tomographic for each other. For simplicity, here we will work with the representation in which states span a vector space V , effects live in the dual space V^* , and the probability rule is simply the evaluation map. The wire labels in the diagrams below describe these vector spaces instead of the GPT systems and fragments, since we consider maps that a priori might not be state/effect maps.

²³In fact, the proof does not rely on the effect space being a polytope.

Because Ω_G is a polytope, it has a finite set of pure (extremal) states, $\{s_\lambda\}_{\lambda=1}^n \in \Omega_G$, which we order in some arbitrary way as s_1, \dots, s_n . Denote the pure states of a classical (simplicial) system as $\{\lambda\}_{\lambda=1}^n$, a unit basis of \mathbb{R}^n , and order them as well. Then, consider the linear map $L: \mathbb{R}^n \rightarrow V$ defined by the action

$$\begin{array}{c} V \\ \downarrow L \\ \mathbb{R}^n \\ \downarrow \lambda \end{array} := \begin{array}{c} V \\ \downarrow s_\lambda \end{array}, \quad \forall s_\lambda \in \Omega_G. \quad (84)$$

Since the classical states are linearly independent—as they form a basis for \mathbb{R}^n —this linear map always exists. This transformation L takes the classical pure states to the pure states in Ω_G , and so can be thought of as a classically controlled preparation procedure with a setting variable that ranges over all of the pure states in the GPT system G . (Thus, it can be thought of as a physical transformation.)

One can represent L in matrix form in an intuitive way using the fact that $\{\lambda\}_\lambda$ is a unit basis of \mathbb{R}^n . Then, if s_λ is the GPT vector representing state $s_\lambda \in \Omega_G$ in some basis, one has

$$L = \begin{bmatrix} | & | & \dots & | \\ s_1 & s_2 & & s_n \\ | & | & & | \end{bmatrix}. \quad (85)$$

Consider now the contravariant²⁴ action of L on the effects in the original GPT system:

$$\begin{array}{c} \epsilon \\ \downarrow V \end{array} \mapsto \begin{array}{c} \xi_e \\ \downarrow \mathbb{R}^n \end{array} := \begin{array}{c} \epsilon \\ \downarrow V \\ \downarrow L \\ \mathbb{R}^n \end{array}. \quad (86)$$

As L is a physical transformation, it is both a state map and an effect map, and so the ξ_e are necessarily valid effects on a classical system.

We can use this map to define a fragment of the simplicial GPT system Λ of dimension n , which we call the Holevo fragment for G and which we denote by \mathcal{H}_G . It contains all the states in the simplicial theory's state space, namely $\omega \in \Delta$, but only contains the effects of the form $\xi_e = e \circ L$ for

²⁴Diagrammatically, this just means one reads the diagram top to bottom instead of bottom to top.

$e \in \mathcal{E}_G$. Diagrammatically²⁵:

$$\mathcal{H}_G := \left(\Delta, \left\{ \begin{array}{c} \xi_e \\ \downarrow \mathbb{R}^n \end{array} \right\}_{e \in \mathcal{E}_G}, \begin{array}{c} \boxed{\mathbb{R}^n} \\ \downarrow \mathbb{R}^n \end{array} \right). \quad (87)$$

Next we will construct the GPT shadow $\mathcal{S}(\mathcal{H}_G)$ of \mathcal{H}_G and show that it is equal to G .

For that, consider a right inverse of L , namely $L_r^{-1}: V \rightarrow \mathbb{R}^n$ satisfying

$$\begin{array}{c} V \\ \downarrow L \\ \mathbb{R}^n \\ \downarrow L_r^{-1} \\ V \end{array} = \begin{array}{c} V \end{array}. \quad (88)$$

This map always exists, due to the fact that L is surjective (i.e., any element $v \in V$ can be written as $L(w)$ for some $w \in \mathbb{R}^n$) and that all surjective functions have right inverses. To see that L is surjective, note that since G is a proper GPT system, $\text{Span}(\Omega_G) = V$, so any vector $v \in V$ can be written as

$$v = \sum_\lambda r_\lambda s_\lambda = \sum_\lambda r_\lambda L(\lambda) = L\left(\sum_\lambda r_\lambda \lambda\right) = L(w), \quad (89)$$

where the last equality follows from the fact that the deterministic distributions $\{\lambda\}_\lambda$ live in \mathbb{R}^n .

