

Classical Explanations in (and of) General Probabilistic Theories

Summary for QPL 2026

John Harding and Alex Wilce

Classical probability theory contains the tacit assumption that all experiments have common refinements, and can, therefore, effectively be performed jointly. QM suggests that this is not a reasonable assumption. A simple and natural way to generalize classical probability theory so as to allow for incompatible experiments takes a *probabilistic model* to be a pair $(\mathcal{M}, \Omega)$, where $\mathcal{M}$ is a test space — a catalogue of outcome-sets of various possible experiments — and Ω is a set of states, understood as probability weights on $\mathcal{M}$. This subsumes both the quantum-logical and (as detailed in [1]), the convex-operational approach.

The literature is replete with theorems showing that various kinds of probabilistic models — quantum logics, compact convex sets, etc. — can be represented, in one way or another, in terms of a classical probabilistic model, provided we are willing to accept (i) failures of contextuality, and (ii) loss of locality. However, these results are obtained within several different formal frameworks, the representations they offer are not all equivalent, and many of them are restricted to the finite-dimensional setting.

Our purpose in this paper is to provide a unified, and somewhat more general, picture; one that we believe is clarifying. This is based on a simple but under-explored idea, which is that our notion of what it means for one probabilistic model B to “explain” another, A , should not be limited to embeddings of A into B . Quotient mappings from B onto A are also explanatory. Moreover, the relation “explains” should be transitive, so one wants to allow a diagram of the form

$$A \xleftarrow{q} C \xrightarrow{e} B$$

where A, B and C are PMs, q is a quotient, and e is an embedding, to count as an explanation of A by B .

After defining the necessary vocabulary, including morphisms (and, in particular, quotients and embeddings) of probabilistic models, we show that spans of the above form do compose, in the expected way, in **Prob**, by showing that the dual diagrams (co-explanations) have pullbacks (our Proposition 2.11). This is slightly nontrivial, since **Prob** does not support arbitrary pullbacks.

This shows that our notion of explanation is indeed transitive, and explains why we need not include co-explanations as explanations. (It also opens the way to defining a bicategory of probabilistic models and explanations, though this is something we leave for another day.)

Having established this, we then ask what is required for a probabilistic model A to have a *classical* explanation, B . The answer is: not much. We show that any locally finite model A (one the tests of which are all of finite cardinality) in which at most one test has cardinality 1, always has a canonical such explanation in terms of a classical model we term its “Borelification”, $\mathbf{Bor}(A)$. This is a re-working of a construction that first appeared in [21]. In brief, every probabilistic model A is a quotient of its *semi-classical cover*, $\tilde{A}$, and this hosts an abundance of dispersion-free states. When A is locally finite, the set $S(\tilde{A})$ of these dispersion-free states is a compact Hausdorff space, and there is a natural embedding of $\tilde{A}$ into the test space $\mathbf{Bor}(A) := \mathcal{B}(S(\tilde{A}))$ of finite partitions of $S(\tilde{A})$ into Borel sets. This can be extended to the context of probabilistic models to yield the desired explanation,

$$A \xleftarrow{q} \tilde{A} \xrightarrow{e} \mathbf{Bor}(A).$$

Moreover, $A \mapsto \mathbf{Bor}(A)$ is the object part of a functor from the category **Prob** of probabilistic models and morphisms of these, to the category of measurable spaces and mappings. It follows that every probabilistic *theory*, understood as a functor $F : \mathcal{C} \rightarrow \mathbf{Prob}$ from a category $\mathcal{C}$ of physical systems and processes into **Prob**, as in [3], has a classical explanation as well, namely $\mathbf{Bor} \circ F$.

We draw some connections with the framework of ontological representations, taking note of the distinction between sharp and unsharp representations. As shown in [1], the unsharp case can largely be reduced to the sharp case. The upshot is that all locally finite models are quotients of models having *sharp* ontological representations.

It is not to be expected that the **Bor** construction will preserve locality. We reformulate Bell-locality in terms of probabilistic models, and prove a non-locality result, but there is a subtlety: a distinction, not previously made, so far as we know, between entangled *probability weights* and entangled *states*. As we show, a bipartite state can be entangled *as a state*, but separable

as a joint probability weight. Entangled probability weights do not have local classical explanations (Proposition 4.7), while entangled states may (Example 4.2).

Our results are quite general: we make no special assumptions about the structure of probabilistic models, other than local finiteness. In particular, we do not assume that our PMs are finite-dimensional. The amount of measure-theoretic and functional-analytic apparatus needed to deal with the infinite-dimensional case turns out to be relatively modest, and is reviewed.

References

- [1] J. Harding and A. Wilce, Order unit spaces and probabilistic models, Submitted to QPL 2026
- [2] A. Wilce, Quantum logic and probability, Stanford Encyclopedia of Philosophy, 2002
- [3] A. Wilce, General probabilistic theories: notes for a short course, [arXiv:2501.00718](https://arxiv.org/abs/2501.00718), 2025

Classical Explanations in (and of) General Probabilistic Theories

John Harding and Alex Wilce

March 5, 2026

Abstract

A probabilistic model consists of a test space M and a designated set of probability weights on M , called the state space of the model. These structures form a category, PROB , in a natural way. A probabilistic theory is then a functor from a category of (say, physical) systems and processes (or mathematical proxies for these) into PROB . We introduce a notion of the explanation of one model by another as a particular kind of span in PROB . More precisely, an explanation of a model A by a model B consists of a model C , a quotient morphism from C to A , and an embedding of B into C (all of which terms we define). We show that co-explanations have pullbacks, which are explanations, allowing us to compose explanations. We then show that, subject to one very minor restriction — that M contain no more than one single-outcome test — every locally-finite model, and in particular, every quantum model, has a canonical classical explanation. The construction is functorial, so every locally-finite, point-closed probabilistic theory has a canonical sharp classical extension. We explain the connections between our framework and that of ontological representations, and also consider the familiar tension between classicality and locality in our terms. Unsurprisingly, our canonical classical explanations are highly non-local. But because we do not assume finite-dimensionality, local tomography, or that all probability weights are states, some subtleties arise. In particular, there is a distinction between entanglement of states (which can occur even classically, as we illustrate) and of probability weights (which can't).

1 Introduction

Classical probability theory contains the tacit assumption that all experiments have common refinements, and can, therefore, effectively be performed jointly. QM suggests that this is not a reasonable assumption. Attempts to generalize probability theory in order to accommodate incompatible experiments have proceeded along two main lines: the “quantum-logic” approach, in which Boolean algebras are replaced by more general orthocomplemented posets (typically, orthomodular lattices or posets), and the “convex” approach, in which the simplex of probability measures on a classical measurable space is replaced by a more-or-less arbitrary convex set (typically compact, or at least radially bounded) with states represented by so-called effects.

The literature is replete with theorems showing that models of either of these types can be represented — in one way or another — in terms of a classical probabilistic model, provided we are willing to accept (i) failures of contextuality and (ii) loss of locality. However, these results are obtained within several different formal frameworks, and the representations they offer are not all equivalent. Our purpose in this paper is to provide a unified, and somewhat more general, picture; one that we believe is ultimately clarifying.

We begin by reviewing a framework for generalized probability theory that includes both the quantum logic and convex approaches as special cases, and, within this, constructing a category **Prob** of concrete probabilistic models, following [9, 24, 25]. We then propose a precise definition of an *explanation* of one probabilistic model by another (briefly: a particular kind of span in **Prob**). Reworking a construction first detailed in [21], we show that there exists an endofunctor $\mathbf{Bor} : \mathbf{Prob} \rightarrow \mathbf{Prob}$ sending each probabilistic model A to a classical *Borel model* $\mathbf{Bor}(A)$, in terms of which A has an explanation.

We explain how this generalizes the classical embeddings and extensions of Beltrametti and Bugajski [3] as well as the framework of ontological models [15]. We also note that any *probabilistic theory* — a functor $M : \mathcal{C} \rightarrow \mathbf{Prob}$, where $\mathcal{C}$ is a category representing physical systems and processes [23, 24, 25] — has a classical representation, namely, $\mathbf{Bor} \circ M$.

Finally, we explore the ways in which this canonical “classicalization” in-

teracts with composite systems. The results are unsurprising: if AB is a nonsignalling composite of A and B , then if AB admits a classical explanation in our sense, then all of its states are separable as joint probability weights, but not necessarily as joint states, of A and B — a distinction we explain in Section 4. It follows that entangled probability weights on AB have no classical explanation, even in our quite generous sense. We also consider how Bell’s Theorem can be formulated in this framework. What novelty there is here consists in the level of generality (e.g., we do not assume local tomography or finite-dimensionality, and allow for a broader notion of classical explanation than is usually dealt with in the literature), and in the detailed development of the mathematics supporting this generality. In Section 5, we conclude with some quasi-philosophical reflections on the different attitudes one can reasonably take about the significance of this result for our understanding of the world.

Somewhat similar ideas are explored in an earlier paper by Gheorgiu and Heunen [11] in the more restricted context of ontological models of quantum theory. Specifically, Gheorgiu and Heunen consider an ontological model as a functor from **Hilb** to the category **Markov** of measurable spaces and Markov processes. However, [11] are more concerned with reproducing existing no-go results; our focus here is on positive results.

