

Article

A Royal Road to Quantum Theory (or Thereabouts)

Alexander Wilce

Department of Mathematics, Susquehanna University, Selinsgrove, PA 17870, USA; wilce@susqu.edu

Received: 15 January 2018; Accepted: 19 March 2018; Published: 26 March 2018



Abstract: This paper fails to derive quantum mechanics from a few simple postulates. However, it gets very close, and does so without much exertion. More precisely, I obtain a representation of finite-dimensional probabilistic systems in terms of Euclidean Jordan algebras, in a strikingly easy way, from simple assumptions. This provides a framework within which real, complex and quaternionic QM can play happily together and allows some (but not too much) room for more exotic alternatives. (This is a leisurely summary, based on recent lectures, of material from the papers arXiv:1206.2897 and arXiv:1507.06278, the latter joint work with Howard Barnum and Matthew Graydon. Some further ideas are also explored, developing the connection between conjugate systems and the possibility of forming stable measurement records and making connections between this approach and the categorical approach to quantum theory.)

Keywords: reconstruction of quantum mechanics; conjugate systems; Jordan algebras

1. Introduction and Overview

Whatever else it may be, Quantum mechanics (QM) is a machine for making probabilistic predictions about the results of measurements. To this extent, QM is, at least in part, about information. Over the last decade or so, it has become clear that the formal apparatus of quantum theory, at least in finite dimensions, can be recovered from constraints on how physical systems store and process information. To this extent, finite-dimensional QM is just about information.

The broad idea of regarding QM in this way, and of attempting to derive its mathematical structure from simple operational or probabilistic axioms, is not new. Efforts in this direction go back at least to the work of von Neumann [1], and include also attempts by Schwinger [2], Mackey [3], Ludwig [4], Piron [5], and many others. However, the consensus is that these were not entirely successful: partly because the results they achieved (e.g., Piron's well-known representation theorem) did not rule out certain rather exotic alternatives to QM, but mostly because the axioms deployed seem, in retrospect, to lack sufficient physical or operational motivation.

More recently, with inspiration from quantum information theory, attention has focused on finite-dimensional systems, where the going is a bit easier. Just as importantly, quantum information theory prompts us to treat properties of composite systems as fundamental, where earlier work focused largely on systems in isolation (a recent exception to this trend is the paper [6] of Barnum, Müller and Ududec). These shifts of emphasis are illustrated by the work of Hardy [7], who presented five simple, broadly information-theoretic postulates governing the states and measurements associated with a physical system, determining a very restricted set of possible theories, parametrized by a positive integer r , with finite-dimensional quantum and classical probability theory corresponding to $r = 1$ and $r = 2$. Following this lead, several papers, notably [8–10], have derived finite-dimensional QM from various packages of axioms governing the information-carrying and information-processing capacity of finite-dimensional systems.

Problems with existing approaches. These recent reconstructive efforts suffer from two related problems. First, they make use of assumptions that seem too strong. Secondly, in trying to derive exactly complex, finite-dimensional quantum theory, they derive too much.

- All of the cited papers assume local tomography. This is the doctrine that the state of a bipartite composite system is entirely determined by the joint probabilities it assigns to outcomes of measurements on the two subsystems. This rules out both real and quaternionic QM, both of which are legitimate quantum theories [11].
- These papers also all make some version of a uniformity assumption: that all systems having the same information-carrying capacity are isomorphic, or that all systems are composed, in a uniform way, from “bits” of a uniform type. Here, “information carrying capacity” means essentially the maximum number of states that can be distinguished from one another with probability one by a single measurement. A bit is a system for which this number is two. This rules out systems involving superselection rules, i.e., those that admit both real and classical degrees of freedom (for example, the quantum system corresponding to $M_2(\mathbb{C}) \oplus M_2(\mathbb{C})$, corresponding to a classical choice between one of two qubits, has the same information-carrying capacity as a single, four-level quantum system). More seriously, it rules out any theory that includes, e.g., real and complex, or real and quaternionic systems, as the state spaces of the bits of these theories have different dimensions. As I will discuss below, one can indeed construct mathematically-reasonable theories that embrace finite-dimensional quantum systems of all three types.
- Another shortcoming, not related to the exclusion of real and quaternionic QM, is the technical assumption (explicit in [10] for bits) that all positive affine functionals on the state space taking values between zero and one correspond to physically-accessible “effects”, i.e., possible measurement results. From an operational point of view, this principle (called the “no-restriction hypothesis” in [12]) seems to call for further motivation.

Another approach. In these notes, I am going to describe an alternative approach that avoids these difficulties. This begins by associating with every physical system a convex set of states and a distinguished set of basic measurements (or experiments) that can be made on the system. We then isolate two striking features shared by classical and quantum probabilistic systems. The first is the possibility of finding a joint state that perfectly correlates a system A with an isomorphic system \bar{A} (call it a conjugate system) in the sense that every basic measurement on A is perfectly correlated with the corresponding measurement on \bar{A} . In finite-dimensional QM, where A is represented by a finite-dimensional Hilbert space \mathcal{H} , \bar{A} , corresponds to the conjugate Hilbert space $\bar{\mathcal{H}}$, and the perfectly-correlating state is the maximally-entangled “EPR” state on $\mathcal{H} \otimes \bar{\mathcal{H}}$.

The second feature is the existence of what I call filters associated with each basic measurement. These are processes that independently attenuate the “response” of each outcome of the measurement by some specified factor. Such a process will generally not preserve the normalization of states, but up to a constant factor, in both classical and quantum theory, one can prepare any desired state by applying a suitable filter to the maximally-mixed state. Moreover, when the target state is not singular (that is, when it does not assign probability zero to any nonzero measurement outcome), one can reverse the filtering process, in the sense that it can be undone by another process with positive probability.

The upshot is that all probabilistic systems having conjugates and a sufficiently lavish supply of (probabilistically) reversible filters can be represented by formally real Jordan algebras, a class of structures that includes real, complex and quaternionic quantum systems, and just two further well-studied additional possibilities, which I will review below.

In addition to leaving room for real and quaternionic quantum mechanics (which I take to be a virtue), this approach has another advantage: it is much easier! The assumptions involved are few and easily stated, and the proof of the main technical result (Lemma 1 in Section 4) is short and straightforward. By contrast, the mathematical developments in the papers listed above are significantly more difficult and ultimately lean on the (even more difficult) classification of compact

groups acting on spheres. My approach, too, leans on a received result, but one that is relatively accessible. This is the Koecher–Vinberg theorem, which characterizes formally real, or Euclidean, Jordan algebras in terms of ordered real vector spaces with homogeneous, self-dual cones. A short and non-taxing proof of this classical result can be found in [13].

These ideas were developed in [14–16] and especially [17], of which this paper is, to an extent, a summary. However, the presentation here is slightly different, and some additional ideas are also explored. In particular, I have spelled out in more detail the connection between conjugate systems and measurement records, only alluded to in the earlier paper. I also link this approach to the categorical approach to quantum theory due to Abramsky, Coecke and others [18], along the way briefly discussing recent work with Howard Barnum and Matthew Graydon [19] on the construction of probabilistic theories in which real, complex and quaternionic quantum systems coexist. Finally, Appendix B presents a uniqueness result for spectral decompositions of states, which may find further application.

A bit of background. At this point, I had better pause to explain some terms. A *Jordan algebra* is a real commutative algebra (a real vector space \mathbf{E} with a commutative bilinear multiplication $a, b \mapsto a \bullet b$) having a multiplicative unit u and satisfying the Jordan identity: $a^2 \bullet (a \bullet b) = a \bullet (a^2 \bullet b)$, for all $a, b, c \in \mathbf{E}$, where $a^2 = a \bullet a$. A Jordan algebra is *formally real* if sums of squares of nonzero elements are always nonzero. The basic, and motivating, example is the space $\mathcal{L}_{\text{sa}}(\mathcal{H})$ of self-adjoint operators on a complex Hilbert space, with the Jordan product given by $a \bullet b = \frac{1}{2}(ab + ba)$. Note that here, $a \bullet a = aa$, so the notation a^2 is unambiguous. To see that $\mathcal{L}_{\text{sa}}(\mathcal{H})$ is formally real, just note that a^2 is always a positive operator.

If \mathcal{H} is finite dimensional, $\mathcal{L}_{\text{sa}}(\mathcal{H})$ carries a natural inner product, namely $\langle a, b \rangle = \text{Tr}(ab)$. This plays well with the Jordan product: $\langle a \bullet b, c \rangle = \langle b, a \bullet c \rangle$ for all $a, b, c \in \mathcal{L}_{\text{sa}}(\mathcal{H})$. More generally, a finite-dimensional Jordan algebra equipped with an inner product having this property is said to be *Euclidean*. For finite-dimensional Jordan algebras, being formally real and being Euclidean are equivalent [13]. In what follows, I will abbreviate “Euclidean Jordan algebra” to EJA.

Jordan algebras were originally proposed, with what now looks like slightly thin motivation, by P. Jordan [20]: if a and b are quantum-mechanical observables, represented by $a, b \in \mathcal{L}_{\text{sa}}(\mathcal{H})$, then while $a + b$ is again self-adjoint, ab and ba are not, unless a and b commute; however, their average, $a \bullet b$, is self-adjoint and, thus, represents another observable. Almost immediately, Jordan, von Neumann and Wigner showed [21] that all formally real Jordan algebras are direct sums of simple such algebras, with the latter falling into just five classes, parametrized by positive integers n : the self-adjoint parts, $M_n(\mathbb{F})_{\text{sa}}$, of matrix algebras $M_n(\mathbb{F})$, where $\mathbb{F} = \mathbb{R}, \mathbb{C}$ or \mathbb{H} (the quaternions) or, for $n = 3$, over \mathbb{O} (the octonions); and also what are called spin factors V_n (closely related to Clifford algebras). There is some overlap: $V_2 \simeq M_2(\mathbb{R})$, $V_3 \simeq M_2(\mathbb{C})$ and $V_5 \simeq M_2(\mathbb{H})$. In all but one case, one can show that a simple Jordan algebra is a Jordan subalgebra of $M_n(\mathbb{C})$ for suitable n . The exceptional Jordan algebra, $M_3(\mathbb{O})_{\text{sa}}$, admits no such representation.

Besides this classification theorem, there is only one other important fact about Euclidean Jordan algebras that is needed for what follows. This is the Koecher–Vinberg (KV) theorem alluded to above. Recall that an ordered vector space is a real vector space, call it \mathbf{E} , spanned by a distinguished convex cone \mathbf{E}_+ having its vertex at the origin. Such a cone induces a translation-invariant partial order on \mathbf{E} , namely $a \leq b$ iff $b - a \in \mathbf{E}_+$. As an example, the space $\mathcal{L}_{\text{sa}}(\mathcal{H})$ is ordered by the cone of positive operators. More generally, any EJA is an ordered vector space, with positive cone $\mathbf{E}_+ := \{a^2 | a \in A\}$. This cone has two special features: first, it is *homogeneous*, i.e., for any points a, b in the interior of \mathbf{E}_+ , there exists an automorphism of the cone (a linear isomorphism $\mathbf{E} \rightarrow \mathbf{E}$, taking \mathbf{E}_+ onto itself) that maps a to b . In other words, the group of automorphisms of the cone acts transitively on the cone’s interior. The other special property is that \mathbf{E}_+ is *self-dual*. This means that \mathbf{E} carries an inner product (in fact, the given one making \mathbf{E} Euclidean) such that $a \in \mathbf{E}_+$ iff $\langle a, b \rangle \geq 0$ for all $b \in \mathbf{E}_+$.

An *order unit* in an ordered vector space \mathbf{E} is an element $u \in \mathbf{E}_+$ such that, for all $a \in \mathbf{E}$, there exists some $n \in \mathbb{N}$ with $a \leq nu$. In finite dimensions, this is equivalent to u 's belonging to the interior of the cone \mathbf{E}_+ [22]. In the following, by a Euclidean order-unit space, I mean an ordered vector space \mathbf{E} equipped with an inner product $\langle \cdot, \cdot \rangle$ with $\langle a, b \rangle \geq 0$ for all $a, b \in \mathbf{E}_+$, and a distinguished order-unit u . I will say that such a space \mathbf{E} is HSD iff \mathbf{E}_+ is homogeneous, and also self-dual with respect to the given inner product.