We now claim that we can define a shadow $\mathcal{S}(\mathcal{H}_G)$ of \mathcal{H}_G by taking $\tau := L_r^{-1}$, taking $\sigma := L$, and taking p_S to be the evaluation map in V . With these definitions, the shadow is

$$\mathcal{H} := \left(\left\{ \begin{array}{c} V \\ \downarrow L \\ \omega \end{array} \right\}_{\omega \in \Delta}, \left\{ \begin{array}{c} \xi_e \\ \downarrow \mathbb{R}^n \\ \downarrow L_r^{-1} \\ V \end{array} \right\}_{e \in \mathcal{E}_G}, \begin{array}{c} \boxed{V} \\ \downarrow V \end{array} \right). \quad (90)$$

It is not hard to see that $\mathcal{H} = G$. That the effect space of Eq. (90) is just \mathcal{E}_G follows from the definition of ξ_e (Eq. (86)). That the state space of Eq. (90) is just Ω_G follows from the fact that the extremal elements of the former are of the form $L \circ \lambda$, and hence are extremal elements of Ω_G (by Eq. (84)).

It follows that \mathcal{H} is a valid GPT system. Consequently, to prove that it is indeed a shadow of \mathcal{H}_G , all that we need to demonstrate is that L

²⁵The dashed box here is just the diagrammatic representation of the evaluation map (the probability rule).

and L_r^{-1} are full state and effect maps, respectively. Clearly, L is a full state map, since by ranging over the extremal states in \mathbb{R}^n , one can reach every extremal state in Ω_G . Moreover, L_r^{-1} is a full effect map, since varying over ξ_e is the same as varying $e \circ L$ over e , and L_r^{-1} takes this to e .

Thus we see that \mathcal{H} is indeed the shadow of \mathcal{H}_G , and that it is indeed equal to G . That is, $\mathcal{H} = \mathcal{S}(\mathcal{H}_G) = G$.

Since this construction is independent of any particularities of the polytopic GPT system G , this establishes Theorem 20. \square

In the rest of the section, we essentially provide an alternative proof to the above, one which makes use of the quotienting definition of a shadow, and so is more explicit and may be more insightful to some readers.

First we prove that L_r^{-1} identifies effects that are not distinguished by the states in $\Omega_{\mathcal{H}}$, and that L identifies states that are not distinguished by the effects in $\mathcal{E}_{\mathcal{H}}$. For the case of L_r^{-1} , it should identify effects $\xi \in \mathcal{E}_{\mathcal{H}}$ that are not distinguished by the states in the fragment; since the fragment contains all states in the simplex Δ , L_r^{-1} should not identify any two different effects. Indeed, consider two Holevo effects ξ_e and $\xi_{e'}$ in $\mathcal{E}^{\mathcal{H}}$. Then:

$$\begin{aligned} \begin{array}{c} \triangle \\ \xi_e \\ \mathbb{R}^n \\ \hline L_r^{-1} \\ \mathbb{V} \end{array} &= \begin{array}{c} \triangle \\ \xi_{e'} \\ \mathbb{R}^n \\ \hline L_r^{-1} \\ \mathbb{V} \end{array} \iff \begin{array}{c} \triangle \\ e \\ \mathbb{R}^n \\ \hline L \\ \mathbb{V} \end{array} = \begin{array}{c} \triangle \\ e' \\ \mathbb{R}^n \\ \hline L \\ \mathbb{V} \end{array} \quad (91) \\ &\iff \begin{array}{c} \triangle \\ e \\ \mathbb{V} \end{array} = \begin{array}{c} \triangle \\ e' \\ \mathbb{V} \end{array}. \quad (92) \end{aligned}$$