2 Probabilistic models and theories

In discrete classical probability theory, a *probabilistic model* is usually understood to consist of the outcome-set E of some experiment, together with a single probability weight thereon. There are two obvious ways in which to generalize: one is to allow for multiple experiments (and thus, multiple outcome-sets); the other is to allow for multiple probability weights. One motivation for the former is to better understand theories like QM in which distinct statistical experiments can be incompatible. The motivation for the latter is simply to allow for some uncertainty as to the true probability weight, e.g., in situations in which many different probability weights can be prepared (either deliberately or “by Nature”).

These ideas motivate the following notions (see, e.g., [6, 9, 22]): A *test space* or *manual*¹ is a collection $\mathcal{M}$ of non-empty sets, understood as outcome-sets for some catalogue of possible experiments. We refer to $X = \bigcup \mathcal{M}$ as the *outcome-space* for $\mathcal{M}$. A *probability weight* on $\mathcal{M}$ is a function $\omega : X \rightarrow [0, 1]$ summing to one on each test. We write $\mathcal{Pr}(\mathcal{M})$ for the set of all probability weights on $\mathcal{M}$. Note that $\mathcal{Pr}(\mathcal{M}) \leq [0, 1]^X$ is convex. If $\mathcal{M}$ is *locally finite*, meaning that all tests are finite, then $\mathcal{Pr}(\mathcal{M})$ is also pointwise closed, and hence, compact.

2.1 Definition. A *probabilistic model* — PM, for short — is simply a pair $(\mathcal{M}, \Omega)$ where $\mathcal{M}$ is a test space and Ω is a set of probability weights on $\mathcal{M}$, referred to as the *state space* of the model.

A very simple example consists of single test E and a single probability weight μ thereon. For another, consider $\mathcal{M} = \{\{a, b\}, \{u, v\}\}$ with $\Omega = \mathcal{Pr}(\mathcal{M})$: it is easy to see that a probability weight on $\mathcal{M}$ is determined uniquely by its values on a and u , so we have an affine isomorphism $\mathcal{Pr}(\mathcal{M}) \rightarrow [0, 1]^2$. Any subset Ω of $[0, 1]^2$ can serve as the state space for a model based on $\mathcal{M}$. We will consider further examples immediately below. At this stage, we impose no conditions on either $\mathcal{M}$ or Ω .

2.1 Classical and Quantum Models

We now outline how standard measure-theoretic probability theory and its non-commutative analogue based on von Neumann algebras both fit into the current framework.

2.2 Example (Borel models). If S is a topological space, the corresponding *Borel test space* is the collection $\mathcal{B}(S)$ of finite or partitions of S by Borel sets. Thus, outcomes of $\mathcal{B}(S)$ are non-empty Borel sets, and events are finite pairwise disjoint collections of Borel sets. Two events are perspective iff they have the same union; using this, it easily follows that probability weights on $\mathcal{B}(S)$ are essentially the same things as finitely-additive Borel probability

¹“Manual” is the older term; it motivates our notation, and survives in the names of some of our examples.

measures on S ². Accordingly, by a *Borel model*, we mean any model of the form $(\mathcal{B}(S), \Omega)$ where Ω is any family of Borel probability measures on S .

A natural modification of this example is to consider the set $\mathcal{B}_\sigma(S)$ of countable partitions of S by Borel sets. Probability weights on $\mathcal{B}_\sigma(S)$ are then countably-additive Borel probability measures on S . One can also attach probabilistic models to arbitrary measurable spaces, or indeed, arbitrary Boolean algebras, in the same way.

2.3 Example (Projective quantum models). If $\mathbf{H}$ is a Hilbert space, let $\mathbb{P}(\mathbf{H})$ be its lattice of projection operators. A *projective test* on $\mathbf{H}$ is simply a finite set of non-zero pairwise orthogonal projections with sum (equivalently, join) equal to 1. We write $\mathcal{M}(\mathbf{H})$ for the collection of such projective tests. Thus, outcomes of $\mathcal{M}(\mathbf{H})$ are non-zero projections, events are finite or countable pairwise disjoint families of such projections, and two events are perspective iff they have the same sum. It follows easily that probability weights on $\mathcal{M}(\mathbf{H})$ are essentially the same things as finitely additive probability measures on $\mathbb{P}(\mathbf{H})$. It is a corollary of Gleason’s Theorem that such a measure extends uniquely to a state on the algebra $\mathcal{B}(\mathbf{H})$ of bounded linear operators on $\mathbf{H}$. In particular, every density operator W on $\mathbf{H}$ yields a probability weight on $\mathcal{M}(\mathbf{H})$, namely $\alpha_W(p) = \text{Tr}(Wp)$. By a *Projective Hilbert model*, we mean a PM of the form $(\mathcal{M}(\mathbf{H}), \Omega(\mathbf{H}))$ where $\Omega(\mathbf{H})$ is some set of probability weights arising from density operators on $\mathbf{H}$ in this way.

As with Borel manuals, we can enlarge $\mathcal{M}(\mathbf{H})$ to the set $\mathcal{M}_\sigma(\mathbf{H})$ of countable partitions of unity in $\mathbb{P}(\mathbf{H})$. More generally, if $\mathfrak{A}$ is a von Neumann algebra, then replacing $\mathbb{P}(\mathbf{H})$ with $\mathbb{P}(\mathfrak{A})$, we can define test spaces $\mathcal{M}(\mathfrak{A})$ and $\mathcal{M}_\sigma(\mathfrak{A})$. The generalized Gleason Theorem [4] again identifies probability weights on the former with states, and those on the latter with σ -additive states, on $\mathfrak{A}$.

Unsharp probabilistic models While our focus in this paper is on probabilistic models in which outcomes are “sharp”, meaning that we can definitively say whether or not a given outcome has occurred, there is a substantial literature on “unsharp” measurements. A few remarks are in order here about how these fit into the present scheme. Let S be a measurable space,

²Indeed, a finitely-additive Borel probability measure *is* simply a probability weight on $\mathcal{B}(S)$, trivially extended to the entire Borel field by setting $\alpha(\emptyset) = 0$.

and let $B(S, \Sigma)$ denote the space of bounded measurable real-valued functions on S . A *response function* on S is simply a positive measurable function $f \in B(S, \Sigma)$ with $0 \leq f(s) \leq 1$ for every $s \in S$. A finite *partition of unity* in $B(S, \Sigma)$ is a list $(f_1, \dots, f_n)$ of response functions with $\sum_{i=1}^n f_i = 1$. If μ is a probability measure on (S, Σ) , then $i \mapsto \int f_i d\mu$ defines a probability weight on the indexing set $\{1, \dots, n\}$, which one understands to be the probability of obtaining the i -th outcome in an “unsharp” experiment with outcomes labelled by $\{1, \dots, n\}$. Notice, however, that we can have $f_i = f_j$ for distinct i, j .

2.4 Definition (Unsharp Borel manuals). *The unsharp Borel manual (or test space) associated with (S, Σ) is the collection $\mathcal{M}^o(S, \Sigma)$ of all graphs of ordered partitions of unity in $B(S, \Sigma)$.*

Thus, outcomes in $\mathcal{M}^o(S, \Sigma)$ are pairs (i, f) where $i \in \mathbb{N}$ and f is a response function. If μ is a probability measure on S , then $\mu^o(i, f) = \int_S f d\mu$ defines a probability weight on $\mathcal{M}^o$. We obtain in this way a probabilistic model $M^o(S, \Sigma) = (\mathcal{M}^o(S, \Sigma), \Delta^o)$, where $\Delta^o = \{\mu^o | \mu \in \Delta\}$. We call this the *unsharp Borel model* associated with (S, Σ) .

A similar approach allows us to construct “unsharp” quantum models, and, more generally, unsharp probabilistic models associated with arbitrary order unit spaces. We can also apply a similar construction above to ordinary probabilistic models, as follows. Suppose $\mathcal{M}$ is any test space. We define the corresponding *ordered test space* $\mathcal{M}^o$ to consist of the graphs of all bijections $f : \{1, \dots, |E|\} \rightarrow E$ for every $E \in \mathcal{M}$. Outcomes of $\mathcal{M}^o$ are ordered pairs (i, x) where $x \in X$ and $i \leq |E|$ for some test E containing x . Any probability weight α on $\mathcal{M}$ extends to one on $\mathcal{M}^o$ by setting $\alpha^o(i, x) = \alpha(x)$. Thus, given a PM $A = (\mathcal{M}, \Omega)$, we obtain a corresponding *ordered model* $A^o = (\mathcal{M}^o, \Omega^o)$ where Ω^o consists of weights α^o with $\alpha \in \Omega$. This construction will be useful later on. For further details, see [13].

2.2 Categories of Probabilistic Models

If we wish to develop a general probability theory on basis of probabilistic models, we need to identify suitable maps between them. These should

behave reasonably both with respect to the models' test-space structure, and with respect to their states.

Test space morphisms Given the generality and comparatively loose structure of test spaces, it should not be surprising that there is more than one plausible way to define a “morphism” of such models [7, 22]. The following is adequate for our purposes³

2.5 Definition. Let $\mathcal{M}$ and $\mathcal{M}'$ be test spaces with outcome-sets X and X' , respectively. A *test space morphism* $\phi : \mathcal{M} \rightarrow \mathcal{M}'$ is an event-valued mapping $\phi : X \rightarrow \mathcal{E}(\mathcal{M}')$ such that

- (i) $a \in \mathcal{E}(\mathcal{M}) \Rightarrow \phi(a) := \bigcup_{x \in a} \phi(x) \in \mathcal{E}(\mathcal{M}')$;
- (ii) $x \perp y$ in X implies $\phi(x) \perp \phi(y)$;
- (iii) if a, b are events of $\mathcal{M}$ with $a \sim b$, then $\phi(a) \sim \phi(b)$. Equivalently, $\phi(E) \sim \phi(F)$ for all tests $E, F \in \mathcal{M}$.