Theorem 1 (Koecher 1958; Vinberg 1961). *Let \mathbf{E} be a finite-dimensional euclidean order-unit space. If \mathbf{E} is HSD, then there exists a unique product \cdot with respect to which \mathbf{E} (with its given inner product) is a euclidean Jordan algebra, u is the Jordan unit, and \mathbf{E}_+ is the cone of squares.*

It seems, then, that if we can motivate a representation of physical systems in terms of HSD order-unit spaces, we will have “reconstructed” what with a little license we might call finite-dimensional Jordan-quantum mechanics. In view of the classification theorem glossed above, this gets us into the neighborhood of orthodox QM, but still leaves open the possibility of taking real and quaternionic quantum systems seriously. (It also leaves the door open to two possibly unwanted guests, namely spin factors and the exceptional Jordan algebra. I will discuss below some constraints that at least bar the latter.)

Some notational conventions. My notation is mostly consistent with the following conventions (more standard in the mathematics than the physics literature, but in places slightly excentric relative to either). Capital Roman letters A, B, C serve as labels for systems. $M_n(\mathbb{F})$ stands for the set of $n \times n$ matrices over $\mathbb{F} = \mathbb{R}$ or \mathbb{H} ; $M_n(\mathbb{F})_{\text{sa}}$ is the set of self-adjoint such matrices. Vectors in a Hilbert space \mathcal{H} are denoted by little Roman letters x, y, z from the end of the alphabet. Operators on \mathcal{H} will usually be denoted by little Roman letters a, b, c, \dots from the beginning of the alphabet. Roman letters t, s typically stand for real numbers. The space of all linear operators on \mathcal{H} is denoted $\mathcal{L}(\mathcal{H})$; as already indicated above, $\mathcal{L}_{\text{sa}}(\mathcal{H})$ is the (real) vector space of self-adjoint operators on \mathcal{H} .

As above, the conjugate Hilbert space is denoted $\overline{\mathcal{H}}$. I will write \bar{x} for the vectors in $\overline{\mathcal{H}}$ corresponding to $x \in \mathcal{H}$. From a certain point of view, this is the same vector; the bar serves to remind us that $\overline{cx} = \bar{c}\bar{x}$ for scalars $c \in \mathbb{C}$. Alternatively, one can regard $\overline{\mathcal{H}}$ as the space of “bra” vectors $\langle x|$ corresponding to the “kets” $|x\rangle$ in \mathcal{H} , i.e., as the dual space of \mathcal{H} .

The inner product of $x, y \in \mathcal{H}$ is written as $\langle x, y \rangle$ and is linear in the first argument (if you like: $\langle x, y \rangle = \langle y|x\rangle$ in Dirac notation). The inner product on $\overline{\mathcal{H}}$ is then $\langle \bar{x}, \bar{y} \rangle = \langle y, x \rangle$. The rank-one projection operator associated with a unit vector $x \in \mathcal{H}$ is p_x . Thus, $p_x(y) = \langle y, x \rangle x$. I denote functionals on $\mathcal{L}_{\text{sa}}(\mathcal{H})$ by little Greek letters, e.g., α, β, \dots , and operators on $\mathcal{L}_{\text{sa}}(\mathcal{H})$ by capital Greek letters, e.g., Φ . Two exceptions to this scheme: a generic density operator on \mathcal{H} is denoted by the capital Roman letter W , and a certain special unit vector in $\mathcal{H} \otimes \overline{\mathcal{H}}$ is denoted by the capital Greek letter Ψ . With luck, context will help keep things straight.

2. Homogeneity and Self-Duality in Quantum Theory

Why should a probabilistic physical system be represented by a Euclidean order-unit space that is either homogeneous or self-dual? One place to start hunting for an answer might be to look at standard quantum probability theory, to see if we can isolate, in operational or probabilistic terms, what makes this self-dual and homogeneous.

Correlation and self-duality. Let \mathcal{H} be a finite-dimensional complex Hilbert space, representing some finite-dimensional quantum system. The system's states are represented by density operators, i.e., positive trace-one operators $W \in \mathcal{L}_{\text{sa}}(\mathcal{H})$; possible measurement-outcomes are represented by effects, i.e., positive operators $a \in \mathcal{L}_{\text{sa}}(\mathcal{H})$ with $a \leq 1$. The Born rule specifies the probability of

observing effect a in state W as $\text{Tr}(Wa)$. If W is a pure state, i.e., $W = p_v$ where v is a unit vector in \mathcal{H} , then $\text{Tr}(Wa) = \langle av, v \rangle$; by the same token, if $a = p_x$, then $\text{Tr}(Wa) = \langle Wx, x \rangle$.

For $a, b \in \mathcal{L}_{\text{sa}}(\mathcal{H})$, let $\langle a, b \rangle := \text{Tr}(ab)$. This is an inner product. By the spectral theorem, $\text{Tr}(ab) \geq 0$ for all $b \in \mathcal{L}_h(\mathcal{H})_+$ iff $\text{Tr}(ap_x) \geq 0$ for all unit vectors x . However, $\text{Tr}(ap_x) = \langle ax, x \rangle$. So $\text{Tr}(ab) \geq 0$ for all $b \in \mathcal{L}_h(\mathcal{H})_+$ iff $a \in \mathcal{L}_h(\mathcal{H})_+$, i.e., the trace inner product is self-dualizing. However, this now leaves us with the following:

Question: *What does the trace inner product represent, operationally or probabilistically?*

Let $\overline{\mathcal{H}}$ be the conjugate Hilbert space to \mathcal{H} . Suppose \mathcal{H} has dimension n . Any unit vector Ψ in $\mathcal{H} \otimes \overline{\mathcal{H}}$ gives rise to a joint probability assignment to effects a on \mathcal{H} and \bar{b} on $\overline{\mathcal{H}}$, namely $\langle (a \otimes \bar{b})\Psi, \Psi \rangle$. Consider the EPR state for $\mathcal{H} \otimes \overline{\mathcal{H}}$ defined by the unit vector:

$$\Psi = \frac{1}{\sqrt{n}} \sum_{x \in E} x \otimes \bar{x} \in \mathcal{H} \otimes \overline{\mathcal{H}},$$

where E is any orthonormal basis for \mathcal{H} . A straightforward computation shows that the joint probability of observing a and b in the state Ψ is:

$$\langle (a \otimes \bar{b})\Psi, \Psi \rangle = \frac{1}{n} \text{Tr}(ab).$$

In other words, the normalized trace inner product just is the joint probability function determined by the pure state vector Ψ !

As a consequence, the state represented by Ψ has a very strong correlational property: if x, y are two orthogonal unit vectors with corresponding rank-one projections p_x and p_y , we have $p_x p_y = 0$, so $\langle (p_x \otimes \bar{p}_y)\Psi, \Psi \rangle = 0$. On the other hand, $\langle (p_x \otimes \bar{p}_x)\Psi, \Psi \rangle = \frac{1}{n} \text{Tr}(p_x) = \frac{1}{n}$. Hence, Ψ perfectly, and uniformly, correlates every basic measurement (orthonormal basis) of \mathcal{H} with its counterpart in $\overline{\mathcal{H}}$.

Filters and homogeneity. Next, let us see why the cone $\mathcal{L}_h(\mathcal{H})_+$ is homogeneous. Recall that this means that any state in the interior of the cone (here, any non-singular density operator) can be obtained from any other by an automorphism of the cone. However, in fact, something better is true: this order-automorphism can be chosen to represent a probabilistically-reversible physical process, i.e., an invertible CP mapping with a CP inverse.

To see how this works, suppose W is a positive operator on \mathcal{H} . Consider the pure CP mapping $\Phi_W : \mathcal{L}_{\text{sa}}(\mathcal{H}) \rightarrow \mathcal{L}_{\text{sa}}(\mathcal{H})$ given by:

$$\Phi_W(a) = W^{1/2} a W^{1/2}.$$

Then, $\Phi_W(\mathbf{1}) = W$. If W is nonsingular, so is $W^{1/2}$, so Φ_W is invertible, with inverse $\Phi_W^{-1} = \Phi_{W^{-1}}$, again a pure CP mapping. Now, given another nonsingular density operator M , we can get from W to M by applying $\Phi_M \circ \Phi_{W^{-1}}$.

All well and good, but we are still left with the following:

Question: *What does the mapping Φ_W represent, physically?*

To answer this, suppose W is a density operator, with spectral expansion $W = \sum_{x \in E} t_x p_x$. Here, E is an orthonormal basis for \mathcal{H} diagonalizing W , and t_x is the eigenvalue of W corresponding to $x \in E$. Then, for each vector $x \in E$,

$$\Phi_W(p_x) = t_x p_x$$

where p_x is the projection operator associated with x . We can understand this to mean that Φ_W acts as a filter on the test E : the response of each outcome $x \in E$ is attenuated by a factor $0 \leq t_x \leq 1$ (my usage

here is slightly non-standard, in that I allow filters that “pass” the system with a probability strictly between zero and one). Thus, if M is another density operator on \mathcal{H} , representing some state of the corresponding system, then the probability of obtaining outcome x after preparing the system in state M and applying the process Φ is t_x times the probability of x in state M . In detail: suppose p_x is the rank-one projection operator associated with x , and note that $W^{1/2}p_x = p_xW^{1/2} = t_x^{1/2}p_x$. Thus,

$$\begin{aligned} \text{Tr}(\Phi_W(M)p_x) &= \text{Tr}(W^{1/2}MW^{1/2}p_x) = \text{Tr}(W^{1/2}Mt_x^{1/2}p_x) = \text{Tr}(t_x^{1/2}p_xW^{1/2}M) \\ &= \text{Tr}(t_x p_x M) = t_x \text{Tr}(Mp_x). \end{aligned}$$

If we think of the basis E as representing a set of alternative channels plus detectors, as in the figure below, we can add a classical filter attenuating the response of one of the detectors (say, x) by a fraction t_x . What the computation above tells us is that we can achieve the same result by applying a suitable CP map to the system’s state. Moreover, this can be done independently for each outcome of E . In Figure 1, this is illustrated for a three-level quantum system: $E = \{x, y, z\}$ is an orthonormal basis, representing three possible outcomes of a Stern–Gerlach-like experiment; the filter Φ acts on the system’s state in such a way that the probability of outcome x is attenuated by a factor of $t_x = 1/2$, while outcomes y and z are unaffected. Returning to the general situation, if we apply a filter Φ_W to the maximally-mixed state $\frac{1}{n}\mathbf{1}$, we obtain $\frac{1}{n}W$. Thus, we can prepare W , up to normalization, by applying the filter Φ_W to the maximally mixed state.

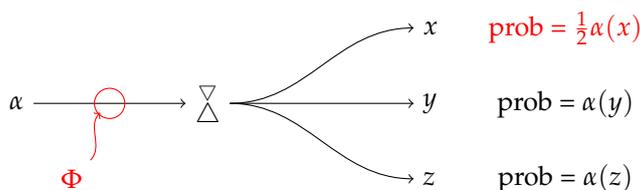


Figure 1. Φ attenuates x ’s sensitivity by 1/2.

Filters are symmetric. Here is a final observation, linking these last two: the filter Φ_W is symmetric with respect to the uniformly-correlating “EPR” state Ψ , in the sense that:

$$\langle (\Phi_W(a) \otimes \bar{b})\Psi, \Psi \rangle = \langle (a \otimes \bar{\Phi}_W(b))\Psi, \Psi \rangle$$

for all effects $a, b \in \mathcal{L}_{\text{sa}}(\mathcal{H})_+$. Remarkably, this is all that is needed to recover the Jordan structure of finite-dimensional quantum theory: the existence of a conjugate system, with a uniformly-correlating joint state, plus the possibility of preparing non-singular states by means of filters that are symmetric with respect to this state, and doing so reversibly when the state is nonsingular.

In a very rough outline, the argument is that states preparable (up to normalization) by symmetric filters have spectral decompositions, and the existence of spectral decompositions makes the uniformly-correlating joint state a self-dualizing inner product. However, to spell this out in a precise way, I need a general mathematical framework for discussing states, effects and processes in abstraction from quantum theory. The next section reviews the necessary apparatus.