That is, two effects in the Holevo fragment are mapped to the same effect in the purported shadow (i.e., mapped by L_r^{-1} to the same effect) if and only if they came from the same effect $e = e'$ in the original theory. But for $e = e'$ one has $\xi_e = e \circ L = e' \circ L = \xi_{e'}$, and so

$$\begin{array}{c} \triangle \\ \xi_e \\ \mathbb{R}^n \\ \hline L_r^{-1} \\ \mathbb{V} \end{array} = \begin{array}{c} \triangle \\ \xi_{e'} \\ \mathbb{R}^n \\ \hline L_r^{-1} \\ \mathbb{V} \end{array} \iff \begin{array}{c} \triangle \\ \xi_e \\ \mathbb{R}^n \end{array} = \begin{array}{c} \triangle \\ \xi_{e'} \\ \mathbb{R}^n \end{array} \quad (93)$$

so no distinct effects in $\mathcal{H}(G)$ are mapped to the same effect in the shadow \mathcal{H} by the action of L_r^{-1} . Now, for the case of states, L must identify those

states in Δ that are not distinguished by the effects in $\mathcal{H}(G)$. Let us show this is indeed the case. First, we have that for all $\xi \in \mathcal{E}_{\mathcal{H}}$,

$$\begin{array}{c} \triangle \\ \xi_e \\ \mathbb{R}^n \\ \hline \omega \end{array} = \begin{array}{c} \triangle \\ \xi_e \\ \mathbb{R}^n \\ \hline \gamma \end{array} \implies \begin{array}{c} \triangle \\ e \\ \mathbb{R}^n \\ \hline L \\ \mathbb{V} \end{array} = \begin{array}{c} \triangle \\ e \\ \mathbb{R}^n \\ \hline L \\ \mathbb{V} \end{array}, \quad \forall e \in \mathcal{E}_G. \quad (94)$$

Above, the condition $\forall e \in \mathcal{E}_G$ follows from the fact that $\{e \circ L\}_{e \in \mathcal{E}_G} = \mathcal{E}_{\mathcal{H}}$, given the definition of the Holevo fragment $\mathcal{H}(G)$. Now, since G is tomographically complete (as it is a proper GPT system), we get

$$\begin{array}{c} \triangle \\ e \\ \mathbb{R}^n \\ \hline L \\ \mathbb{V} \end{array} = \begin{array}{c} \triangle \\ e \\ \mathbb{R}^n \\ \hline L \\ \mathbb{V} \end{array}, \quad \forall e \in \mathcal{E}_G \implies \begin{array}{c} \triangle \\ e \\ \mathbb{R}^n \\ \hline \omega \end{array} = \begin{array}{c} \triangle \\ e \\ \mathbb{R}^n \\ \hline \gamma \end{array} \quad (95)$$

showing that all states that are not distinguished by the effects in $\mathcal{H}(G)$ get identified by L . This implies that $\text{Ker}_L[p_{\mathcal{H}}] \subseteq \text{Ker}[L]$, where $p_{\mathcal{H}}$ is the probability rule of the Holevo fragment (the evaluation map). In other words, all states that are not distinguished by the effects in \mathcal{H} get mapped to the same state by the map L .

Finally, note that the probabilities predicted by \mathcal{H} are the same as those predicted by $\mathcal{H}(G)$, since

$$\begin{array}{c} \triangle \\ \xi_e \\ \mathbb{R}^n \\ \hline L_r^{-1} \\ \mathbb{V} \end{array} = \begin{array}{c} \triangle \\ e \\ \mathbb{R}^n \\ \hline L \\ \mathbb{V} \end{array} \quad (96)$$

$$= \begin{array}{c} \triangle \\ e \\ \mathbb{R}^n \\ \hline L \\ \mathbb{V} \end{array} \quad (97)$$

$$= \begin{array}{c} \triangle \\ \xi_e \\ \mathbb{R}^n \\ \hline \omega \end{array}. \quad (98)$$

In the first and last equality, we used that $\xi_e = e \circ L$ and in the second equality we used that $L_r^{-1} \circ L = \mathbb{1}_V$.