In the case in which $\phi(x) = \{y\}$ for an outcome $y \in Y$, we will ordinarily write $\phi(x) = y$.

We say that a morphism is *test-preserving* iff $\phi(E)$ is a test in $\mathcal{M}'$ for every test $E \in \mathcal{M}$, *positive* iff $\phi(x) \neq \emptyset$ for every $x \in X$, and *outcome-preserving* when $\phi(x)$ is a singleton for every $x \in X$. When ϕ is both positive and outcome-preserving, we naturally regard it as a mapping $X \rightarrow X'$. An *interpretation* is a positive, test-preserving morphisms.

For a simple example, if $\mathcal{M} \subseteq \mathcal{M}'$, then inclusion mapping from $X = \bigcup \mathcal{M}$ to $X' = \bigcup \mathcal{M}'$ is an interpretation. For another, consider a subset $S \subseteq X = \bigcup \mathcal{M}$. We say that S is a *support* iff the set $\mathcal{M}_S := \{E \cap S \mid E \in \mathcal{M}\}$ is irredundant, and thus a test space in its own right. In this case, the mapping $\phi : X \rightarrow \mathcal{E}(\mathcal{M})$ given by $\phi(x) = \{x\}$ if $x \in S$ and $\phi(x) = \emptyset$ if $x \notin S$ is a test-preserving but non-positive morphism. On the other hand, the inclusion mapping $x \mapsto \{x\}$ from S to X is generally not a test space morphism at all, as simple examples show that condition (iii) can fail.

³What we are calling *morphisms* of test spaces, Foulis and Randall called ..., using the term morphism for a much broader class of mappings.

Turning to more specific cases, we can characterize morphisms between Borel and Quantum models as follows.

2.6 Example (Morphisms of Borel manuals). Let (S, Σ) and (T, Ξ) be measurable spaces, and let $\mathcal{M}(S, \Sigma)$ and $\mathcal{M}(T, \Xi)$ be the corresponding manuals of countable measurable partitions of S and T , respectively. If $f : S \rightarrow T$ is a measurable mapping, then we have $\phi := f^{-1}\Xi \rightarrow \Sigma$, a Boolean homomorphism. Let J be the kernel of ϕ , i.e., the ideal

$$J = \{b \in \Xi \mid f^{-1}(b) = \emptyset\} = \{b \in \Xi \mid b \cap f(S) = \emptyset\}.$$

Then $\Xi \setminus J$ is a support of $\mathcal{M}(T, \Xi)$ ⁴ We have, for $a, b \in \Xi \setminus J$, that $f^{-1}(a), f^{-1}(b)$ are non-empty, thus elements of $\Sigma \setminus \{\emptyset\}$, and $f^{-1}(a) \cap f^{-1}(b) = \emptyset$. In other words, ϕ preserves orthogonality. If $\{a_i\}$ and $\{b_j\}$ are events of $\mathcal{M}(T, \Xi)$ — that is, pairwise disjoint families of non-empty sets in Ξ — then $\{a_i\} \sim \{b_j\}$ iff $\bigcup a_i = \bigcup b_j$ in Ξ , so $\bigcup f^{-1}(\{a_i\}) = f^{-1}(\bigcup a_i) = f^{-1}(\bigcup b_j) = \bigcup f^{-1}(\{b_j\})$, so $f^{-1}(\{a_i\}) \sim f^{-1}(\{b_j\})$. Thus, f^{-1} defines a test-space morphism $\mathcal{M}(T, \Xi) \rightarrow \mathcal{M}(S, \Sigma)$. Note that this is positive iff f is surjective, but is always test-preserving.

2.7 Example (Morphisms of projection manuals). Let $\mathfrak{A}$ and $\mathfrak{B}$ be von Neumann algebras, with corresponding projection manuals $\mathcal{M}(\mathfrak{A})$ and $\mathcal{M}(\mathfrak{B})$ as in 2.3. It is not hard to show that any outcome-preserving morphism $\phi : \mathcal{M}(\mathfrak{A}) \rightarrow \mathcal{M}(\mathfrak{B})$, regarded as a mapping $\phi : \mathbb{P}(\mathfrak{A}) \rightarrow \mathbb{P}(\mathfrak{B})$ (setting $\phi(0) = 0$), is *ortho-additive*, in that $\phi(a + b) = \phi(a) + \phi(b)$ whenever a and b are orthogonal projections; conversely, any ortho-additive mapping $\mathbb{P}(\mathfrak{A}) \rightarrow \mathbb{P}(\mathfrak{B})$ defines an outcome-preserving morphism $\mathcal{M}(\mathfrak{A}) \rightarrow \mathcal{M}(\mathfrak{B})$. In particular, any isomorphism or anti-isomorphism $\phi : \mathfrak{A} \rightarrow \mathfrak{B}$ defines a morphism of the corresponding projection manuals. As a special case, if $\mathfrak{A} = \mathcal{B}(\mathbf{H})$ and $\mathfrak{B} = \mathcal{B}(\mathbf{K})$ for Hilbert spaces $\mathbf{H}$ and $\mathbf{K}$, then any isometry or anti-isometry $W : \mathbf{H} \rightarrow \mathbf{K}$ defines a morphism $\phi(a) = WaW^*$ yields such a morphism. In [26], R. Wright, generalizing Dye's Theorem, showed that if $\mathfrak{A}$ and $\mathfrak{B}$ are von Neumann algebras acting on Hilbert spaces $\mathbf{H}_1$ and

⁴since if E, F are two partitions of S by members of Ξ , and $E \setminus J \subsetneq F \setminus J$, then

$$\bigcup_{a \in E \setminus J} f^{-1}(a) \subsetneq \bigcup_{b \in F \setminus J} f^{-1}(b).$$

The first union is all of S , and so must be the second, a contradiction.

$\mathbf{H}_2$, and $\mathfrak{A}$ has no type I_2 summand, then any completely ortho-additive, unit-preserving map $\phi : \mathbb{P}(\mathcal{A}) \rightarrow \mathbb{P}(\mathcal{B})$ must have the structure

$$\phi(P) = W_1(P \otimes \mathbf{1}_1)W_1^* + W_2(P \otimes \mathbf{1}_2)W_2^*$$

where $\mathbf{K} = \mathbf{K}_1 \oplus \mathbf{K}_2$ and where $W_i : \mathbf{H} \otimes \mathbf{H}_i \rightarrow \mathbf{K}_i$ are an isometry and an anti-isometry.

If $\phi : \mathcal{M} \rightarrow \mathcal{M}'$ and $\psi : \mathcal{M}' \rightarrow \mathcal{M}''$ are morphisms, then

$$(\psi \circ \phi)(x) := \bigcup_{y \in \phi(x)} \psi(y)$$

is again a morphism. It's easily checked that this rule of composition is associative. The mapping $x \mapsto \{x\}$ provides an identity morphism from $\mathcal{M}$ to itself. Thus, we have a category, **Tesp**, of test spaces and morphisms. Note that every morphism $\phi : X = \bigcup \mathcal{M} \rightarrow \mathcal{E}(\mathcal{M}')$ lifts to a mapping $\tilde{\phi} : \mathcal{E}(\mathcal{M}) \rightarrow \mathcal{E}(\mathcal{M}')$, given by $\tilde{\phi}(a) = \bigcup_{x \in a} \phi(x)$, and then $\widetilde{\psi \circ \phi} = \tilde{\psi} \circ \tilde{\phi}$; thus, we can regard $\mathcal{E}$ as a functor from **Tesp** to **Set**.

PM Morphisms It is straightforward that if ϕ is a test space morphism and β is a probability weight on $\mathcal{M}'$, then $\phi^*(\beta) := \beta \circ \phi$ is a sub-normalized positive weight on $\mathcal{M}$. This gives us an affine mapping $\beta \mapsto \phi^*(\beta)$ from $\mathcal{Pr}(\mathcal{M}')$ to $\mathcal{Pr}(\mathcal{M})$. Moreover, if $\mathcal{M}'$ is locally finite, this mapping is continuous with respect to the relative product topologies on $\mathcal{Pr}(\mathcal{M})$ and $\mathcal{Pr}(\mathcal{M}')$.

2.8 Definition. If A and B are probabilistic models, a *PM -morphism* $\phi : A \rightarrow B$ is a test-space morphism $\phi : \mathcal{M}(A) \rightarrow \mathcal{M}(B)$ such that for every $\beta \in \Omega(B)$, $\phi^*(\beta)$ is a *possibly sub-normalized* state of A .

This last condition tells us that ϕ^* defines an affine mapping $\Omega(B) \subseteq \mathbb{V}(A)$ with $u_A(\phi^*(\beta)) \leq 1$ for every $\beta \in \Omega(B)$. This extends uniquely to a positive linear mapping $\phi^* : \mathbb{V}(B) \rightarrow \mathbb{V}(A)$ (it is generally harmless here to use the same symbol for this extension). This dualizes again to give a positive linear mapping $\mathcal{E}(\phi) : \mathcal{E}(A) \rightarrow \mathcal{E}(B)$ with $\mathcal{E}(\phi)(u_A) \leq u_B$.