3. General Probabilistic Theories

A characteristic feature of quantum mechanics is the existence of incompatible, or non-measurable, observables. This suggests the following simple, but very fruitful, notion:

Definition 1. A test space is a collection \mathcal{M} of non-empty sets E, F, \dots , each representing the outcome-set of some measurement, experiment, or test. At the outset, one makes no special assumptions about the combinatorial structure of \mathcal{M} . In particular, distinct tests are permitted to overlap. Let $X := \bigcup \mathcal{M}$ denote the set of all

outcomes of all tests in \mathcal{M} : a probability weight on \mathcal{M} is a function $\alpha : X \rightarrow [0, 1]$ such that $\sum_{x \in E} \alpha(x) = 1$ for every $E \in \mathcal{M}$.

Test spaces were introduced and studied by D. J. Foulis and C. H. Randall in a long series of papers beginning around 1970. The original term for a test was an operation, which has the advantage of signaling that the concept has wider applicability than simply reading a number off a meter: anything an agent can do that leads to a well-defined, exhaustive set of mutually-exclusive outcomes defines an operation. Accordingly, test spaces were originally called “manuals of operations”.

It can happen that a test space admits no probability weights at all. However, to serve as a model of a real family of experiments associated with an actual physical system, a test space should obviously carry a lavish supply of such weights. One might want to single out some of these as describing physically (or otherwise) possible states of the system. This suggests the following:

Definition 2. A probabilistic model is a pair $A = (\mathcal{M}, \Omega)$, where \mathcal{M} is a test space and Ω is some designated convex set of probability weights, called the states of the model.

The definition is deliberately spare. Nothing prohibits us from adding further structure (a group of symmetries, say, or a topology on the space of outcomes). However, no such additional structure is needed for the results I will discuss below. I will write $\mathcal{M}(A)$, $X(A)$ and $\Omega(A)$ for the test space, associated outcome space and state space of a model A . The convexity assumption on $\Omega(A)$ is intended to capture the possibility of forming mixtures of states. To allow the modest idealization of taking outcome-wise limits of states to be states, I will also assume that $\Omega(A)$ is closed as a subset of $[0, 1]^{X(A)}$ (in its product topology). This makes $\Omega(A)$ compact and, so, guarantees the existence of pure states, that is, extreme points of $\Omega(A)$. If $\Omega(A)$ is the set of all probability weights on $\mathcal{M}(A)$, I will say that A has a full state space.

Two bits. Here is a simple, but instructive illustration of these notions. Consider a test space $\mathcal{M} = \{\{x, x'\}, \{y, y'\}\}$. Here, we have two tests, each with two outcomes. We are permitted to perform either test, but not both at once. A probability weight is determined by the values it assigns to x and to y , and since the sets $\{x, x'\}$ and $\{y, y'\}$ are disjoint, these values are independent. Thus, geometrically, the space of all probability weights is the unit square in \mathbb{R}^2 (Figure 2a, below). To construct a probabilistic model, we can choose any closed, convex subset of the square for Ω . For instance, we might let Ω be the convex hull of the four probability weights $\delta_x, \delta_{x'}, \delta_y$ and $\delta_{y'}$ corresponding to the midpoints of the four sides of the square, as in Figure 2b, that is,

$$\delta_x(x) = 1, \delta_x(x') = 0, \delta_x(y) = \delta_x(y') = 1/2,$$

$$\delta_{x'}(x) = 0, \delta_{x'}(x') = 1, \delta_{x'}(y) = \delta_{x'}(y') = 1/2,$$

and similarly for δ_y and $\delta_{y'}$.

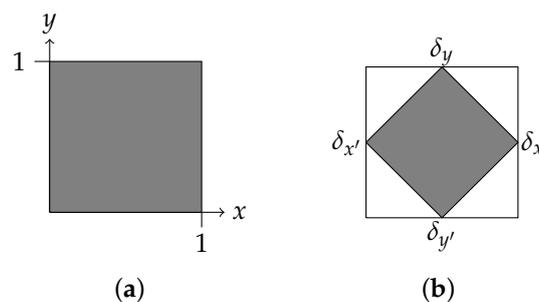


Figure 2. The state spaces of two bits. (a) The square bit; (b) The diamond bit.

The model of Figure 2a, in which we take Ω to be the entire set of probability weights on $\mathcal{M} = \{\{x, x'\}, \{y, y'\}\}$, is sometimes called the *square bit*. I will call the model of Figure 2b the *diamond bit*.

Classical, quantum and Jordan models. If E is a finite set, the corresponding *classical model* is $A(E) = (\{E\}, \Delta(E))$ where $\Delta(E)$ is the simplex of probability weights on E . If \mathcal{H} is a finite-dimensional complex Hilbert space, let $\mathcal{M}(\mathcal{H})$ denote the set of orthonormal bases of \mathcal{H} : then $X = \bigcup \mathcal{M}(\mathcal{H})$ is the unit sphere of \mathcal{H} , and any density operator W on \mathcal{H} defines a probability weight α_W , given by $\alpha_W(x) = \langle Wx, x \rangle$ for all $x \in X$. Letting $\Omega(\mathcal{H})$ denote the set of states of this form, we obtain the *quantum model*, $A(\mathcal{H}) = (\mathcal{M}(\mathcal{H}), \Omega(\mathcal{H}))$, associated with \mathcal{H} (Gleason's theorem tells us that $A(\mathcal{H})$ has a full state space for $\dim(\mathcal{H}) > 2$, but we will not need this fact).

More generally, every Euclidean Jordan algebra \mathbf{E} gives rise to a probabilistic model as follows. A minimal or primitive idempotent of \mathbf{E} is an element $p \in \mathbf{E}$ with $p^2 = p$ and, for $q = q^2 < p$, $q = 0$. A Jordan frame is a maximal pairwise orthogonal set of primitive idempotents. Let $X(\mathbf{E})$ be the set of primitive idempotents; let $\mathcal{M}(\mathbf{E})$ be the set of Jordan frames; and let $\Omega(\mathbf{E})$ be the set of probability weights of the form $\alpha(p) = \langle a, p \rangle$ where $a \in \mathbf{E}_+$ with $\langle a, u \rangle = 1$. These data define the *Jordan model* $A(\mathbf{E})$ associated with \mathbf{E} . In the case where $\mathbf{E} = \mathcal{L}_h(\mathcal{H})$ for a finite-dimensional Hilbert space \mathcal{H} , this almost gives us back the quantum model $A(\mathcal{H})$: the difference is that we replace unit vectors by their associated projection operators, thus conflating outcomes that differ only by a phase.

Sharp models. Jordan models enjoy many special features that the generic probabilistic model lacks. I want to take a moment to discuss one such feature, which will be important below.

Definition 3. A model A is *unital* iff, for every outcome $x \in X(A)$, there exists a state $\alpha \in \Omega(A)$ with $\alpha(x) = 1$, and *sharp* if this state is unique (from which it follows easily that it must be pure). If A is sharp, I will write δ_x for the unique state making $x \in X(A)$ certain.

If A is sharp, then there is a sense in which each test $E \in \mathcal{M}(A)$ is maximally informative: if we are certain which outcome $x \in E$ will occur, then we know the system's state exactly, as there is only one state in which x has probability 1.

Classical and quantum models are obviously sharp. More generally, every Jordan model is sharp. To see this, note first that every state α on a Euclidean Jordan algebra \mathbf{E} has the form $\alpha(x) = \langle a, x \rangle$ where $a \in \mathbf{E}_+$ with $\langle a, u \rangle = 1$ and where $\langle \cdot, \cdot \rangle$ is the given inner product on \mathbf{E} , normalized so that $\|x\| = 1$ for all primitive idempotents (equivalently, so that $\|u\| = n$, the rank of \mathbf{E}). The spectral theorem for EJAs [13] shows that $a = \sum_{p \in E} t_p p$ where E is a Jordan frame and the coefficients t_p are non-negative and sum to one (since $\langle a, u \rangle = 1$). If $\langle a, x \rangle = 1$, then $\sum_{p \in E} t_p \langle p, x \rangle = 1$ implies that, for every $p \in E$ with $t_p > 0$, $\langle p, x \rangle = 1$. However, $\|p\| = \|x\| = 1$, so this implies that $\langle p, x \rangle = \|p\| \|x\|$, which in turn implies that $p = x$.

In general, a probabilistic model need not even be unital, much less sharp. On the other hand, given a unital model A , it is often possible to construct a sharp model by suitably restricting the state space. This is illustrated in Figure 2b above: the full state space of the square bit is unital, but far from sharp; however, by restricting the state space to the convex hull of the barycenters of the faces, we obtain a sharp model. This is possible whenever A is unital and carries a group of symmetries acting transitively on the outcome-set $X(A)$. For details, see Appendix A. The point here is that sharpness is not, by itself, a very stringent condition: since we should expect to find highly symmetric, unital models represented abundantly "in nature", we can also expect to encounter an abundance of systems represented by sharp models.

The spaces $\mathbf{V}(A), \mathbf{V}^*(A)$. Any probabilistic model gives rise to a pair of ordered vector spaces in a canonical way. These will be essential in the development below, so I am going to go into a bit of detail here.

Definition 4. Let A be any probabilistic model. Let $\mathbf{V}(A)$ be the span of the state space $\Omega(A)$ in $\mathbb{R}^{X(A)}$, ordered by the cone $\mathbf{V}(A)_+$ consisting of non-negative multiples of states, i.e.,

$$\mathbf{V}(A)_+ = \{t\alpha \mid \alpha \in \Omega(A), t \geq 0\}.$$

Call the model A finite-dimensional iff $\mathbf{V}(A)$ is finite-dimensional. From now on, I assume that all models are finite-dimensional.

Let $\mathbf{V}^*(A)$ denote the dual space of $\mathbf{V}(A)$, ordered by the dual cone of positive linear functionals, i.e., functionals f with $f(\alpha) \geq 0$ for all $\alpha \in \mathbf{V}(A)_+$. Any measurement-outcome $x \in X(A)$ yields an evaluation functional $\hat{x} \in \mathbf{V}^*(A)$, given by $\hat{x}(\alpha) = \alpha(x)$ for all $\alpha \in \mathbf{V}(A)$. More generally, an effect is a positive linear functional $f \in \mathbf{V}^*(A)$ with $0 \leq f(\alpha) \leq 1$ for every state $\alpha \in \Omega(A)$. The functionals \hat{x} are effects. One can understand an arbitrary effect a to represent a mathematically possible measurement outcome, having probability $a(\alpha)$ in state α . I stress the adjective mathematically because, a priori, there is no guarantee that every effect will correspond to a physically-realizable measurement outcome. In fact, at this stage, I make no assumption at all about what, apart from the tests $E \in \mathcal{M}(A)$, is or is not physically realizable. (Later, it will follow from further assumptions that every element of $\mathbf{V}^*(A)$ represents a random variable associated with some $E \in \mathcal{M}(A)$ and is, therefore, operationally meaningful. However, this will be a theorem, not an assumption.)

The *unit effect* is the functional $u_A := \sum_{x \in E} \hat{x}$, where E is any element of $\mathcal{M}(A)$. This takes the constant value of one on $\Omega(A)$, and, thus, represents a trivial measurement outcome that occurs with probability one in every state. This is an order unit for $\mathbf{V}^*(A)$ (to see this, let $a \in \mathbf{V}(A)^*$, and let N be the maximum value of $|a(\alpha)|$ for $\alpha \in \Omega(A)$, remembering that the latter is compact: then $a \leq Nu$).

For both classical and quantum models, the ordered vector spaces $\mathbf{V}^*(A)$ and $\mathbf{V}(A)$ are naturally isomorphic. If $A(E)$ is the classical model associated with a finite set E , both are isomorphic to the space \mathbb{R}^E of all real-valued functions on E , ordered pointwise. If $A = A(\mathcal{H})$ is the quantum model associated with a finite-dimensional Hilbert space \mathcal{H} , $\mathbf{V}(A)$ and $\mathbf{V}^*(A)$ are both naturally isomorphic to the space $\mathcal{L}_h(\mathcal{H})$ of Hermitian operators on \mathcal{H} , ordered by its usual cone of positive semi-definite operators. More generally, if \mathbf{E} is a Euclidean Jordan algebra and $A = A(\mathbf{E})$ is the corresponding Jordan model, then $\mathbf{V}(A) \simeq \mathbf{E} \simeq \mathbf{V}^*(A)$, with \mathbf{E} ordered as usual, i.e., by its cone of squares. The first of these isomorphisms is due to the definition of the model $A(\mathbf{E})$ and the second to \mathbf{E} 's self-duality.