This also implies that any two states that are distinguished by the effects in \mathcal{H} should be distinct in $\Omega_{\mathcal{H}}$. This means that $\text{Ker}[L] \subseteq \text{Ker}_L[p_{\mathcal{H}}]$ which then implies $\text{Ker}[L] = \text{Ker}_L[p_{\mathcal{H}}]$. By the universal

property of quotienting, $L = \tilde{L} \circ \sim$, where \sim is the quotienting relation of Def. 17, and \tilde{L} is some injective map. What is important here is that if $L(\omega) \neq L(\gamma)$, then there exists some effect $f \in \mathcal{E}_H$ that distinguishes them. Therefore, the left kernel of p_H is trivial. Given that the right kernel is also trivial, the evaluation map p_H is tomographic. Therefore, \mathcal{H} is indeed a shadow of $\mathcal{H}(G)$.

We now prove that $\mathcal{H} = G$. We start by showing that the state space Ω_H of the shadow is equal to Ω_G and that the effect space \mathcal{E}_H of the shadow is equal to \mathcal{E}_G . First, note that all states $\nu \in \Omega_H$ can be written as $L(\omega)$ where $\omega \in \Delta$. Since a classical state $\omega \in \Delta$ is always a convex combination of the deterministic ones, $\{\lambda\}_{\lambda \in \Delta}$, we get

$$\begin{array}{c} \downarrow \nu \\ \omega \end{array} = \begin{array}{c} \downarrow L \\ \omega \end{array} = \sum_{\lambda} \alpha_{\lambda} \begin{array}{c} \downarrow L \\ \lambda \end{array} = \sum_{\lambda} \alpha_{\lambda} \begin{array}{c} \downarrow s_{\lambda} \\ \lambda \end{array} \in \Omega_G, \quad (99)$$

where $\{\alpha_{\lambda}\}$ are coefficients of the convex combination. This tells us that every $\nu \in \Omega_H$ belongs to Ω_G . Now, any $\omega \in \Omega_G$ is a convex combination of the pure states $\{s_{\lambda}\}_{\lambda}$, so

$$\begin{array}{c} \downarrow \omega \\ \omega \end{array} = \sum_{\lambda} \beta_{\lambda} \begin{array}{c} \downarrow s_{\lambda} \\ \lambda \end{array} = \sum_{\lambda} \beta_{\lambda} \begin{array}{c} \downarrow L \\ \lambda \end{array} = \begin{array}{c} \downarrow L \\ \gamma \end{array} \in \Omega_H, \quad (100)$$

where we have defined

$$\begin{array}{c} \downarrow \gamma \\ \gamma \end{array} := \sum_{\lambda} \beta_{\lambda} \begin{array}{c} \downarrow s_{\lambda} \\ \lambda \end{array}. \quad (101)$$

This shows that every $\omega \in \Omega_G$ is also an element of Ω_H . Putting both facts together gives $\Omega_H = \Omega_G$.

Now, let us turn to the effects. First, using Eq. 88, any $e \in \mathcal{E}_G$ can be written as

$$\begin{array}{c} \uparrow e \\ \downarrow \nu \end{array} = \begin{array}{c} \uparrow e \\ \downarrow L \\ \downarrow L_r^{-1} \\ \downarrow \nu \end{array} = \begin{array}{c} \uparrow \xi_e \\ \downarrow L_r^{-1} \\ \downarrow \nu \end{array} \in \mathcal{E}_H. \quad (102)$$

Conversely, for every effect $f \in \mathcal{E}_H$ in the shadow, there is (by Eq. 90) a unique effect $e \in \mathcal{E}_G$ in the original GPT system such that $f = \xi_e \circ L_r^{-1}$, so we denote the shadow effects as f_e . Given the

definition of ξ_e as $e \circ L$, this implies that

$$\begin{array}{c} \uparrow f_e \\ \downarrow \nu \end{array} = \begin{array}{c} \uparrow \xi_e \\ \downarrow L_r^{-1} \\ \downarrow \nu \end{array} = \begin{array}{c} \uparrow e \\ \downarrow L \\ \downarrow L_r^{-1} \\ \downarrow \nu \end{array} = \begin{array}{c} \uparrow e \\ \downarrow \nu \end{array} \in \mathcal{E}_G, \quad (103)$$

which proves that $f \in \mathcal{E}_H \implies f \in \mathcal{E}_G$. This proves that $\mathcal{E}_H = \mathcal{E}_G$ and that $f_u = u$.