We shall say that a model A is *full* iff $\Omega(A) = \mathcal{Pr}(\mathcal{M}(A))$. If so, then any test-space morphism $\phi : A \rightarrow B$ defines a PM -morphism. In particular, for the full models associated with measurable spaces, or those associated

with Hilbert spaces of dimension > 2 , all test-space morphisms are PM -morphisms. However, the quantum model associated with $\mathbb{C}^2$ is not full.

It's clear that PM -morphism compose to give PM -morphisms, so we have a category, **Prob**.

2.3 Embeddings, Quotients, and Explanations

We can now explain what we mean by an *explanation* of one model by another. We begin with two cases in which this usage seems relatively obvious.

Let $(\mathcal{M}, \Omega)$ and $(\mathcal{M}', \Omega')$ be probabilistic models with outcome-spaces X and X' , respectively. A mapping $\phi : X \rightarrow X'$ is a quotient mapping of test spaces if it is surjective on both outcomes and tests. $\mathcal{A}$ onto the tests of $\mathcal{A}'$, and $x \perp y \Rightarrow \phi(x) \perp \phi(y)$. Equivalently, ϕ is an outcome- and test-preserving morphism with $\phi(\mathcal{M}) = \mathcal{M}'$ (this implies surjectivity on outcomes). Note that in this case, the mapping $\phi^* : \mathcal{P}r(\mathcal{M}') \rightarrow \mathcal{P}r(\mathcal{M})$ is automatically injective.

2.9 Definition. A PM -morphism $\phi : A \rightarrow B$ is an *embedding* if it is a test space embedding of $\mathcal{A}$ into $\mathcal{A}'$ and ϕ^* is surjective, that is, $\phi^*(\Omega(\mathcal{A}')) = \Omega(\mathcal{A})$. A PM -morphism $q : A \rightarrow A'$ is a *quotient morphism* iff q is a quotient morphism of test spaces (in which case, q^* is automatically injective).

If $\phi : A \rightarrow B$ is an embedding, there is an obvious sense in which B “explains” A , while if ϕ is a quotient mapping, then there is a sense in which A explains B . Indeed, if $\phi : A \rightarrow B$ is an embedding, then B explains A : we obtain A from B by restricting access to certain tests, and identifying states that can no longer be distinguished by this restricted set. If ϕ is a quotient mapping, then A explains B : we obtain B from A by neglecting to distinguish between certain outcomes of A , and restrict the set of states to (some of) those that do not distinguish between outcomes we’ve identified.

It is natural to want “explains” to be transitive. Thus, if A is a quotient of B and B is embedded in C , we should like to say that C (or rather, the quotient map and the embedding) provide an explanation of A . In fact, we will adopt this as our definition.

2.10 Definition (Explanations). An *explanation* of a model A by a model B is a triple (C, ϕ, ψ) where C is model, $\phi : C \rightarrow A$ is a quotient mapping, and $\psi : C \rightarrow B$ is an embedding:

$$\begin{array}{ccc} C & \xrightarrow{\phi} & B \\ \downarrow q & & \\ A & & \end{array}$$

As we will see, this notion of explanation is indeed transitive. First, we address the natural question of why we do not consider the reverse (dual) construction, in which $\psi : A \rightarrow C$ is an embedding and $\phi : B \rightarrow C$ is a quotient mapping. The reason is that any “explanation” of this form gives rise to one of the above type. Let us call a diagram of the form

$$\begin{array}{ccc} & & B' \\ & & \downarrow q \\ A' & \xrightarrow{e} & B \end{array}$$

(still with q a quotient and e , an embedding) a *sub-quotient*.

Thus, explanations are a particular kind of *span* in **Prob**, while subquotients are a kind of co-span. In any category with pullbacks, spans compose according to the rule

$$\begin{array}{ccccc} Z & \dashrightarrow & Y & \longrightarrow & C \\ \downarrow & \lrcorner & \downarrow & & \\ X & \longrightarrow & B & & \\ \downarrow & & & & \\ A & & & & \end{array}$$

To make composites unique, we must choose a pullback for each pair X, Y . This makes composition of spans associative only up to isomorphism: spans serve as (horizontal) morphisms in a bicategory. See [14] for details.

Pullbacks do not generally exist in **Tesp** or **Prob**, essentially, because they don't exist in the category of non-empty sets and mappings. But they do exist in special cases. In particular,

2.11 Proposition. *Sub-quotients have pullbacks, which are explanations.*

Proof: Let $A \xrightarrow{e} C \xleftarrow{q} B$ be a sub-quotient; thus, $e : A \rightarrow C$ and $q : B \rightarrow C$ are respectively an embedding and a quotient of models. Letting A, B and C have X, Y and Z , respectively, we can view e and f as mappings $e : X \rightarrow Z$ and $q : Y \rightarrow Z$, with e injective and q surjective. Define

$$Y' = q^{-1}(e(X))$$

and let $\mathcal{M}' = \{F \subseteq Y' \mid q(F) \in e(\mathcal{M}(A))\}$. Thus, $\mathcal{M}'$ is a sub-test space of $\mathcal{M}(B)$, precisely the one carried by q onto $\mathcal{M}(A)$'s image in $\mathcal{M}(C)$. Note that since q is a quotient, every outcome in $e(X)$ arises as $q(y)$ for some $y \in Y$, so $\bigcup \mathcal{M}' = q^{-1}(e(X)) = Y'$. Since e is injective, we have an inverse mapping $e^{-1} : e(X) \rightarrow X$. By construction, $q' := e^{-1} \circ q|_{Y'}$ is a quotient mapping from $\mathcal{M}'$ onto $\mathcal{M}(A)$, and $e' =$ the inclusion map $Y' \rightarrow Y$ defines a test-space embedding from $\mathcal{M}'$ to $\mathcal{M}(B)$.

Now construct a model $B' = (\mathcal{M}', \Omega')$ where $\Omega' = \Omega(B)|_{Y'}$, that is, Ω' consists of restrictions to Y' of states of B . The inclusion mapping $e' : Y' \rightarrow Y$ then defines an embedding $B' \rightarrow B$. We need to show that $q' : Y' \rightarrow X(A)$ is also a PM-morphism. To this end, let $\alpha \in \Omega(A)$. Since e is an embedding, there is some $\gamma \in \Omega(C)$ with $\alpha = e^*(\gamma)$. Let $\beta' = q^*(\gamma)|_{Y'}$, a state of B' . We claim that $\beta' = (q')^*(\alpha)$. For if $y \in Y'$ and $z = q(y)$, then $z \in e(X)$, i.e., there is a unique $x = e^{-1}(z)$ with $q(y) = e(x)$. Now we have

$$\begin{aligned} (q')^*(\alpha) &= \alpha \circ e^{-1} \circ q|_{Y'} \\ &= \gamma \circ e \circ e^{-1} \circ q|_{Y'} \\ &= \gamma \circ q|_{Y'} \\ &= (\gamma \circ q)|_{Y'} \\ &= q^*(\gamma)|_{Y'} \\ &= \beta'. \end{aligned}$$

This shows that q' is also a morphism of models.

To see that that (B', e', q') is the pullback of q and e , apply the $\mathcal{E}$ functor: we find that $\mathcal{E}(B') = \tilde{q}^{-1}(\tilde{e}(\mathcal{E}(A)))$. Since $\tilde{e}$ and $\tilde{q}$ are still injective and surjective, respectively, we see that $\mathcal{E}(B')$ is the pullback of these in **Set**. Thus, given morphisms $\phi : D \rightarrow A$ and $\psi : D \rightarrow B$ with $e \circ \phi = q \circ \psi$, and regarding these as mappings $U \rightarrow \mathcal{E}(A)$ and $U \rightarrow \mathcal{E}(B)$ where $U = \bigcup \mathcal{M}(D)$, we have a unique mapping $\xi : U \rightarrow \mathcal{E}(D)$ with $\tilde{q}' \circ \xi = \tilde{\phi}$ and $\tilde{e}' \circ \xi = \tilde{\psi}$. This is given by ψ 's co-restriction to $\mathcal{E}(B') \subseteq \mathcal{E}(B)$ (as ψ 's range lies here). Since ψ is a morphism of test spaces, so is ξ . It remains to check that ξ is a PM-morphism. Let $\beta' \in \Omega'$. Then $\beta' = \beta|_{Y'}$ for some $\beta \in \Omega(B)$. Thus,

$$\xi^*(\beta') = (e' \circ \xi)^*(\beta) = \psi^*(\beta);$$

since ψ is a PM-morphism, $\psi^*(\beta) \in \Omega(D)$. Thus, ξ is a morphism of models, and we are done. $\square$

It follows that explanations compose associatively (up to a canonical isomorphism), and we can, if we wish, form a bicategory of probabilistic models and explanations. We will not pursue this here. For our purposes, the important points are that our notion of explanation is indeed transitive, and that quotients and embeddings are (in their dual ways) instances of explanations.

3 Classical Explanations

We will now show that every locally finite probabilistic model has an explanation, and indeed a *canonical* one, in terms of a Borel model. This is largely a repackaging of well-known results in the present framework, but we think this offers some slight additional generality, and a more-than-slight increment of clarity.