The space $\mathbf{E}(A)$. It is going to be technically useful to introduce a third ordered vector space, which I will denote by $\mathbf{E}(A)$. This is the span of the evaluation-effects \hat{x} , associated with measurement outcomes $x \in X(A)$, in $\mathbf{V}^*(A)$, ordered by the cone:

$$\mathbf{E}(A)_+ := \left\{ \sum_i t_i \hat{x}_i \mid t_i \geq 0 \right\}.$$

That is, $\mathbf{E}(A)_+$ is the set of linear combinations of effects \hat{x} having non-negative coefficients. It is important to note that this is, in general, a proper sub-cone of $\mathbf{V}(A)_+^*$. To see this, we can revisit the example of the “diamond bit” of Figure 2b. Letting x and y be the outcomes corresponding to the right face and the top face of the larger (full) state space pictured below in Figure 3a, consider the functional $f := \hat{x} + \hat{y} - \frac{1}{2}u$. This takes positive values on the smaller state space of the diamond bit, but is negative on, for example, the state γ corresponding to the lower-left corner of the full state space (see Figure 3b). Thus, $f \in \mathbf{V}(A)_+$, but $f \notin \mathbf{E}(A)_+$.

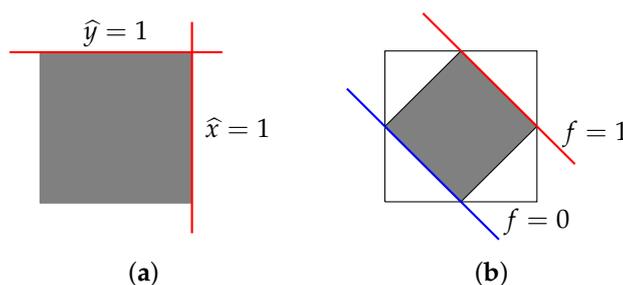


Figure 3. (a) Two outcome-effects for the square bit; (b) An effect for the diamond bit not positive on the square bit.

Since we are working in finite dimensions, the outcome-effects \hat{x} span $\mathbf{V}^*(A)$. Thus, as vector spaces, $\mathbf{E}(A)$ and $\mathbf{V}^*(A)$ are the same. However, as the diamond bit illustrates, they can have quite different positive cones and, thus, need not be isomorphic as ordered vector spaces.

Processes and subnormalized states. A *subnormalized state* of a model A is an element α of $\mathbf{V}(A)_+$ with $u(\alpha) < 1$. These can be understood as states that allow a nonzero probability $1 - u(\alpha)$ of some generic “failure” event, (e.g., the destruction of the system), represented by the zero functional in $\mathbf{V}^*(A)$.

More generally, we may wish to regard two systems, represented by models A and B , as the input to and output from some process, whether dynamical or purely information-theoretic, that has some probability to destroy the system or otherwise “fail”. Since such a process should preserve probabilistic mixtures, it should be represented mathematically by an affine mapping $T : \Omega(A) \rightarrow \mathbf{V}(B)_+$, taking each normalized state α of A to a possibly sub-normalized state $T(\alpha)$ of B . One can show that such a mapping extends uniquely to a positive linear mapping:

$$T : \mathbf{V}(A) \rightarrow \mathbf{V}(B),$$

so from now on, this is how I represent processes.

Even if a process T has a nonzero probability of failure, it may be possible to reverse its effect with nonzero probability.

Definition 5. A process $T : A \rightarrow B$ is *probabilistically reversible* iff there exists a process S such that, for all $\alpha \in \Omega(A)$, $(S \circ T)(\alpha) = p\alpha$, where $p \in (0, 1]$.

This means that there is a probability $1 - p$ of the composite process $S \circ T$ failing, but a probability p that it will leave the system in its initial state (note that, since $S \circ T$ is linear, p must be constant); where T preserves normalization, so that $T(\Omega(A)) \subseteq \Omega(B)$, S can also be taken to be normalization-preserving and will undo the result of T with probability one. This is the more usual meaning of “reversible” in the literature.

Given a process $T : \mathbf{V}(A) \rightarrow \mathbf{V}(B)$, there is a dual mapping $T^* : \mathbf{V}^*(B) \rightarrow \mathbf{V}^*(A)$, also positive, given by $T^*(b)(\alpha) = b(T(\alpha))$ for all $b \in \mathbf{V}^*(B)$ and $\alpha \in \mathbf{V}(A)$. The assumption that T takes normalized states to subnormalized states is equivalent to the requirement that $T^*(u_B) \leq u_A$, that is that T^* maps effects to effects.

Remark 1. Since we are attaching no special physical interpretation to the cone $\mathbf{E}_+(A)$, we do not require a physical process $T : \mathbf{V}(A) \rightarrow \mathbf{V}(B)$ to have a dual process T^* that maps $\mathbf{E}_+(B)$ to $\mathbf{E}_+(A)$. That is, we do not require T^* to be positive as a mapping $\mathbf{E}(B) \rightarrow \mathbf{E}(A)$.

Joint probabilities and joint states. If \mathcal{M}_1 and \mathcal{M}_2 are two test spaces, with outcome-spaces X_1 and X_2 , we can construct a space of product tests (note here the savage abuse of notation: $\mathcal{M}_1 \times \mathcal{M}_2$ is not the Cartesian product of \mathcal{M}_1 and \mathcal{M}_2):

$$\mathcal{M}_1 \times \mathcal{M}_2 = \{ E \times F \mid E \in \mathcal{M}_1, F \in \mathcal{M}_2 \}$$

This models a situation in which tests from \mathcal{M}_1 and from \mathcal{M}_2 can be performed separately, and the results collated. Note that the outcome-space for $\mathcal{M}_1 \times \mathcal{M}_2$ is $X_1 \times X_2$. A joint probability weight on \mathcal{M}_1 and \mathcal{M}_2 is just a probability weight on $\mathcal{M}_1 \times \mathcal{M}_2$, that is a function $\omega : X_1 \times X_2 \rightarrow [0, 1]$ such that $\sum_{(x,y) \in E \times F} \omega(x, y) = 1$ for all tests $E \in \mathcal{M}_1$ and $F \in \mathcal{M}_2$. One says that ω is non-signaling iff the marginal (or reduced) probability weights ω_1 and ω_2 , given by:

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

are well-defined, i.e., independent of the choice of the tests E and F , respectively. One can understand this to mean that the choice of which test to measure on \mathcal{M}_1 has no observable, i.e., no statistical, influence on the outcome of tests made of \mathcal{M}_2 , and vice versa. In this case, one also has well-defined conditional probability weights:

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

(with, say, $\omega_{2|x} = 0$ if $\omega_1(x) = 0$, and similarly for $\omega_{1|y}$). This gives us the following bipartite version of the law of total probability [23]: for any choice: of $E \in \mathcal{M}_1$ or $F \in \mathcal{M}_2$,

$$\omega_2 = \sum_{x \in E} \omega_1(x) \omega_{2|x} \quad \text{and} \quad \omega_1 = \sum_{y \in F} \omega_2(y) \omega_{1|y}. \tag{1}$$

Definition 6. A joint state on a pair of probabilistic models A and B is a non-signaling joint probability weight ω on $\mathcal{M}(A) \times \mathcal{M}(B)$ such that, for every $x \in X(A)$ and every $y \in X(B)$, the conditional probability weights $\omega_{2|x}$ and $\omega_{1|y}$ belong to $\Omega(A)$ and $\Omega(B)$, respectively. It follows from (1) that the marginal weights ω_1 and ω_2 are also states of A and B , respectively.

This naturally suggests that one should define, for models A and B , a composite model AB , the states of which would be precisely the joint states on A and B . If one takes $\mathcal{M}(AB) = \mathcal{M}(A) \times \mathcal{M}(B)$, this is essentially the “maximal tensor product” of A and B [24]. However, this does not coincide with the usual composite of quantum-mechanical systems. In Section 6, I will discuss composite systems in more detail. Meanwhile, for the main results of this paper, the idea of a joint state is sufficient.

For a simple example of a joint state that is neither classical, nor quantum, let B denote the “square bit” model discussed above. That is, $B = (\mathcal{B}, \Omega)$ where $\mathcal{B} = \{ \{x, x'\}, \{y, y'\} \}$ is a test space with two non-overlapping, two-outcome tests, and Ω is the set of all probability weights thereon, amounting to the unit square in \mathbb{R}^2 . The joint state on $\mathcal{B} \times \mathcal{B}$ given by Table 1 (a variant of the “non-signaling box” of Popescu and Rohrlich [25]) is clearly non-signaling. Notice that it also establishes a perfect, uniform correlation between the outcomes of any test on the first system and its counterpart on the second.

Table 1. A joint state for two square bits.

	x	x'	y	y'
x	1/2	0	1/2	0
x'	0	1/2	0	1/2
y	0	1/2	1/2	0
y'	1/2	0	0	1/2

Conditioning maps. If ω is a joint state on A and B , define the associated *conditioning maps* $\hat{\omega} : X(A) \rightarrow \mathbf{V}(B)$ and $\hat{\omega}^* : X(B) \rightarrow \mathbf{V}(A)$ by:

$$\hat{\omega}(x)(y) = \omega(x, y) = \hat{\omega}^*(y)(x)$$

for all $x \in X(A)$ and $y \in X(B)$. Note that $\hat{\omega}(x) = \omega_1(x)\omega_{2|x}$ for every $x \in X(A)$, i.e., $\hat{\omega}(x)$ can be understood as the un-normalized conditional state of B given the outcome x on A . Similarly, $\hat{\omega}^*(y)$ is the unnormalized conditional state of A given outcome y on B .

The conditioning map $\hat{\omega}$ extends uniquely to a positive linear mapping $\mathbf{E}(A) \rightarrow \mathbf{V}(B)$, which I also denote by $\hat{\omega}$, such that $\hat{\omega}(\hat{x}) = \hat{\omega}(x)$ for all outcomes $x \in X(A)$. To see this, consider the linear mapping $T : \mathbf{V}^*(A) \rightarrow \mathbb{R}^{X(B)}$ defined, for $f \in \mathbf{V}^*(A)$, by $T(f)(y) = f(\hat{\omega}^*(y))$ for all $y \in X(B)$. If $f = \hat{x}$, we have $T(\hat{x}) = \omega_1(x)\omega_{2|x} \in \mathbf{V}(B)_+$, whence, for all $y \in X(B)$, $T(\hat{x})(y) = \omega(x, y) = \hat{\omega}(x)(y)$. Since the evaluation functionals \hat{x} span $\mathbf{E}(A)$, the range of T lies in $\mathbf{V}(B)$, and moreover, T is positive on the cone $\mathbf{E}(A)_+$. Hence, as advertised, T defines a positive linear mapping $\mathbf{E}(A) \rightarrow \mathbf{V}(B)$, extending $\hat{\omega}$. In the same way, $\hat{\omega}^*$ defines a positive linear mapping $\mathbf{E}(B) \rightarrow \mathbf{V}(A)$.

An immediate and important corollary is that any joint state ω on A and B defines a bilinear form, which by abuse of notation I also call ω , on $\mathbf{E}(A) \times \mathbf{E}(B)$, given by $\omega(a, b) := \hat{\omega}(a)(b)$ for all $a, b \in \mathbf{E}(A)$. Note that $\omega(\hat{x}, \hat{y}) = \omega(x, y)$ for all $x \in X(A), y \in X(B)$ and also that the bilinear form ω is positive, in the sense that $\omega(a, b) \geq 0$ for all $a \in \mathbf{E}(A)_+$ and all $b \in \mathbf{E}(B)_+$.

4. Conjugates and Filters

We are now in a position to abstract the two features of QM discussed earlier. Call a test space (X, \mathcal{M}) *uniform* iff all tests $E \in \mathcal{M}$ have the same size, which we then call the *rank* of the test space. The test spaces associated with quantum models are uniform, and it is quite easy to generate many other examples (see Appendix A).

A uniform test space of rank n always admits at least one probability weight, namely the maximally-mixed probability weight $\rho(x) = 1/n$ for all $x \in X$. I will say that a probabilistic model A is uniform if the test space $\mathcal{M}(A)$ is uniform and the maximally-mixed state ρ belongs to $\Omega(A)$.