In addition to having identical state spaces and effect spaces and unit effects, the probability rules for \mathcal{H} and G are identical, so they are in fact identical GPT systems.

E Proof of Theorem 22

We now prove Theorem 22.

Theorem 22. *Consider the data table generated by some GPT fragment. The output of theory-agnostic tomography (as in Definition 21) applied to this data table is the GPT shadow of the GPT fragment. That is, for all fragments f we have that $\mathcal{S}^{TA}(f) = G_{D_f}$ is a valid shadow of f .*

Proof. Note that in this proof, the diagram wires are labelled by the domain (codomain) of the map instead of by the GPT systems and fragments.

In theory-agnostic tomography, one factorizes the matrix of empirical data $D \in M_{m \times n}(\mathbb{R})$ generated by the states and effects in one's fragment f by finding the smallest integer k such that there exist $E \in M_{m \times k}(\mathbb{R})$ and $S \in M_{k \times n}(\mathbb{R})$ satisfying

$$\begin{array}{c} \mathbb{R}^m \\ \boxed{D} \\ \mathbb{R}^n \end{array} = \begin{array}{c} \mathbb{R}^m \\ \boxed{E} \\ \mathbb{R}^k \\ \boxed{S} \\ \mathbb{R}^n \end{array}. \quad (104)$$

Consider now the states $\Theta := \{s_1, \dots, s_n\}$ and effects $\Phi := \{e_1, \dots, e_m\}$ defining²⁶ the original (unknown) fragment f , and define a map $d: \Theta \rightarrow \mathbb{R}^n$ mapping every $s_i \in \Theta$ to one of the n orthonormal basis vectors in \mathbb{R}^n :

$$\begin{array}{c} \mathbb{R}^n \\ \boxed{d} \\ \Theta \\ \downarrow s_i \end{array} := \begin{array}{c} \mathbb{R}^n \\ \downarrow i \end{array}. \quad (105)$$

²⁶Recall from Section 6.1 that these may just be taken to be the convexly extremal states and effects in the fragment.

Note that this map has a set (Θ) , not a vector space, as its domain; consequently, one cannot speak about whether or not it is linear. The map $D \circ d : \Theta \rightarrow \mathbb{R}^m$ therefore takes states in Θ to columns of the data table D .

Due to the minimality of k , the process E must admit of a left-inverse, i.e., there is E_l^{-1} such that

$$\begin{array}{c} \mathbb{R}^k \\ \hline \text{trapezoid } E_l^{-1} \\ \hline \mathbb{R}^m \\ \hline \text{trapezoid } E \\ \hline \mathbb{R}^k \end{array} = \begin{array}{c} \mathbb{R}^k \\ \hline \end{array}. \quad (106)$$

Consequently, let us define the map

$$\begin{array}{c} \mathbb{R}^k \\ \hline \text{box } \hat{\sigma} \\ \hline \Theta \end{array} := \begin{array}{c} \mathbb{R}^k \\ \hline \text{trapezoid } E_l^{-1} \\ \hline \mathbb{R}^m \\ \hline \text{box } D \\ \hline \mathbb{R}^n \\ \hline \text{box } d \\ \hline \Theta \end{array} = \begin{array}{c} \mathbb{R}^k \\ \hline \text{trapezoid } S \\ \hline \mathbb{R}^n \\ \hline \text{box } d \\ \hline \Theta \end{array}, \quad (107)$$

which takes states $s_i \in \Omega$ to columns of the matrix S .