3.1 Weak-* integrals and barycenters

We will need a few facts from analysis. First, recall that if μ is a finitely additive probability measure on an algebra Σ of subsets of a set S , and if $f : S \rightarrow \mathbb{R}$ is a bounded Σ -measurable function, then $\int_S f d\mu$ exists, with $|\int_S f d\mu| \leq \|f\| \|\mu\|$ where $\|f\|$ and $\|\mu\|$ are the supremum and variation

norms of f and μ , respectively. This gives us a bounded linear mapping $M(S) \rightarrow B(S)^*$, where $B(S)$ is the space of bounded measurable functions on (S, Σ) and $M(S)$ is the space of bounded finitely-additive measures on Σ ; in fact, one can show that this is a Banach-space isomorphism [5, Theorem IV.5.1]. Accordingly, from now on we will identify measures $\mu \in M(S)$ with the corresponding functionals, writing $\mu(f)$ for $\int_S f(s) d\mu(s)$ whenever convenient.

Weak-* integrals V is any normed space, a function $\phi : S \rightarrow V^*$ is *weak-* measurable* iff the function $\phi^*(v) : s \mapsto \phi(s)(v)$ is measurable for every $v \in V$. If ϕ is bounded, so is $\phi^*(v)$ for every $v \in V$, and thus if μ is a bounded finitely-additive measure on S , we can form

$$\widehat{\mu}(\phi) : v \mapsto \int_S \phi(s)(v) d\mu(s)$$

which defines an element of V^* . We also write

$$\widehat{\mu}(\phi) = \int_S \phi(s) d\mu(s) :$$

this is called the *weak-* integral* of ϕ with respect to μ . If μ is a probability measure, we can regard this as the average of the functionals $\phi(s)$ with respect to μ .

Now consider the case in which $V = B(S)$ and $V^* = M(S)$. If δ_s is the point-mass at s , the mapping $s \mapsto \delta_s$ is bounded and weak-* measurable (the variation norm of δ_s is 1, and for $f \in B(S, \Sigma)$, $\delta_s(f) = f(s)$.) Hence, for every probability measure μ on S , we can average the point masses δ_s against μ , in the weak-* sense, to obtain

$$\widehat{\mu} := \widehat{\mu}(\delta) = \int \delta_s d\mu(s) \in B(S)^* = M(S). \quad (1)$$

In fact, $\widehat{\mu} = \mu$. To see this, let $a \in \Sigma$ with indicator function I_a : we have

$$\widehat{\mu}(I_a) = \int_S \delta_s(I_a) d\mu(s) = \int_S I_a(s) d\mu(s) = \mu(a).$$

Barycenters If K is a compact convex subset of a topological vector space V . Let $A(K)$ be the space of all continuous affine (convex-linear) functionals

on K , and let μ be a Baire probability measure on K . Then there is a unique element $\hat{\mu} \in K$, the *barycenter* of K , with $f(\hat{\mu}) = \int_K f d\mu$ for all $f \in A(K)$. The mapping $\mu \mapsto \hat{\mu}$ is clearly bounded and linear. Moreover, we have the following: [18, Theorems 9.1 and 9.2]:

3.1 Theorem. *Every point $\alpha \in K$ can be obtained as $\hat{\mu}$ for some measure μ supported on the closure of K 's extreme points.*

3.2 Borelification

A version of the following construction appears in [21]

3.2 Lemma. *If $\mathcal{A}$ is a locally finite, semi-classical space with outcome-space X , then the set $\mathcal{Pr}(\mathcal{A})$ of all probability weights and the set S of df probability weights on $\mathcal{A}$ are both closed, hence compact, in $[0, 1]^X$, and $S = \text{Ext}(\mathcal{Pr}(\mathcal{A}))$.*

Proof. Let $K = \mathcal{Pr}(\mathcal{A})$, and suppose $\alpha \in [0, 1]^X$ is not in K . Then for some test $E \in \mathcal{A}$, $\sum\{\alpha(x) \mid x \in E\} \neq 1$. Hence, if $\beta \in [0, 1]^X$ with $\beta(x)$ sufficiently close to $\alpha(x)$ for each x in the finite set E , β will also fail to be a probability weight, and thus K is closed. Since S is the intersection of K with the closed set $\{0, 1\}^X$, it is also closed. To see that $S = \text{Ext}(K)$, note that since $\mathcal{A}$ is semiclassical, the probability weights $\alpha \in \mathcal{A}$ amount to indexed families (α_E) where α_E is a probability weight on E for each $E \in \mathcal{A}$. If one of these, say α_F , is a non-trivial convex combination $\alpha_F = \lambda\beta_F + (1 - \lambda)\gamma_F$ where β_F and γ_F are probability weights on F , then define $\beta = (\beta_E)$ by setting $\beta_E = \alpha_E$ for all $E \neq F$, with β_F the given probability weight. Define $\gamma = (\gamma_E)$ similarly. Then $\alpha = \lambda\beta + (1 - \lambda)\gamma$ is also a non-trivial convex combination. So if (α_E) belongs to $\text{Ext}(K)$, each α_E must be an extreme point of the set of probability weights on E , and hence, a point mass, and thus α is df. It is clear that df probability weight is an extreme point. $\square$

We are now going to show that *almost* every semi-classical test space $\mathcal{A}$ can be embedded into a Borel test space. The only potential obstacle is that if x and y are distinct outcomes with $\{x\}, \{y\} \in \mathcal{A}$, then every df state assigns probability one to both.

3.3 Theorem. *Every locally finite, semiclassical PM $A = (\mathcal{A}, \Omega)$ having at most one single-outcome test, has a canonical embedding into a classical model.*

Proof. Let S be the set of d.f. probability weights on $\mathcal{A}$. By Lemma 3.2, S and $\mathcal{Pr}(\mathcal{A})$ are compact with $S = \text{Ext}(\mathcal{Pr}(\mathcal{A}))$. Let Σ be the Borel algebra of S and $\mathcal{B}(S)$ be the corresponding Borel test space. That is, $\mathcal{B}(S)$ consists of finite partitions of S by Borel sets. Define a map ϕ from the outcomes of $\mathcal{A}$ to the outcomes of $\mathcal{B}$ — the non-empty Borel sets — by setting

$$\phi(x) = \{s \in S \mid s(x) = 1\}.$$

Note that $\phi(x)$ is the intersection of a closed set of $[0, 1]^X$ with S , hence, is indeed a Borel set of S . Now, if $x \perp y$, $\phi(x) \cap \phi(y) = \emptyset$, and $\{\phi(x) \mid x \in E, \phi(x) \neq \emptyset\}$ is a partition of S . It is straightforward that if $a, b \in \mathcal{E}(\mathcal{A})$ with $a \sim b$, $\bigcup \phi(a) = \bigcup \phi(b)$, so $\phi(a) \sim \phi(b)$ in $\mathcal{B}(S)$. Thus, ϕ is a test-preserving test-space morphism. This much holds regardless of $\mathcal{A}$ -s structure. But since $\mathcal{A}$ is semiclassical, we also have $\phi(x) \neq \emptyset$ for every $x \in X$. If x and y are distinct outcomes, at most constitutes a one-outcome test, so there exists some $s \in S$ with $s(x) \neq s(y)$. Thus, ϕ is injective on outcomes, and thus, a test-space embedding. Finally, suppose $\alpha \in \Omega$. By Theorem 3.1, α is the barycenter, $\hat{\mu}$, of a regular Borel probability measure μ on S . This means that for the continuous linear functional $\hat{x} : \mathbb{R}^X \rightarrow \mathbb{R}$ of evaluation at x , we have

$$\hat{x}(\alpha) = \int_S \hat{x} d\mu.$$

For $s \in S$ we have $\hat{x}(s) = 1$ iff $s \in \phi(x)$ and is 0 otherwise, so $\int_S \hat{x} d\mu = \mu(\phi(x))$. Thus

$$\alpha(x) = \mu(\phi(x)),$$

or, equivalently, $\alpha = \varphi^*(\mu)$. Taking Ω' to be the set of all probability measures μ on (S, Σ) with $\hat{\mu}$ in Ω , and letting $B = (\mathcal{B}, \Omega')$, ϕ is a PM-embedding in the sense of Definition 2.9. $\square$

Semiclassical covers We have just seen that locally finite semiclassical PMs can always be interpreted as sub-models of classical models. On the other hand, every locally finite PM is a “quotient” of a locally finite semi-classical model in an obvious way ([8, p. 184], see also [21, Section 6.2]). Let $\mathcal{A}$ be a

test space. For each $E \in \mathcal{A}$, let $\tilde{E} = \{(x, E) \mid x \in E\}$. The *semiclassical cover* of $\mathcal{A}$ is the test space $\tilde{\mathcal{A}}$ given by

$$\tilde{\mathcal{A}} = \{\tilde{E} \mid E \in \mathcal{A}\}.$$

We write $\tilde{X}$ for $\bigcup \mathcal{A} = \{(x, E) \mid x \in E \in \mathcal{A}\}$. Every probability weight $\alpha \in \mathcal{Pr}(\mathcal{A})$ gives rise to a probability weight $\tilde{\alpha}$ on $\tilde{\mathcal{A}}$ defined by $\tilde{\alpha}(x, E) = \alpha(x)$. By the semiclassical cover of a PM $A = (\mathcal{A}, \Omega)$, we mean the model $\tilde{A} = (\tilde{\mathcal{A}}, \tilde{\Omega})$ where

$$\tilde{\Omega} = \{\tilde{\alpha} \mid \alpha \in \Omega\}.$$

It is then straightforward that the mapping $q : \tilde{X} \rightarrow X$ given by $q(x, E) = x$ is a quotient morphism of models.