By an *isomorphism* $\gamma : A \rightarrow B$ from a probabilistic model A to a probabilistic model B , I mean the obvious thing: a bijection $\gamma : X(A) \rightarrow X(B)$ taking $\mathcal{M}(A)$ onto $\mathcal{M}(\bar{A})$, and such that $\beta \mapsto \beta \circ \gamma$ maps $\Omega(\bar{A})$ onto $\Omega(A)$.

Definition 7. Let A be uniform probabilistic model with tests of size n . A *conjugate* for A is a model \bar{A} , plus a chosen isomorphism $\gamma_A : A \simeq \bar{A}$ and a joint state η_A on A and \bar{A} such that for all $x, y \in X(A)$,

- (a) $\eta_A(x, \bar{x}) = 1/n$
- (b) $\eta_A(x, \bar{y}) = \eta(y, \bar{x})$

where $\bar{x} := \gamma_A(x)$.

This corresponds to what is called a “weak conjugate” in [17]. Note that if $E \in \mathcal{M}(A)$, we have $\sum_{x,y \in E \times E} \eta_A(x, \bar{y}) = 1$ and $|E| = n$. Hence, $\eta_A(x, \bar{y}) = 0$ for $x, y \in E$ with $x \neq y$. Thus, η_A establishes a perfect, uniform correlation between any test $E \in \mathcal{M}(A)$ and its counterpart, $\bar{E} := \{\bar{x} | x \in E\}$, in $\mathcal{M}(\bar{A})$.

The symmetry condition (b) is pretty harmless. If η is a joint state on A and \bar{A} satisfying (a), then so is $\eta^t(x, \bar{y}) := \eta(y, \bar{x})$; thus, $\frac{1}{2}(\eta + \eta^t)$ satisfies both (a) and (b). In fact, if A is sharp, (b) is automatic: if η satisfies (a), then the conditional state $(\eta_A)_{1|\bar{x}}$ assigns probability one to the outcome x . If A is sharp, this implies that $\eta_{1|\bar{x}} = \delta_x$ is uniquely defined, whence $\eta(x, \bar{y}) = n\delta_y(x)$ is also uniquely defined. In other words, for a sharp model A and a given isomorphism $\gamma : A \simeq \bar{A}$, there exists at most one joint state η satisfying (a); whence, in particular, $\eta = \eta^t$.

If $A = A(\mathcal{H})$ is the quantum-mechanical model associated with an n -dimensional Hilbert space \mathcal{H} , then we can take $\bar{A} = A(\bar{\mathcal{H}})$ and define $\eta_A(x, \bar{y}) = |\langle \Psi, x \otimes \bar{y} \rangle|^2$, where Ψ is the EPR state on $\mathcal{H} \otimes \bar{\mathcal{H}}$, as discussed in Section 3.

So much for conjugates. We generalize the filters associated with pure CP mappings as follows:

Definition 8. A filter associated with a test $E \in \mathcal{M}(A)$ is a positive linear mapping $\Phi : \mathbf{V}(A) \rightarrow \mathbf{V}(A)$ such that for every outcome $x \in E$, there is some coefficient $t_x \in [0, 1]$ with $\Phi(\alpha)(x) = t_x \alpha(x)$ for every state $\alpha \in \Omega(A)$.

Equivalently, Φ is a filter iff the dual process $\Phi^* : \mathbf{V}^*(A) \rightarrow \mathbf{V}^*(A)$ satisfies $\Phi^*(\hat{x}) = t_x \hat{x}$ for each $x \in E$. Just as in the quantum-mechanical case, a filter independently attenuates the “sensitivity” of the outcomes $x \in E$. (The extreme case is one in which the coefficient t_x corresponding to a particular outcome is one, and the other coefficients are all zero. In that case, all outcomes other than x are, so to say, blocked by the filter. Conversely, given such an “all or nothing” filter Φ_x for each $x \in E$, we can construct an arbitrary filter with coefficients t_x by setting $\Phi = \sum_{x \in E} t_x \Phi_x$.)

Call a filter Φ reversible iff Φ is an order-automorphism of $\mathbf{V}(A)$; that is, iff it is probabilistically reversible as a process. Evidently, this requires that all the coefficients t_x be nonzero. We will eventually see that the existence of a conjugate, plus the preparability of arbitrary nonsingular states by symmetric reversible filters, will be enough to force A to be a Jordan model. Most of the work is done by the easy Lemma 1, below. First, some terminology.

Definition 9. Suppose $\Delta = \{\delta_x | x \in X(A)\}$ is a family of states indexed by outcomes $x \in X(A)$ and such that $\delta_x(x) = 1$. Say that a state α is spectral with respect to Δ iff there exists a test $E \in \mathcal{M}(A)$ such that $\alpha = \sum_{x \in E} \alpha(x) \delta_x$. Say that the model A itself is spectral with respect to Δ if every state of A is spectral with respect to Δ .

If A has a conjugate \bar{A} , then the bijection $\gamma_A : X(A) \rightarrow X(\bar{A})$ extends to an order-isomorphism $\mathbf{E}(A) \simeq \mathbf{E}(\bar{A})$. It follows that every non-signaling joint probability weight ω on A and \bar{A} defines a bilinear form $a, b \mapsto \omega(a, \bar{b})$ on $\mathbf{E}(A)$.

The following is essentially proven in [17], but the presentation here is somewhat different.

Lemma 1. Let A have a conjugate (\bar{A}, η_A) . Suppose A is spectral with respect to the states $\delta_x := \eta_{1|\bar{x}}$, $x \in X(A)$. Then:

$$\langle a, b \rangle := n \eta_A(a, \bar{b}),$$

where n is the rank of A , defines a self-dualizing inner product on $\mathbf{E}(A)$, with respect to which $\mathbf{V}(A)_+ \simeq \mathbf{E}(A)_+$. Moreover, A is sharp, and $\mathbf{E}(A)_+ = \mathbf{V}^*(A)_+$.

Proof. That $\langle \cdot, \cdot \rangle$ is symmetric and bilinear follows from η_A 's being symmetric and non-signaling. Note that $\langle \hat{x}, \hat{x} \rangle = 1$ for every $x \in X(A)$ and $\langle \hat{x}, \hat{y} \rangle = 0$ for any distinct $x, y \in X(A)$ lying in a common test. We need to show that $\langle \cdot, \cdot \rangle$ is positive-definite. Since $\hat{A} \simeq A$ and the latter is spectral, so is the former. It follows that $\hat{\eta}$ takes $\mathbf{E}(A)_+$ onto $\mathbf{V}(\bar{A})_+$ and, hence, is an order-isomorphism. From this, it follows that every $a \in \mathbf{E}(A)_+$ has a “spectral” decomposition of the form $\sum_{x \in E} t_x x$ for some coefficients $t_x \geq 0$ and some test $E \in \mathcal{M}(A)$. In fact, any $a \in \mathbf{E}(A)$, positive or otherwise, has such a decomposition (albeit with possibly negative coefficients). If $a \in \mathbf{E}(A)$ is arbitrary, with $a = a_1 - a_2$ for some $a_1, a_2 \in \mathbf{E}(A)_+$, we can find $N \geq 0$ with $a_2 \leq Nu$. Thus, $b := a + Nu = a_1 + (Nu - a_2) \geq 0$, and so, $b := \sum_{x \in E} t_x x$ for some $E \in \mathcal{M}(A)$, and hence, $a = b - Nu = \sum_{x \in E} t_x x - N(\sum_{x \in E} x) = \sum_{x \in E} (t_x - N)x$.

Now, let $a \in \mathbf{E}(A)$. Decomposing $a = \sum_{x \in E} t_x x$ for some test E and some coefficients t_x , we have:

$$\langle a, a \rangle = \sum_{x, y \in E \times E} t_x t_y \langle \hat{x}, \hat{y} \rangle = \sum_{x \in E} t_x^2 \geq 0.$$

This is zero only where all coefficients t_x are zero, i.e., only for $a = 0$. Therefore, $\langle \cdot, \cdot \rangle$ is an inner product, as claimed.

We need to show that $\langle \cdot, \cdot \rangle$ is self-dualizing. Clearly $\langle a, b \rangle = n\eta_A(a, \bar{b}) \geq 0$ for all $a, b \in \mathbf{E}(A)_+$. Suppose $a \in \mathbf{E}(A)$ is such that $\langle a, b \rangle \geq 0$ for all $b \in \mathbf{E}(A)_+$. Then, $\langle a, \hat{y} \rangle \geq 0$ for all $y \in X$. Now, $a = \sum_{x \in E} t_x \hat{x}$ for some test E ; thus, for all $y \in E$, we have $\langle a, \hat{y} \rangle = t_y \geq 0$, whence, $a \in \mathbf{E}(A)_+$.

Next, we want to show that $\mathbf{E}(A)_+ = \mathbf{V}(A)_+^*$. Since $\hat{\eta} : \mathbf{E}(A) \rightarrow \mathbf{V}(\bar{A})$ is an order-isomorphism, for every $\alpha \in \mathbf{V}(A)$, there exists a unique $a \in \mathbf{E}(A)$ with $\hat{\eta}(a) = \frac{1}{n}\bar{\alpha}$. In particular,

$$\langle a, x \rangle = n\eta_A(a, \bar{x}) = \bar{\alpha}(\bar{x}) = \alpha(x).$$

It follows that if $b \in \mathbf{E}(A) = \mathbf{V}^*(A)$,

$$b(\alpha) = \bar{b}(\bar{\alpha}) = \bar{b}n\hat{\eta}_A(a) = n\eta(a, \bar{b}) = \langle a, b \rangle.$$

Since every $a \in \mathbf{E}(A)_+$ has the form $a = \hat{\eta}^{-1}(\frac{1}{n}\bar{\alpha})$ for some $\alpha \in \mathbf{V}(A)_+$, if $b \in \mathbf{V}^*(A)_+$, we have $\langle a, b \rangle \geq 0$ for all $a \in \mathbf{E}(A)_+$, whence, by the self-duality of the latter cone, $b \in \mathbf{E}(A)_+$. Thus, $\mathbf{V}^*(A) = \mathbf{E}(A)_+$.

Finally, let us see that A is sharp. If $\alpha \in \Omega(A)$, let a be the unique element of $\mathbf{E}(A)_+$ with $\langle a, x \rangle = \alpha(x)$. In particular, $\langle a, u \rangle = 1$. If a has spectral decomposition $a = \sum_{x \in E} t_x \hat{x}$, where $E \in \mathcal{M}(A)$, then for all $x \in E$, $\langle a, x \rangle = t_x$; hence, $\sum_{x \in E} t_x = \sum_{x \in E} \langle a, x \rangle = \langle a, u \rangle = 1$. Thus, $\|a\|^2 = \sum_{x \in E} t_x^2 \leq 1$, whence, $\|a\| \leq 1$. Now, suppose $\alpha(x) = 1$ for some $x \in X(A)$: then, $1 = \langle a, x \rangle \leq \|a\| \|x\|$; as $\|x\| = 1$, we have $\|a\| = 1$. However, now $\langle a, \hat{x} \rangle = \|a\| \|\hat{x}\|$, whence, $a = \hat{x}$. Hence, there is only one weight α with $\alpha(x) = 1$, namely, $\alpha = \langle x, \cdot \rangle$, so A is sharp. \square

If A is sharp, then we say that A is *spectral* iff it is spectral with respect to the pure states δ_x defined by $\delta_x(x) = 1$. If A is sharp and has a conjugate \bar{A} , then, as noted earlier, the state $\eta_{1|\bar{x}}$ is exactly δ_x , so the spectrality assumption in Lemma 1 is fulfilled if we simply say that A is spectral. Hence, a sharp, spectral model with a conjugate is self-dual.

For the simplest systems, this is already enough to secure the desired representation in terms of a Euclidean Jordan algebra.

Definition 10. Call A a *bit* iff it has rank two (that is, all tests have two outcomes) and if every state $\alpha \in \Omega(A)$ can be expressed as a mixture of two sharply distinguishable states; that is, $\alpha = t\delta_x + (1 - t)\delta_y$ for some $t \in [0, 1]$ and states δ_x and δ_y with $\delta_x(x) = 1$ and $\delta_y(y) = 1$ for some test $\{x, y\}$.

Corollary 1. If A is a sharp bit, then $\Omega(A)$ is a ball of some finite dimension d .

The proof is given in Appendix C. If d is 2, 3 or 5, we have a real, complex or quaternionic bit. For $d = 4$ or $d \geq 6$, we have a non-quantum spin factor.