We now define a map σ , which (as we will show later) is the linear extension of $\hat{\sigma}$, acting on $\text{Span}[\Theta]$. For each vector $s \in \text{Span}[\Theta]$, pick one specific linear combination $s = \sum_i \alpha_i s_i$ of that vector in terms of the states s_i , and define the action of σ by

$$\begin{array}{c} \mathbb{R}^k \\ \hline \text{box } \sigma \\ \hline \sum_i \alpha_i s_i \end{array} := \sum_i \alpha_i \begin{array}{c} \mathbb{R}^k \\ \hline \text{box } \hat{\sigma} \\ \hline \Theta \\ \hline s_i \end{array}, \quad \alpha_i \in \mathbb{R}, \forall i \quad (108)$$

In principle, this map could be non-linear, as it might depend on the particular choice of decomposition for each vector s . However, below we will show that it forms part of a valid shadow map taking f to G_{Df} , so by Lemma 26, it is indeed linear.

Similarly to what was done for the states, one can define a map d' taking effects $e_i \in \Phi$ to one of the m orthonormal basis vectors in \mathbb{R}^m , such that the composition $d' \circ D$ maps the effects $e_i \in \Phi$ to rows of the matrix D :

$$\begin{array}{c} \text{triangle } e_i \\ \hline \Phi \\ \hline \text{box } d' \\ \hline \mathbb{R}^m \end{array} := \begin{array}{c} \text{triangle } i \\ \hline \mathbb{R}^m \end{array}. \quad (109)$$

And by minimality of k , the map S must admit of

a right-inverse, S_r^{-1} , such that

$$\begin{array}{c} \mathbb{R}^k \\ \hline \text{trapezoid } S \\ \hline \mathbb{R}^n \\ \hline \text{trapezoid } S_r^{-1} \\ \hline \mathbb{R}^k \end{array} = \begin{array}{c} \mathbb{R}^k \\ \hline \end{array}. \quad (110)$$

So, analogous to the above, we define the map

$$\begin{array}{c} \Phi \\ \hline \text{box } \hat{\tau} \\ \hline \mathbb{R}^k \end{array} := \begin{array}{c} \Phi \\ \hline \text{box } d' \\ \hline \mathbb{R}^m \\ \hline \text{box } D \\ \hline \mathbb{R}^n \\ \hline \text{trapezoid } S_r^{-1} \\ \hline \mathbb{R}^k \end{array} \quad (111)$$

which takes effects $e_i \in \Phi$ to rows of the matrix E . Finally, we define the linear extension of $\hat{\tau}$, namely $\tau : \text{Span}(\Phi) \rightarrow \mathbb{R}^k$, just as we did for the map σ above. (Here too, defining τ by a particular choice for the decomposition of each vector in its domain might give the appearance that the map could be nonlinear, but we will justify below that it is in fact linear.)

It now follows that $\sigma(\text{Conv}[\Theta])$ and $\tau(\text{Conv}[\Phi])$, together with the evaluation map on \mathbb{R}^k , constitutes a shadow for the fragment f that generated the data table. Indeed, the evaluation map in \mathbb{R}^k is tomographic, due to minimality of k , and the probabilities of f (encoded in D) are reproduced:

$$\begin{array}{c} \text{triangle } e_j \\ \hline \Phi \\ \hline \text{box } \tau \\ \hline \mathbb{R}^k \\ \hline \text{box } \sigma \\ \hline \Theta \\ \hline \text{triangle } s_i \end{array} = \begin{array}{c} \text{triangle } e_j \\ \hline \Phi \\ \hline \text{box } d' \\ \hline \mathbb{R}^m \\ \hline \text{box } D \\ \hline \mathbb{R}^n \\ \hline \text{box } d \\ \hline \Theta \\ \hline \text{triangle } s_i \end{array} = \begin{array}{c} \text{triangle } j \\ \hline \mathbb{R}^m \\ \hline \text{box } D \\ \hline \mathbb{R}^n \\ \hline \text{triangle } i \end{array} = D(i,j). \quad (112)$$

Since σ and τ constitute shadow maps, Lemma 26 tell us that they are linear. Therefore, the particular choice of linear combination for a vector $s \in \text{Span}[\Theta]$ in terms of $\{s_i\}$, made when defining σ , is irrelevant (and similarly for τ).

Therefore, we have established that the output of theory-agnostic tomography applied to the data table generated by a GPT fragment gives the shadow of that fragment as its output.

□