Thus, every PM is a quotient of its semi-classical cover. Since every locally finite semiclassical model has a canonical classical embedding, we have the

3.4 Theorem. *Every locally finite PM A with at most one one-outcome test, has a canonical classical explanation*

$$\begin{array}{ccc} \tilde{A} & \xrightarrow{\phi} & \mathcal{B}(S) \\ \downarrow q & & \\ A & & \end{array}$$

where S is the set of dispersion-free probability weights on $\tilde{A}$, understood as a subspace of $[0, 1]^{\tilde{X}}$.

Borelification as a functor Let us write $\mathbf{Bor}(A)$ for the Borel model arising in the canonical classical explanation of a probabilistic model A . Thus, $\mathbf{Bor}(A)$'s test space is the space of finite partitions of the set S of dispersion-free states on $\tilde{S}$ by Borel sets, and $\Omega(\mathbf{Bor}(A))$ is the set of finitely-additive Borel probability measures on S having barycenter in $\Omega(A)$. We will refer to $\mathbf{Bor}(A)$ as the *Borelification* of A . It is not difficult to show that $A \mapsto \mathbf{Bor}(A)$ is the object part of an endofunctor $\mathbf{Bor} : \mathbf{Prob} \rightarrow \mathbf{Prob}$.

As discussed elsewhere [2, 25], one can view a *probabilistic theory* as a functor $F : \mathcal{C} \rightarrow \mathbf{Prob}$, where $\mathcal{C}$ is a category representing systems and processes.

The idea is that to each system A , M attaches a concrete probabilistic mode, $M(A) = (\mathcal{M}(A), \Omega(A))$, and to each process $f : A \rightarrow B$, M attaches a morphism $Mf : M(A) \rightarrow M(B)$ of probabilistic models. Thus, if $M : \mathcal{C} \rightarrow \mathbf{Prob}$ is any probabilistic theory, we have a canonical “Borelification” $\mathbf{Bor} \circ M$ from $\mathcal{C}$ to the category of Borel models and morphisms between these.

3.3 Ontological Representations

A related formulation of these ideas can be given in terms of so-called *ontological representations* [11, 15, 19]. Adapted to our current language, the definition is as follows:

3.5 Definition (Ontological representations). An *ontological representation* of a PM A is measurable space (S, Σ) of “ontic states” with a function p assigning to each $s \in S$ and each test $E \in \mathcal{A}$ a probability weight $x \mapsto p(x | E, s)$ on E so that

- (a) For each $x \in E$, the function $s \mapsto p(x | E, s)$ is measurable, and
- (b) For each $\alpha \in \Omega_A$, there is a probability measure μ on S with

$$\alpha(x) = \int_S p(x | E, s) d\mu.$$

If p takes only values 0, 1 we say the representation is *sharp* and if $p(x | E, s)$ is independent of E for every x and s , it is *non-contextual*.

Perhaps the simplest example is the *Beltrametti-Bugajski representation* [3, 12, 15]: given a model $A = (\mathcal{M}, \Omega)$ with Ω compact, let S be the set of extreme points of Ω , i.e., the pure states of A , with Borel structure inherited from $[0, 1]^{X(A)}$. For every $x \in E \in \mathcal{M}$, and every $\alpha \in S$, let $p(x | E, \alpha) = \alpha(x)$. This is evidently non-contextual. As long as Ω is separating, it yields an embedding of A into the unsharp model associated with S .

More generally, an ontological representation of A based on a measurable space (S, Σ) yields an embedding of $\tilde{A}$ into the unsharp Borel model $\mathcal{B}^o(S, \Sigma)$ associated with (S, Σ) , as discussed in Section 2.1.

Recall that the tests of this model consist of (graphs of) lists of positive bounded measurable functions summing to one, and that the states are induced by finitely additive probability measures on (S, Σ) in an obvious way. Given an ontological representation p of $A = (\mathcal{A}, \Omega)$, define $\phi : \tilde{X} \rightarrow B(S)$ by $\phi(x, E)(s) = p(x|E, s)$. This extends uniquely to a morphism $\phi^o : \tilde{A}^o \rightarrow \mathcal{B}^o(S)$, which will be injective iff Ω is a separating set of states. For any probability measure μ on (S, Σ) ,

$$\phi^{o*}(\mu)(x, E) = \widehat{\mu}(f_{x,E}) = \int_S f_{x,E}(s) d\mu(s),$$

so condition (c) tells us that ϕ^* is surjective on states. Thus, we have an embedding $\phi^o : (\tilde{A})^o \rightarrow \mathcal{M}^o(B(S, \Sigma))$, where $B(S, \Sigma)$ is the order-unit space of bounded measurable real-valued functions on (S, Σ) , as in 2.1. Moreover, we have a natural quotient mapping $q : (\tilde{A})^o \rightarrow A$, given by $q(i, (x, E)) = x$. Thus, we have an explanation

$$\begin{array}{ccc} (\tilde{A})^o & \xrightarrow{\phi^o} & \mathcal{M}^o(B) \\ \downarrow q & & \\ A & & \end{array}$$

The unsharp classical model $\mathcal{M}^o(B(S, \Sigma))$ can in turn be explained in terms of the sharp classical model $\mathbf{Bor}(\mathcal{M}^o(B(S, \Sigma)))$.

In fact, the mapping $(x, E) \mapsto f_{x,E}$ is an embedding. By (c), if $f_{x,E} = f_{y,F}$, then $\tilde{\alpha}(x, E) = \widehat{(y, F)}$ for all $\tilde{\alpha} \in \Omega(\tilde{A})$, i.e., $\alpha(x) = \alpha(y)$ for all $x, y \in X_A$. We are (or should be) assuming a separating set of states here, so $x = y$. Condition (a) tells us that this is test-preserving, so we have an embedding of test spaces. Finally, for any probability measure μ on (S, Σ) ,

$$\phi^*(\mu)(x, E) = \widehat{\mu}(f_{x,E}) = \int_S f_{x,E}(s) d\mu(s),$$

so condition (c) tells us that ϕ^{o*} is surjective on states.

Since A is a quotient of its semiclassical cover, this yields a unsharp classical explanation of A . Conversely, any unsharp classical explanation of A obtained by embedding its semiclassical cover $\tilde{A}$ into an unsharp classical model, can be seen to yield an ontological representation.

In summary: a sharp, respectively unsharp, contextual ontological representation of a probabilistic model A is simply an embedding $\phi : \tilde{A} \rightarrow C$ where C is an sharp, respectively unsharp, Borel model. Thus, we can paraphrase Theorem 3.4 as saying that every probabilistic model arises as the quotient of one admitting a sharp ontological model.

4 Locality

By Theorem 3.4, every locally finite PM admits a classical explanation, hence, by the discussion in Section 3.3, it admits a sharp ontological representation. The situation for composite systems is more delicate: while such a composite, if locally finite, admits an ontological representation, this will generally not satisfy a reasonable locality criterion with respect to the component models. We return to this below.

4.1 Composites and entanglement

We begin with an overview of the theory of non-signalling composites of probabilistic models. For a more leisurely account, see [22].

Let A and B be test spaces, with outcome-spaces $X(A)$ and $X(B)$. We abuse notation slightly to write $\mathcal{M}(A) \times \mathcal{M}(B)$ for the set of all *product tests* $E \times F$, with $E \in \mathcal{M}(A)$ and $F \in \mathcal{M}(B)$. We understand these to represent experiments in which E and F are performed on separate systems, and their outcomes, collated. Note that the outcome-space of $\mathcal{M}(A) \times \mathcal{M}(B)$ is $X(A) \times X(B)$. A *joint probability weight* ω on $\mathcal{M}(A)$ and $\mathcal{M}(B)$ is simply a probability weight on $\mathcal{M}(A) \times \mathcal{M}(B)$. We say that ω is *non-signalling* iff

the *marginal weights*

$$\omega_1(x) = \sum_{y \in F} \omega(x, y) \quad \text{and} \quad \omega_2(y) = \sum_{x \in E} \omega(x, y)$$

are independent of the choice of $E \in \mathcal{M}(A)$ and $F \in \mathcal{M}(B)$. In this case, we define the *conditional states* $\omega_{1|y} \in \mathcal{Pr}(\mathcal{M}(A))$ and $\omega_{2|x} \in \mathcal{Pr}(\mathcal{M}(B))$ by

$$\omega_{1|y}(x) = \omega(x, y)/\omega_2(y) \quad \text{and} \quad \omega_{2|x}(y) = \omega(x, y)/\omega_1(x),$$

respectively. If $\omega_{1|y} \in \Omega(A)$ and $\omega_{2|x} \in \Omega(A)$ for every $x \in X(A)$ and $y \in X(B)$, then we say that ω is a non-signalling joint *state* of A and B . We write $\Omega_{NS}(A, B)$ for the set of all non-signalling joint states of A and B . The *minimal non-signalling composite*⁵ of A and B is then

$$A \times_{NS} B := (\mathcal{M}(A) \times \mathcal{M}(B), \Omega_{NS}(A, B)).$$

The simplest example of a non-signalling probability weight on A and B is a *product weight*, $\alpha \otimes \beta$, given by

$$(\alpha \otimes \beta)(x, y) = \alpha(x)\beta(y)$$

for all $x \in X(A)$ and $y \in X(B)$, where α and β are probability weights on A and B , respectively. If α and β are states, then of course we speak of $\alpha \otimes \beta$ as a product state on A and B .