For systems of higher rank (higher “information capacity”), we need to assume a bit more. Suppose A satisfies the hypotheses of Lemma 1. Appealing to the Koecher–Vinberg theorem, we see that if $\mathbf{V}(A)$ and, hence, $\mathbf{V}^*(A)$ are also homogeneous, then $\mathbf{V}^*(A)$ carries a canonical Jordan structure. In fact, we can say something a little stronger.

Theorem 2. Let A be spectral with respect to a conjugate system \bar{A} . If $\mathbf{V}(A)$ is homogeneous, then there exists a canonical Jordan product on $\mathbf{E}(A)$ with respect to which u_A is the Jordan unit. Moreover, with respect to this product, $X(A)$ is exactly the set of primitive idempotents, and $\mathcal{M}(A)$ is exactly the set of Jordan frames.

The first part is almost immediate from the Koecher–Vinberg theorem, together with Lemma 1. The KV theorem gives us an isomorphism between the ordered vector spaces $\mathbf{V}(A)$ and $\mathbf{E}(A)$, so if one is homogeneous, so is the other. Since $\mathbf{E}(A)$ is also self-dual by Lemma 1, the KV theorem yields the requisite unique Euclidean Jordan structure having u as the Jordan unit. One can then show without

much trouble that every outcome $x \in X(A)$ is a primitive idempotent of $\mathbf{E}(A)$ with respect to this Jordan structure and that every test is a Jordan frame. The remaining claims (that every minimal idempotent belongs to $X(A)$ and every Jordan frame, to $\mathcal{M}(A)$) take a little bit more work. I will not reproduce the proof here; the details (which are not especially difficult, but depend on some facts concerning Euclidean Jordan algebras) can be found in [17].

The homogeneity of $\mathbf{V}(A)$ can be understood as a preparability assumption: it is equivalent to saying that every state in the interior of $\Omega(A)$ can be obtained, up to normalization, from the maximally-mixed state by a reversible process. That is, if $\alpha \in \Omega(A)$, there is some such process ϕ such that $\phi(\rho) = p\alpha$ where $0 < p \leq 1$. One can think of the coefficient p as the probability that the process ϕ will yield a nonzero result (more dramatically: will not destroy the system). Thus, if we prepare an ensemble of identical copies of the system in the maximally-mixed state ρ and subject them all to the process ϕ , the fraction that survives will be about p , and these will all be in state α .

In fact, if the hypotheses of Lemma 1 hold, the homogeneity of $\mathbf{E}(A)$ follows directly from the mere existence of reversible filters with arbitrary non-zero coefficients. To see this, suppose $a \in \mathbf{E}(A)_+$ has a spectral decomposition $\sum_{x \in E} t_x \hat{x}$ for some $E \in \mathcal{M}(A)$, with $t_x > 0$ for all x when a belongs to the interior of $\mathbf{E}(A)_+$. Now, if we can find a reversible filter for E with $\Phi(x) = t_x \hat{x}$ for all $x \in E$, then applying this to the order-unit $u = \sum_{x \in E} \hat{x}$ yields a . Thus, $\mathbf{V}^*(A)$ is homogeneous.

Two paths to spectrality. Some axiomatic treatments of quantum theory have taken one or another form of spectrality as an axiom [6,26]. If one is content to do this, then Lemma 1 above provides a very direct route to the Jordan structure of quantum theory. However, spectrality can actually be derived from assumptions that, on their face, seem a good deal weaker, or anyway more transparent (a different path to spectrality is charted in a recent paper [27] by G. Chiribella and C. M. Scandolo).

I will call a joint state on models A and B *correlating* iff it sets up a perfect correlation between some pair of tests $E \in \mathcal{M}(A)$ and $F \in \mathcal{M}(B)$. More exactly:

Definition 11. A joint state ω on probabilistic models A and B correlates a test $E \in \mathcal{M}(A)$ with a test $F \in \mathcal{M}(B)$ iff there exist subsets $E_0 \subseteq E$ and $F_0 \subseteq F$, and a bijection $f : E_0 \rightarrow F_0$ such that $\omega(x, y) = 0$ for $(x, y) \in E \times F$ unless $y = f(x)$. In this case, say that ω correlates E with F along f . A joint state on A and B is correlating iff it correlates some pair of tests $E \in \mathcal{M}(A)$, $F \in \mathcal{M}(B)$.

Note that ω correlates E with F along f iff $\omega(x, f(x)) = \omega_1(x) = \omega_2(f(x))$, which, in turn, is equivalent to saying that $\omega_{2|x}(f(x)) = 1$ for $\omega_1(x) \neq 0$.

Lemma 2. Suppose A is sharp and that every state α of A arises as the marginal of a correlating joint state between A and some model B . Then, A is spectral.

Proof. Suppose $\alpha = \omega_1$, where ω is a joint state correlating a test $E \in \mathcal{M}(A)$ with a test $F \in \mathcal{M}(B)$, say along a bijection $f : E_0 \rightarrow F_0$, where $E_0 \subseteq E$ and $F_0 \subseteq F$. Then, for any $x \in E$ with $\alpha(x) \neq 0$, $\omega_{1|f(x)}(x) = 1$, whence, as A is sharp, $\omega_{1|f(x)} = \delta_x$, the unique state making x certain. It follows from the law of total probability that $\alpha = \sum_{x \in E} \alpha(x) \delta_x$. \square

In principle, the model B can vary with the state α . Lemma 2 suggests the following language:

Definition 12. A model A satisfies the correlation condition iff every state $\alpha \in \Omega(A)$ is the marginal of some correlating joint state of A and some model B .

This has something of the same flavor as the *purification postulate* of [8], which requires that all states of a given system arise as marginals of a pure state on a larger, composite system, unique up to symmetries on the purifying system. However, note that we do not require the correlating joint state to be either pure (which, in classical probability theory, it will not be) or unique.

If A is sharp and satisfies the correlation condition, then every state of A is spectral. If, in addition, A has a conjugate, then for every $x \in X(A)$, we have $\eta_{1|\bar{x}} = \delta_x$. In this case, A is spectral with respect to the family of states $\eta_{1|\bar{x}}$, and the hypotheses of Lemma 1 are satisfied.

Here is another, superficially quite different, way of arriving at spectrality. Suppose A has a conjugate, \bar{A} . Call a transformation Φ symmetric with respect to η_A iff, for all $x, y \in X(A)$,

$$\eta_A(\Phi^* x, \bar{y}) = \eta_A(x, \bar{\Phi}^* y).$$

Say that a state α is preparable by a filter Φ iff $\alpha = \Phi(\rho)$, where ρ is the maximally-mixed state.

Lemma 3. *Let A have a conjugate, \bar{A} , and suppose every state of A is preparable by a symmetric filter. Then, A is spectral.*

Proof. Let $\alpha = \Phi(\rho)$ where Φ is a filter on a test $E \in \mathcal{M}(A)$, say $\Phi(x) = t_x x$ for all $x \in E$. Then:

$$\alpha = \Phi(\hat{\eta}^*(\bar{u})) = \eta(\Phi^*(\cdot), \bar{u}) = \eta(\cdot, \bar{\Phi}^*(\bar{u})) = \sum_{x \in E} \eta(\cdot, t_x \bar{x}) = \sum_{x \in E} t_x \frac{1}{n} \delta_x.$$

□

Thus, the hypotheses of either Corollary 2 or Lemma 3 will supply the needed spectral assumption that makes Lemma 1 work (in fact, it is not hard to see that these hypotheses are actually equivalent, an exercise I leave for the reader).

To obtain a Jordan model, we still need homogeneity. This is obviously implied by the preparability condition in Lemma 3, provided the preparing filters Φ can be taken to be reversible whenever the state to be prepared is non-singular. On the other hand, as noted above, in the presence of spectrality, it is enough to have arbitrary reversible filters, as these allow one to prepare the spectral decompositions of arbitrary non-singular states. Thus, conditions (a) and (b) below both imply that A is a Jordan model. Conversely, one can show that any Jordan model satisfies both (a) and (b), closing the loop [17]:

Theorem 3. *The following are equivalent:*

- (a) A has a conjugate, and every non-singular state can be prepared by a reversible symmetric filter;
- (b) A is sharp, has a conjugate, satisfies the correlation condition and has arbitrary reversible filters;
- (c) A is a Jordan model.

5. Measurement, Memory and Correlation

Of the spectrality-underwriting conditions given in Lemmas 2 and 3, the one that seems less transparent (to me, anyway) is the correlation condition, i.e., that every state arises as the marginal of a correlating bipartite state. While surely less ad hoc than spectrality, this still calls for further explanation. Suppose we hope to implement a measurement of a test $E \in \mathcal{M}(A)$ dynamically. This would involve bringing up an ancilla system B (also uniform, suppose; and which we can suppose, by suitable coarse-graining, if necessary, to have tests of the same cardinality as A 's) in some "ready" state β_0 . We would then subject the combined system AB to some physical process, at the end of which, AB is in some final joint state ω , and B is (somehow!) in one of a set of record states, β_x , each corresponding to an outcome $x \in X(A)$. (This way of putting things takes us close to the usual formulation of the quantum-mechanical "measurement problem", which I certainly do not propose to discuss here. The point is only that, if any dynamical process, describable within the theory, can account for measurement results, it should be consistent with this description.)

We would like to insist that:

- (a) The states β_x are distinguishable, or readable, by some test $F \in \mathcal{M}(B)$. This means that for each $x \in E$, there is a unique $y \in F$ such that $\beta_x(y) = 1$. Note that this sets up an injection $f : E \rightarrow F$.

- (b) The record states must be accurate, in the sense that if we were to measure E on A , and secure $x \in E$, the record state β_x should coincide with the conditional state $\omega_{2|x}$ (if this is not the case, then a measurement of A cannot correctly calibrate the system B as a measuring device for E).

It follows from (a) and (b) that, for $x \in E$ and $y \neq f(x) \in F$,

$$\omega(x, y) = \omega_1(x)\omega_{2|x}(y) = \omega_1(x)\beta_x(y) = 0.$$

In other words, ω must correlate E with F , along the bijection $f : E \rightarrow F_0 \subseteq F$. If the measurement process leaves α undisturbed, in the sense that $\omega_1 = \alpha$, then α dilates to a correlating state. This suggests the following *non-disturbance principle*: every state can be measured, by some test $E \in \mathcal{M}(A)$, without disturbance. Lemma 2 then tells us that if A is sharp and satisfies the non-disturbance principle, every state of A is spectral.

Here is a slightly different, but possibly more compelling, version of this story. Suppose we can perform a test E on A directly (setting aside, that is, any issue of whether or not this can be achieved through some dynamical process): this will result in an outcome x occurring. To do anything with this, we need to record its having occurred. This means we need a storage medium, B and a family of states β_x , one for each $x \in E$, such that if, on performing the test E , we obtain x , then B will be in state β_x . Moreover, these record states need to be readable at a later time, i.e., distinguishable by a later measurement on B . To arrange this, we need A and B to be in a joint state, associated with a joint probability weight ω , such that $\omega_1 = \alpha$ (because we want to have prepared A in the state α) and $\beta_x = \omega_{2|x}$ for every $x \in E$. We then measure E on A ; upon our obtaining outcome $x \in E$, B is in the state β_x . Since the ensemble of states β_x is readable by some $F \in \mathcal{M}(B)$ with $|F| \geq |E|$, we have correlation, and α must also be spectral.

Of course, these desiderata cannot always be satisfied. What is true, in QM, is that for every choice of state α , there will exist some test that is recordable in that state, in the foregoing sense. If we promote this to the general principle, we again see that every state is the marginal of a correlating state, and hence spectral, if A is sharp.

6. Composites and Categories

Thus far, we have been referring to the correlator η_A as a joint state, but dodging the question: *state of what?* Mathematically, nothing much hangs on this question: it is sufficient to regard η_A as a bipartite probability assignment on A and \bar{A} . However, it would surely be more satisfactory to be able to treat it as an actual physical state of some composite system $A\bar{A}$. How should this be chosen? As mentioned above, one possibility is to take $A\bar{A}$ to be the maximal tensor product of the models A and \bar{A} [24]. By definition, this has for its states all non-signaling probability assignments with conditional states belonging to A and \bar{A} . However, we might want composite systems, in particular $A\bar{A}$, to satisfy the same conditions we are imposing on A and \bar{A} , i.e., to be a Jordan model. If so, we need to work somewhat harder: the maximal tensor product will be self-dual only if A is classical.

In order to be more precise about all this, the first step is to decide what ought to count as a composite of two probabilistic models. If we mean to capture the idea of two physical systems that can be acted upon separately, but which cannot influence one another in any observable way (e.g., two spacelike-separated systems), the following seems to capture the minimal requirements:

Definition 13. A non-signaling composite of models A and B is a model AB , together with a mapping $\pi : X(A) \times X(B) \rightarrow \mathbf{V}^*(AB)_+$ such that:

$$\sum_{x \in E, y \in F} \pi(x, y) = u_{AB}$$

and, for $\omega \in \Omega(AB)$, $\omega \circ \pi$ is a joint state on A and B , as defined in Section 2.

The idea here, expressed in Alice-and-Bob language (Alice controlling system A , Bob controlling system B), is that $\pi(x, y)$ is an effect of the composite system AB , corresponding to x being observed by Alice and y , by Bob. In many cases, $\pi(x, y)$ will actually be an outcome in $X(AB)$. Indeed, we usually have $\pi : X(A) \times X(B) \rightarrow X(AB)$ injective, and for $E \in \mathcal{M}(A), F \in \mathcal{M}(B)$, $\pi(E \times F) = \{\pi(x, y) | x \in E, y \in F\}$ a test in $\mathcal{M}(AB)$. The rank of AB will then be the product of the ranks of A and B . Accordingly, let us call a non-signaling composite with these these properties *multiplicative*. Composites in real and complex quantum mechanics are multiplicative; in quaternionic quantum mechanics, with the most plausible definition of tensor product, they are not [28].

Therefore, the question becomes: can one construct, for Jordan models A and B , a non-signaling composite AB that is also a Jordan model? At present, and in this generality, this question seems to be open, but some progress is made in [28]: if neither A , nor B contain the exceptional Jordan algebra as a summand, such a composite can indeed be constructed, and in multiple ways. Moreover, under a considerably more restrictive definition of “Jordan composite”, no Jordan composite AB can exist if either factor has an exceptional summand.

Categories of Self-Dual Probabilistic Models. It is natural to interpret a physical theory as a category, in which objects represent physical systems and morphisms represent physical processes having these systems (or their states) as inputs and outputs. In order to discuss composite systems, this should be a symmetric monoidal category. That is, for every pair of objects A, B , there should be an object $A \otimes B$, and for every pair of morphisms $f : A \rightarrow A'$ and $g : B \rightarrow B'$, there should be a morphism $f \otimes g : A \otimes B \rightarrow A' \otimes B'$, representing the two processes f and g occurring “in parallel”. One requires that \otimes be associative and commutative, and have a unit object I , in the sense that there exist canonical isomorphisms $\alpha_{A,B,C} : A \otimes (B \otimes C) \simeq (A \otimes B) \otimes C$, $\sigma_{A,B} : A \otimes B \simeq B \otimes A$, $\lambda_A : I \otimes A \simeq A$ and $\rho_A : A \otimes I \rightarrow A$. These must satisfy various “naturality conditions”, guaranteeing that they interact correctly; see [29] for details. One also requires that \otimes be bifunctorial, meaning that $\text{id}_A \otimes \text{id}_B = \text{id}_{A \otimes B}$, and if $f : A \rightarrow A', f' : A' \rightarrow A'', g : B \rightarrow B'$ and $g' : B' \rightarrow B''$, then:

$$(f' \otimes g') \circ (f \otimes g) = (f' \circ f) \otimes (g' \circ g).$$

By a *probabilistic theory*, I mean a category of probabilistic models and processes; that is, objects of \mathcal{C} are models, and a morphism $A \rightarrow B$, where $A, B \in \mathcal{C}$, is a process $\mathbf{V}(A) \rightarrow \mathbf{V}(B)$. A *monoidal probabilistic theory* is such a category, \mathcal{C} , carrying a symmetric monoidal structure $A, B \mapsto AB$, where AB is a non-signaling composite in the sense of the definition above. I also assume that the monoidal unit, I , is the trivial Model 1 with $\mathbf{V}(1) = \mathbb{R}$, and that, for all $A \in \mathcal{C}$,

- (a) $\alpha \in \Omega(A)$ iff the mapping $\alpha : \mathbb{R} \rightarrow \mathbf{V}(A)$ given by $\alpha(1) = \alpha$ belongs to $\mathcal{C}(I, A)$;
- (b) The evaluation functional \hat{x} belongs to $\mathcal{C}(A, I)$ for all outcomes $x \in X(A)$.

Call \mathcal{C} *locally tomographic* iff AB is a locally tomographic composite for all $A, B \in \mathcal{C}$. Much of the qualitative content of (finite-dimensional) quantum information theory can be formulated in purely categorical terms [11,18,30]. In particular, in the work of Abramsky and Coecke [18], it is shown that a range of quantum phenomena, notably gate teleportation, is available in any *dagger-compact* category. For a review of this notion, as well as a proof of the following result, see Appendix D:

Theorem 4. *Let \mathcal{C} be a locally-tomographic monoidal probabilistic theory, in which every object $A \in \mathcal{C}$ is sharp, spectral and has a conjugate $\bar{A} \in \mathcal{C}$, with $\eta_A \in \Omega(A\bar{A})$. Assume also that, for all $A, B \in \mathcal{C}$,*

- (i) $\bar{\bar{A}} = A$, with $\eta_{\bar{A}}(\bar{a}, b) = \eta_A(a, \bar{b})$;
- (ii) If $\phi \in \mathcal{C}(A, B)$, then $\bar{\phi} \in \mathcal{C}(\bar{A}, \bar{B})$.

Then, \mathcal{C} has a canonical dagger-compact structure, in which \bar{A} is the dual of A with $\eta_A : \mathbb{R} \rightarrow \mathbf{V}(A\bar{A})$ as the co-unit.

Jordan composites. The local tomography assumption in Theorem 4 is a strong constraint. As is well known, the standard composite of two real quantum systems is not locally tomographic, yet the category of finite-dimensional real mixed-state quantum systems is certainly dagger-compact and satisfies the other assumptions of Theorem 4, so local tomography is definitely not a necessary condition for dagger-compactness.

This raises some questions. One is whether local tomography can simply be dropped in the statement of Theorem 4. At any rate, at present, I do not know of any non-dagger-compact monoidal probabilistic theory satisfying the other assumptions.

Another question is whether there exist examples other than real QM of non-locally-tomographic, but still dagger-compact, monoidal probabilistic theories satisfying the assumptions of Theorem 2. The answer to this is yes. Without going into detail, the main result of [28] is that one can construct a dagger-compact category in which the objects are Hermitian parts of finite-dimensional real, complex and quaternionic matrix algebras, that is the Euclidean Jordan algebras corresponding to finite-dimensional real, complex or quaternionic quantum-mechanical systems, and morphisms are certain completely positive mappings between enveloping complex $*$ -algebras for these Jordan algebras. The monoidal structure gives *almost* the expected results: the composite of two real quantum systems is the real system corresponding to the usual (real) quantum-mechanical composite of the two components (and, in particular, is not locally tomographic). The composite of two quaternionic systems is a real system (see [11] for an account of why this is just what one wants). The composite of a real and a complex, or a quaternionic and a complex, system is again complex. The one surprise is that the composite of two standard complex quantum systems, in this category, is not the usual thing, but rather, comes with an extra superselection rule. This functions to make time-reversal a legitimate physical operation on complex systems, as it is for real and quaternionic systems. This is part of the price one pays for the dagger-compactness of this category.

7. Conclusions

As promised, we have here an easy derivation of something close to orthodox, finite-dimensional QM, from operationally or probabilistically transparent assumptions. As discussed earlier, this approach offers, in addition to its relative simplicity, greater latitude than the locally-tomographic axiomatic reconstructions of [7–10], putting us in the slightly less constrained realm of formally real Jordan algebras. This allows for real and quaternionic quantum systems, superselection rules and even theories, such as the ones discussed in Section 6, in which real, complex and quaternionic quantum systems coexist and interact.

There remains some mystery as to the proper interpretation of the conjugate system \bar{A} . Operationally, the situation is clear enough: if we understand A as controlled by Alice and \bar{A} , by Bob, then if Alice and Bob share the state η_A , then they will always obtain the same result, as long as they perform the same test. However, what does it mean physically that this should be possible (in a situation in which Alice and Bob are still able to choose their tests independently)? In fact, there is little consensus (that I can find, anyway) among physicists as to the proper interpretation of the conjugate of the Hilbert space representing a given quantum-mechanical system. One popular idea is that the conjugate is a time-reversed version of the given system; but why, then, should we expect to find a state that perfectly correlates the two? At any rate, finding a clear physical interpretation of conjugate systems, even (or especially!) in orthodox quantum mechanics, seems to me an urgently important problem.

I would like to close with another problem, this one of mainly mathematical interest. The hypotheses of Theorem 2 yield a good deal more structure than just a homogeneous, self-dual cone. In particular, we have a distinguished set $\mathcal{M}(A)$ of orthonormal observables in $\mathbf{V}^*(A)$, with respect to which every effect has a spectral decomposition. Moreover, with a bit of work, one can show that this decomposition is essentially unique. More exactly, if $a = \sum_i t_i p_i$ where the coefficients t_i are all distinct and the effects p_1, \dots, p_k are associated with a coarse-graining of a test $E \in \mathcal{M}(A)$, then both

the coefficients and the effects are uniquely determined. The details are in Appendix B. Using this, we have a functional calculus on $\mathbf{V}^*(A)$, i.e., for any real-valued function f of a real variable and any effect a with spectral decomposition $\sum_i t_i p_i$ as above, we can define $f(a) = \sum_i f(t_i) p_i$. This gives us a unique candidate for the Jordan product of effects a and b , namely,

$$a \bullet b = \frac{1}{2}((a + b)^2 - a^2 - b^2).$$

We know from Theorem 2 (and thus, ultimately, from the KV theorem) that this is bilinear. The challenge is to show this without appealing to the KV theorem (the fact that the state spaces of “bits” are always balls, as shown in Appendix C, is perhaps relevant here).

Acknowledgments: This paper is partly based on talks given in workshops and seminars in Amsterdam, Oxford, in 2014 and 2015, and was largely written while the author was a guest of the Quantum Group at the Oxford Computing Laboratory, supported by a grant (FQXi-RFP3-1348) from the FQXi foundation. I would like to thank Sonja Smets (in Amsterdam) and Bob Coecke (in Oxford) for their hospitality on these occasions. I also wish to thank Carlo Maria Scandolo for his careful reading of, and useful comments on, two earlier drafts of this paper.

Conflicts of Interest: The author declares no conflict of interest.

Appendix A. Models with Symmetry

Recall that a probabilistic model A is sharp iff, for every measurement outcome $x \in X(A)$, there exists a unique state $\delta_x \in \Omega(A)$ with $\delta_x(x) = 1$. While this is clearly a very strong condition, it is not an unreasonable one. In fact, given the test space $\mathcal{M}(A)$, we can often choose the state space $\Omega(A)$ in such a way as to guarantee that A is sharp. In particular, this is the case when $\mathcal{M}(A)$ enjoys enough symmetry.

Definition A1. Let G be a group. A G -test space is a test space (X, \mathcal{M}) where X is a G -space, that is, where X comes equipped with a preferred G -action $G \times X \rightarrow X, (g, x) \mapsto gx$, such that $gE \in \mathcal{M}$ for all $E \in \mathcal{M}$. A G -model is a probabilistic model A such that (i) $\mathcal{M}(A)$ is a G -test space and (ii) $\Omega(A)$ is invariant under the action of G on probability weights given by $\alpha \mapsto g\alpha := \alpha \circ g^{-1}$ for $g \in G$.

Lemma A1. Let A be a finite-dimensional G -model, and suppose G acts transitively on the outcome space $X(A)$. Suppose also that A is unital, i.e., for every $x \in X(A)$, there exists at least one state α with $\alpha(x) = 1$. Then, there exists a G -invariant convex subset $\Delta \subseteq \Omega(A)$ such that $A' = (\mathcal{M}(A), \Delta)$ is a sharp G -model.