4.1 Definition. A *separable joint state* of A and B is a joint state lying in the closed convex hull of the product states on A and B . A joint state that is not separable is an *entangled joint state*.

Note that the distinction between separability and entanglement of *states* depends on the models, A and B , and, in particular, their state spaces as well as their test spaces. In particular, if we replace $\Omega(A)$ and $\Omega(B)$ by the full sets $\mathcal{Pr}(A)$ and $\mathcal{Pr}(B)$ of all probability weights, we will speak of separable and entangled *weights* on A and B . Thus, a joint state can be separable as a joint probability weight — that is, separable with respect to $\mathcal{Pr}(A)$ and $\mathcal{Pr}(B)$ — while being entangled as a joint state. This is even possible for classical models:

⁵For *cognoscenti*: the state space of $A \times B$ is the *maximal* tensor product of the state spaces of A and B ; thus, the effect cone of $A \times_{NS} B$ is the minimal tensor product of those of A and B

4.2 Example (Classical entanglement). Let $\mathcal{B}$ be the test space of finite partitions of $[0, 1]$ by Borel sets. Outcomes of $\mathcal{B}$ are non-empty Borel sets, and probability weights on $\mathcal{B}$ are simply finitely-additive Borel probability measures on $[0, 1]$, as restricted to non-empty Borel sets. We will ignore this tiny distinction, and simply identify $\mathcal{Pr}(\mathcal{B})$ with the simplex Δ of all finitely additive Borel measures on $[0, 1]$. Now let λ denote Lebesgue measure on $[0, 1]$, and let Δ_λ be the collection of all Borel measures on $[0, 1]$ absolutely continuous with respect to λ . This is a closed face of Δ , and thus, itself also a simplex. Now consider the probabilistic model $B = (\mathcal{B}, \Delta_\lambda)$. Every separable probability weight on $B \times_{NS} B$ — that is, every convex combination of product states, and every pointwise limit of such states in $[0, 1]^{\Sigma \times \Sigma}$ — is absolutely continuous with respect to Lebesgue measure $\lambda \otimes \lambda$ on $[0, 1]^2$. Let ω be the measure on $[0, 1]^2$ given by $\omega(a, b) = \lambda(a \cap b)$. This is not absolutely continuous with respect to $\lambda \otimes \lambda$, but is non-signalling, with conditional weights that do belong to Δ_λ . Hence, ω is a non-separable, non-signalling state. On the other hand, every finitely-additive measure on $[0, 1]^2$ has the form $\mu = \int_s \delta_s \otimes \delta_t d\mu(s, t)$ and is therefore separable with respect to the set $\mathcal{Pr}(\mathcal{M}_A)$ of all probability weights.

General non-signalling composites The minimal non-signalling composite $A \times_{NS} B$ allows only product tests. To allow for a greater abundance of joint tests, it's natural to introduce the following

4.3 Definition. A *non-signalling composite* of A and B is a probabilistic model AB , together with a positive, injective, test-preserving morphism of models $\pi : A \times_{NS} B \rightarrow AB$.

Thus, every pair of outcomes $x \in X(A)$ and $y \in X(B)$ correspond to an event $\pi(x, y) \in \mathcal{E}(AB)$, and every state in $\Omega(AB)$ pulls back to a non-signalling joint state $\pi^*(\omega) \in \Omega_{NS}(A, B)$, given by

$$\pi^*(\omega)(x, y) = \omega(\pi(x, y)).$$

In a slight abuse of terminology, we will refer to $\pi^*(\omega)(x, y)$ as the *restriction* of ω to $A \times_{NS} B$.

When π^* is injective, so that every state $\omega \in \Omega(AB)$ is determined by its restriction joint state $\pi^*(\omega)$ to $A \times_{NS} B$, we say that (AB, π) is *locally tomographic*. When every product state $\alpha \otimes \beta$ arises as the restriction $\pi^*(\omega)$ of some state $\omega \in \Omega(AB)$, we say that (AB, π) is *strong*.

In more physical terms, a composite is locally tomographic iff local measurements suffice to determine the global state, and strong, if all pairs of local states can be independently prepared (possibly in multiple ways) as the marginals of a suitable state of the global composite system. Where either of these properties is lacking, the definition of entanglement requires some care.

4.4 Definition. *Let (AB, π) be a non-signaling composite of A and B . A joint state of the form $\pi^*(\omega)$ for some $\omega \in \Omega(AB)$ is preparable in AB . A state $\omega \in \Omega(AB)$ is separable iff $\pi^*(\omega)$ lies in the closed convex hull of the preparable product states.*

If (AB, π) is locally tomographic, then AB is strong iff $\Omega(AB)$ contains all product states. If (AB, π) is strong, then $\omega \in \Omega(AB)$ is separable iff $\pi^*(\omega)$ is a separable joint state in the sense of Definition 4.4. But non-locally tomographic and non-strong composites do arise, even in quantum theory, as we now illustrate.

4.5 Example (Quantum Composites). Let A and B be the full projective quantum model corresponding to Hilbert spaces $\mathbf{H}_A$ and $\mathbf{H}_B$ — both real, or both complex — as in Example 2.3. Thus, $\mathcal{M}(A) = \mathcal{M}(\mathbf{H}_A)$ consists of all maximal pairwise orthogonal projections on $\mathbf{H}_A$, and $\Omega(A) = \Omega(\mathbf{H}_A)$ is the set of probability weights α_W corresponding to density operators W on $\mathbf{H}_A$, according to $\alpha_W(p) = \text{Tr}(Wp)$, and similarly for B . The mapping $\pi : \mathcal{M}(\mathbf{H}_A) \times \mathcal{M}(\mathbf{H}_B) \rightarrow \mathcal{M}(\mathbf{H}_A \otimes \mathbf{H}_B)$ given by $\pi(p, q) = p \otimes q$ is easily seen to be a test-space embedding. To obtain a non-signalling composite AB , we must equip $\mathcal{M}(AB) = \mathcal{M}(\mathbf{H}_A \otimes \mathbf{H}_B)$ with a suitable state-space. There are multiple options. For instance,

- (a) One might take $\Omega(AB)$ to be the full set $\Omega(\mathbf{H}_A \otimes \mathbf{H}_B)$ of density operators on $\mathbf{H}_A \otimes \mathbf{H}_B$, as in standard QM.
- (b) One could take $\Omega(AB)$ to be the set of separable bipartite states.
- (c) If $\mathbf{H}_A = \mathbf{H}_B$, one could take $\Omega(AB)$ to consist of Bosonic or Fermionic mixed states — that is, density operators W on $\mathbf{H}_A \otimes \mathbf{H}_B$ satisfying $\sigma W \sigma = W$, or those satisfying $\sigma W \sigma = -W$, where σ is the swap operator $\sigma : x \otimes y \rightarrow y \otimes x$.

As is well-known, the composite in (a) is locally tomographic if $\mathbf{H}_A$ and $\mathbf{H}_B$ are complex, but not if they are real. In both (a) and (b), the composites are strong, but in (b), there are, by fiat, no entangled states. Neither of the composites in (c) are strong. In the Fermionic case, no product state is preparable, so all states are entangled; in the Bosonic case, products of identical states are preparable, and mixtures of these are separable.

4.2 Bell-locality

We now come to the crucial idea of Bell-locality. The following is a direct paraphrase of the standard definition in terms of ontological models (e.g., [15]) into our language:

4.6 Definition. Let $\pi : A \times_{NS} B \rightarrow AB$ be a non-signalling composite and $\phi : \widetilde{AB} \rightarrow B(S, \Sigma)$ be a test-space embedding into the finitely-additive Borel model associated with a measurable space (S, Σ) . We say that ϕ is *Bell-local* with respect to π iff for each point mass δ_s , $s \in S$, the probability weight γ_s with $\widetilde{\gamma}_s := \phi^*(\delta_s)$ on AB restricts to a product weight, that is, there are probability weights α_s on $\mathcal{M}(A)$ and β_s on $\mathcal{M}(B)$ with $\pi^*(\gamma_s) = \alpha_s \otimes \beta_s$.

It is important to notice here that γ_s , α_s , and β_s need not be *states* of AB , A , or B . Thus, locality, as defined above, is entirely a property of the test-spaces $\mathcal{M}(A)$, $\mathcal{M}(B)$, and $\mathcal{M}(AB)$, and the test-space embeddings π and ϕ : the chosen state spaces $\Omega(A)$, $\Omega(B)$ and $\Omega(AB)$ do not enter into it. Or, equivalently, the definition refers only to full models, in which $\Omega(A) = \mathcal{Pr}(\mathcal{M}(A))$, $\Omega(B) = \mathcal{Pr}(\mathcal{M}(B))$, and $\Omega(AB) = \mathcal{Pr}(\mathcal{M}(AB))$.

At the same time, even if all probability *weights* on $\mathcal{M}(AB)$ are separable, it is possible for Ω_{AB} to contain states that are *not* convex combinations of product *states* in Ω_A and Ω_B . These would then be entangled as states, though not as probability weights, as in Example 4.2.

Returning to Definition 4.6, suppose that the models A , B and AB therein are locally finite, so that their state spaces are compact. For every probability weight μ on S , we can write

$$\mu = \int_S \delta_s d\mu(s)$$

as in Equation 1. Since C is locally finite, $\pi^* : \mathcal{Pr}(C) \rightarrow \mathcal{Pr}(AB)$ is continuous (see remarks above in 2.2), and thus, Borel measurable. Since by Bell locality $\pi^* \circ \phi^*(\delta_s) = \alpha_s \otimes \beta_s$, we have

$$\pi^* \phi^*(\mu) = \int_S \pi^* \phi^*(\delta_s) d\mu(s) = \int_S \alpha_s \otimes \beta_s d\mu(S)$$

which belongs to the closed convex hull of the product weights, i.e., to the set of separable probability weights, on $A \times B$. Since $\phi \circ \pi$ is an embedding, it follows that all states of AB are separable *as probability weights*; that is, they are all (integral) convex combinations of weights restricting to products of probability weights. In summary, we have the following

4.7 Proposition. *If a non-signalling composite $\pi : A \times_{NS} B \rightarrow AB$ of locally finite probabilistic models admits a Bell-local classical embedding, then the restrictions of all states of AB are separable as joint probability weights.*

For a given classical embedding ϕ to be local, therefore, all states on AB must be separable as probability weights. If AB supports even one state that is not separable in this sense, then AB admits no local classical embedding.

Since $\tilde{\pi} : \tilde{A} \times_{NS} \tilde{B} \rightarrow \tilde{AB}$ is a non-signalling composite, Proposition 4.7 applies to the canonical classical embedding of $\tilde{AB}$ into the Borel test space over the set S of dispersion-free probability weights on $\tilde{AB}$. It is not difficult to see that a dispersion-free probability weight on $\mathcal{M}(A) \times \mathcal{M}(B)$ is non-signalling iff it is a product weight. Thus unless $\mathcal{M}(A)$ and $\mathcal{M}(B)$ consist of single tests, the vast majority of the d.f. probability weights $\phi^*(\delta_s)$ arising from points $s \in S$ are signalling. The classical embedding, then, represents non-signalling states of $A \times_{NS} B$ as weighted averages of dispersion-free, *signalling*, probability weights. As a simple but instructive illustration of this, consider $A = B = \{\{x, y\}, \{u, v\}\}$. The d.f. joint probability weight defined by

$$\omega(x, x) = \omega(x, u) = \omega(u, x) = \omega(u, v) = 1$$

is signalling (since, e.g., $\omega_{2|E}(u) = 1$ while $\omega_{2|F}(u) = 0$). Similarly,

$$\omega'(y, y) = \omega(y, v) = \omega(v, y) = \omega(v, u) = 1$$

is signalling. But $\frac{1}{2}(\omega + \omega')$ is the well-known entangled PR-box state [17]. ,

4.8 Remark. It is not difficult to extend Definition 4.6 to apply to arbitrary classical explanations (C, q, e) of a composite AB . One can also amend it to require that $\pi^*(\delta_s) = \gamma_s$ have the form $q^*(\alpha_s \otimes \beta_s)$ for *states* α_s and β_s in $\Omega(A)$ and $\Omega(B)$, respectively (in which case, of course, $\gamma_s \in \Omega(C)$). In that case, we might say that the explanation is *state-local*. The same argument as give above then shows that if AB has a state-local classical explanation, its states are all separable *as states*.

5 Conclusion

Bell’s Theorem and subsequent tests of the BCHSH inequality establish that one can prepare joint states of quantum systems that do not have local classical explanations, even in our broad sense. In some quarters, the significance of this is a matter of contention. According to one point of view, Bell’s Theorem sets up a tension between locality and (some version of) “classicality”. According to another, the theorem simply establishes that the world is non-local, full stop. See [16, 20] for a particularly sharp exchange of views on this matter. However, there nothing sacred about our definition of classicality, or that of locality. What has been shown above is that *one* notion of locality is in conflict with *one* notion of classicality. The following attitudes all seem to us consistent with this, and reasonable on their own terms (though we lean towards the first one).

- (1) The non-product dispersion-free probability weights on $\widetilde{AB}$ are signalling, and hence, unphysical. The world is non-classical, since non-commuting quantum observables are *in fact* not jointly testable. Whether it is local or not depends on whether or not one understands entanglement *per se* to be a form of non-locality. If so, the world is non-local *because* it is non-classical.
- (2) The dispersion-free probability weights on $\widetilde{AB}$ provide a coherent ontology of the world; the fact that the states we observe are all non-signalling reflects a degree of fine-tuning that is simply part of the world’s architecture. The world is classical, but non-local.
- (3) Locality is part of what we *mean* by classicality. The world is non-classical *because* it is non-local.

References

- [1] H. Barnum, M. Graydon, and A. Wilce, Composites and categories of euclidean Jordan algebras, *Quantum* **4** (2020)
- [2] H. Barnum, M. Graydon, and A. Wilce, Locally tomographic shadows, in S. Mansfield, B. Valiron, V. Zamdzhiev (eds.): *Quantum Physics and Logic 2023 (QPL 2023) EPTCS* **384** (2023),
- [3] E. Beltrametti and S. Bugajski, A classical extension of quantum mechanics, *J. Phys. A* **28** (1995)
- [4] J. L. Bunce and J. D. M. Wright, The Mackey-Gleason problem, *Bull. Amer. Math. Soc.* **26** (1992)
- [5] N. Dunford and J. Schwarz, *Linear Operators*, vol. I., Wiley, 1988
- [6] D. Foulis and C. Randall, The empirical logic approach to the physical sciences, in A. Härtkammer and H. Neumann (eds.), *Foundations of Quantum Mechanics in Ordered Linear Spaces*, *Lecture Notes in Physics* **29**, Springer, 1974
- [7] D. Foulis and C. Randall, Manuals, morphisms and quantum mechanics, in A. R. Marlowe (ed.), *Mathematical Foundations of Quantum Theory*, AP 1978
- [8] D. Foulis and C. Randall, The operational approach to quantum mechanics, in C. A. Hooker (Ed.), *Physical Theory as Logico-Operational Structure*, University of Western Ontario Series in Philosophy of Science, Reidel, 1979 Reidel,
- [9] D. Foulis and C. Randall, Dirac Revisited, in P. Lahti and P. Mittelstaedt, eds., *Symposium on the foundations of modern physics. 50 years of the Einstein-Podolsky-Rosen Gedankenexperiment*, World Scientific, 1985.
- [10] D. Foulis, R. Greechie, and G. Rüttimann, Logico-algebraic structures II: supports in test spaces, *Int. J. Theor. Physics* **32** (1993)
- [11] A. Gheorgiu and C. Heunen, Ontological models for quantum theory as functors, in B. Coecke and M. Leifer (eds.), *Quantum Physics and Logic 2019 (QPL) EPTCS* **318**, 2020, pp. 196–21 [arXiv: 1905.09055](#)

- [12] A. Holevo, *Probabilistic and Statistical Aspects of Quantum Theory*, North-Holland 1982 (reissued by Edizioni della Normale, 2011)
- [13] J. Harding and A. Wilce, Probabilistic theories and order-unit spaces, to appear
- [14] N. Johnson and D. Yau, *2-Dimensional Categories*, <https://arxiv.org/pdf/2002.06055> (2020)
- [15] M. Leifer, Is the quantum state real? An extended review of ψ -ontology theorems, *Quanta* **3** 2014;
- [16] T. Maudlin, What Bell did, *J. Phys. A* **47**, 2014 [arXiv:1408.1826](https://arxiv.org/abs/1408.1826)
- [17] S. Popescu and D. Rohrlich, Quantum nonlocality as an axiom, *Foundations of Physics* **24** (1994)
- [18] B. Simon, *Convexity: an analytic viewpoint*, Cambridge, 2011
- [19] R. Spekkens, Contextuality for preparations, transformations, and unsharp measurements, *Phys. Rev. A* **71** (2005)
- [20] R. Werner, Comment on Maudlin’s Paper “What Bell did”, *J. Phys. A* **47**, 2014; see also R. Werner, What Maudlin replied to, <https://arxiv.org/abs/1411.2120> (2014)
- [21] A. Wilce, Quantum logic and probability theory, *Stanford Encyclopedia of Philosophy* 2002
- [22] A. Wilce, Test spaces, in D. Gabbay, K. Engesser and D. Lehman (eds.) *Handbook of Quantum Logic: Quantum Structures*, North-Holland, 2006 [check]
- [23] A. Wilce, A shortcut from categorical quantum theory to convex operational theories, in Bob Coecke and Aleks Kissinger (eds.): 14th International Conference on Quantum Physics and Logic (QPL) EPTCS **266**, 2018, pp. 222–236
- [24] Coarse-graining and compounding as monads, [arXiv:2410.08818](https://arxiv.org/abs/2410.08818) (2024)

- [25] A. Wilce, *Generalized Probability Theory: notes for a short course* [arXiv:2501.00718](#) (2025)
- [26] R. Wright, The structure of projection-valued states: A generalization of Wigner's theorem, *Int. J. Theor. Physics* **16W**, 1977