Proof. For each $x \in X(A)$, let F_x denote the face of $\Omega(A)$ consisting of states α with $\alpha(x) = 1$. Let β_x be the barycenter of F_x . It is easy to check that $F_{gx} = gF_x$ for every $g \in G$. Thus, $g\beta_x = \beta_{gx}$, i.e., the set of barycenters β_x is an orbit. Let Δ be the convex hull of these barycenters. Then, Δ is invariant under G . If $\alpha \in \Delta$ with $\alpha(x) = 1$, then $\alpha \in F_x \cap \Delta = \{\beta_x\}$, so $(\mathcal{M}(A), \Delta)$ is sharp. \square

Appendix B. Uniqueness of Spectral Decompositions

Let A be a model satisfying the conditions of Lemma 1. In particular, every $a \in \mathbf{E}(A) = \mathbf{V}^*(A)$ has a spectral representation $a = \sum_{x \in E} t_x \hat{x}$ for some test $E \in \mathcal{M}(A)$. In general, this expansion is highly non-unique. For instance, the unit u_A can be expanded as $\sum_{x \in E} \hat{x}$ for any test $E \in \mathcal{M}(A)$. The aim in this Appendix is to obtain a form of spectral expansion for effects that is unique.

Call a subset of a test an *event*. That is, $D \subseteq X(A)$ is an event iff there exists a test $E \in \mathcal{M}(A)$ with $D \subseteq E$. The probability of an effect D in a state α is $\alpha(D) = \sum_{x \in E} \alpha(x)$. Thus, any event gives rise to an effect, \hat{D} , given by $\hat{D}(\alpha) = \alpha(D)$. Evidently,

$$\hat{D} := \sum_{x \in E} \hat{x}.$$

A test is a maximal event, and for any test $E \in \mathcal{M}(A)$, $\hat{D} = u$.

Definition A2. An effect $p \in \mathbf{V}^*(A)$ is sharp iff it has the form $p = \widehat{D}$ for some event D . A set of sharp effects $p_1, \dots, p_n \in \mathbf{V}^*(A)$ is jointly orthogonal with respect to $\mathcal{M}(A)$ iff there exists a test $E \in \mathcal{M}(A)$ and pairwise disjoint events $D_1, \dots, D_n \subseteq E$ with $p_i = \widehat{D}_i$ for $i = 1, \dots, n$.

Given an arbitrary element $a \in \mathbf{V}^*(A)$ with spectral decomposition $a = \sum_{x \in E} t_x \widehat{x}$, we can isolate distinct values $t_0 > t_1 > \dots > t_k$ of the coefficients t_x . Letting $E_i = \{x \in E | t_x = t_i\}$ and setting $p_i = p(E_i) = \sum_{x \in E_i} \widehat{x}$, we have $a = \sum_i t_i p_i$, with p_1, \dots, p_n jointly orthogonal. Suppose there is another such decomposition, say $a = \sum_j s_j q_j$, with $q_j = \widehat{F}_j = \sum_{y \in F_j} \widehat{y}$, where $F_1, \dots, F_l \subseteq F \in \mathcal{M}(A)$ are pairwise disjoint, and again, with the coefficients in descending order, say $s_0 > s_1 > \dots > s_l$.

Lemma A2. In the situation described above, $t_0 = s_0$ and $p_0 = q_0$.

Proof. Normalize the inner product on $\mathbf{E}(A)$ so that $\|x\| = 1$ for all outcomes x . Then, for any sharp effect $p = \widehat{D}$, D an event, we have $\|D\|^2 = |D|$, the cardinality of D . Choosing any outcome $x_0 \in E_0$, set $\alpha = |x_0\rangle$, i.e., $\alpha(\widehat{x}) = \langle \widehat{x}, \widehat{x}_0 \rangle$ for all $x \in X(A)$. Then, $\alpha \in \Omega(A)$, $\alpha(p_0) = 1$ and $\alpha(p_i) = 0$ for $i > 0$. Thus,

$$t_0 = \alpha(a) = \sum_j s_j \alpha(q_j).$$

Since the coefficients $\alpha(q_j)$ are sub-convex, the right-hand side is no larger than the largest of the values s_j , namely, s_0 . Thus, $t_0 \leq s_0$. The same argument, with the roles of the two decompositions reversed, shows that $s_0 \leq t_0$. Thus, $s_0 = t_0$.

Now again, let $x \in E_0$: then,

$$\langle \widehat{x}, p_0 \rangle = \sum_{y \in E_0} \langle \widehat{x}, \widehat{y} \rangle = \langle x, x \rangle = 1,$$

whence, $\langle \widehat{x}, a \rangle = t_0$. However, we then have (using the fact that $s_0 = t_0$):

$$t_0 = \langle \widehat{x}, a \rangle = \left\langle \widehat{x}, t_0 p_0 + \sum_{j=1}^l s_j q_j \right\rangle = t_0 \langle \widehat{x}, p_0 \rangle + \sum_{j=1}^l s_j \langle \widehat{x}, q_j \rangle.$$

Since $\sum_{j=0}^l \langle \widehat{x}, q_j \rangle \leq \langle \widehat{x}, u \rangle \leq 1$, the sum in the last expression above is a sub-convex combination of the distinct values $s_0 > \dots > s_l$. This can equal $t_0 = s_0$, the maximum of these values, only if $\langle \widehat{x}, q_0 \rangle = 1$ and $\langle \widehat{x}, q_j \rangle = 0$ for the remaining q_j . It follows that $\langle p_0, q_0 \rangle = \sum_{x \in E_0} \langle \widehat{x}, q_0 \rangle = |E_0| = \|p_0\|^2$. The same argument, with p 's and q 's interchanged, shows that $\langle p_0, q_0 \rangle = \|q_0\|^2$. Hence, $\|p_0\| = \|q_0\|$, and $\langle p_0, q_0 \rangle = \|p_0\|^2 = \|p_0\| \|q_0\|$, whence, $p_0 = q_0$ □

Proposition A1. Every $a \in \mathbf{V}^*(A)$ has a unique expansion of the form $a = \sum_{i=0}^k t_i p_i$ where $t_0 > t_1 > \dots > t_k$ are non-zero coefficients and p_1, \dots, p_n are jointly orthogonal sharp effects.

Proof. Suppose $a = \sum_{i=1}^k t_i p_i$, as above, and also $a = \sum_{j=1}^l s_j q_j$, with $s_0 > \dots > s_l > 0$ and q_j pairwise orthogonal sharp effects. We shall show that $k = l$, and that $t_i = s_i$ and $p_i = q_i$ for each $i = 1, \dots, k$. Lemma A2 tells us that $t_0 = s_0$ and $p_0 = q_0$. Hence,

$$\sum_{i=1}^k t_i p_i = a - t_0 p_0 = a - s_0 q_0 = \sum_{j=1}^l s_j q_j.$$

Applying Lemma A2 recursively, we find that $t_i = s_i$ and $p_i = q_i$ for $i = 1, \dots, \min(k, l)$. If $k \neq l$, say $k < l$; we then have:

$$t_k p_k = s_k q_k + \sum_{j=k+1}^l s_j q_j = t_k p_k + \sum_{j=k+1}^l s_j q_j$$

whence, $\sum_{j=k+1}^l s_j q_j = 0$, which is impossible since all q_j are sharp and the coefficients s_j are strictly positive. Hence, $l = k$, and the proof is complete. \square

Appendix C. Bits Are Balls

In most other reconstructions of QM [8–10], the first step is to show that the state space of a bit, that is, a system in which every state is the mixture of two sharply-distinguishable pure states, is a ball. In our approach, this fact is an easy consequence of Lemma 1. In our framework, we will define a *bit* to be a sharp, uniform model A with rank two, in which every state has the form $t\delta_x + (1 - t)\delta_{x'}$, where $\{x, x'\} \in \mathcal{M}(A)$. Note that this implies that A is spectral.

Lemma A3. *Let A be a bit with conjugate \bar{A} . Then, $\Omega(A)$ is a Euclidean ball, the extreme points of which are the states $\delta_x, x \in X(A)$.*

Proof. By Lemma 1, $\mathbf{E}(A)$ carries a self-dualizing inner product such that $\langle \hat{x}, \hat{y} \rangle = 0$ for $\{x, y\} \in \mathcal{M}(A)$, and which we can normalize so that $\|\hat{x}\| = 1$ for each outcome $x \in X(A)$, so that $\langle u, \hat{x} \rangle = \langle \hat{x}, \hat{x} \rangle = 1$ and $\|u\|^2 = 2$. Every state $\alpha \in \Omega(A)$ corresponds to a unique vector $a \in \mathbf{E}(A)_+$ with $\langle a, u \rangle = 1$, where $\alpha(x) = \langle a, \hat{x} \rangle$ for all $x \in X(A)$; conversely, every vector $a \in \mathbf{E}(A)_+$ with $\langle a, u \rangle = 1$ corresponds in this way to a state. In particular, the state δ_x corresponds to the unit vectors \hat{x} , and the maximally-mixed state corresponds to the vector $\frac{1}{n}u$. To simplify the notation, let us agree for the moment to write ρ for this vector. Thus, $\langle \rho, \hat{x} \rangle = \frac{1}{2}$, $\|\rho\|^2 = \frac{1}{4}\langle u, u \rangle = \frac{1}{2}$, and hence,

$$\|\rho - \hat{x}\|^2 = \|\rho\|^2 - 2\langle \rho, \hat{x} \rangle + \|\hat{x}\|^2 = \frac{1}{2}.$$

Thus, $\hat{X}(A) := \{ \hat{x} \mid x \in X(A) \}$ lies on the sphere of radius $1/\sqrt{2}$ about the state ρ . I now claim that any $a \in \mathbf{E}(A)$ with $\langle a, u \rangle = 1$ (in effect, any state) such that $\|\rho - a\| \leq 1/\sqrt{2}$ belongs to the positive cone $\mathbf{E}(A)_+$. To see this, use spectrality to decompose a as $s\hat{x} + t\hat{y}$ where $\{x, y\} \in \mathcal{M}(A)$ and $s, t \in \mathbb{R}$. Consider now the two-dimensional subspace $\mathbf{E}_{x,y}$ spanned by \hat{x} and \hat{y} . With respect to the inner product inherited from \mathbf{E} , we can regard this as a two-dimensional Euclidean space, in which a is represented by the Cartesian coordinate pair (s, t) . Expanding ρ as $\rho = \frac{1}{2}(\hat{x} + \hat{y})$, we see that $\rho \in \mathbf{E}_{x,y}$ with coordinates $(1/2, 1/2)$. The point (t, s) lies, therefore, in the disk of radius $1/\sqrt{2}$ centered at $(1/2, 1/2)$ in $\mathbf{E}_{x,y}$. Moreover, as $\langle a, u \rangle = 1$, we see that $s + t = 1$, i.e., (s, t) lies on the line of slope -1 through $(1/2, 1/2)$. This puts (s, t) in the positive quadrant of this plane, i.e., $s \geq 0$ and $t \geq 0$. However, then $a \in \mathbf{E}(A)_+$, as claimed. \square

It follows that, for rank-two models, we do not even need to invoke homogeneity: they all correspond to spin factors. Letting d denote the dimension of the state space (that is, $d = \dim(\mathbf{E}) - 1$), we see that if $d = 1$, we have the classical bit; $d = 2$ gives the real quantum-mechanical bit, $d = 3$ gives the familiar Bloch sphere, i.e., the usual qubit of complex QM; while $d = 5$ corresponds to the quaternionic unit sphere, giving us the quaternionic bit. The generalized bits with $d = 4$ and $d \geq 6$ are more exotic “post-quantum” possibilities.

Appendix D. Locally-Tomographic and Dagger-Compactness

A *dagger* on a category \mathcal{C} is a contravariant functor $\dagger : \mathcal{C} \rightarrow \mathcal{C}$ that is the identity on objects and satisfies $\dagger \circ \dagger = \text{id}_{\mathcal{C}}$. That is, if $A \xrightarrow{f} B$ is a morphism in \mathcal{C} , then $A \xleftarrow{f^\dagger} B$, with $f^{\dagger\dagger} = f$ and $(f \circ g)^\dagger = g^\dagger \circ f^\dagger$ whenever $f \circ g$ is defined. An isomorphism $f : A \simeq B$ in \mathcal{C} is then said to be *unitary* iff $f^\dagger = f^{-1}$. One says that \mathcal{C} is *†-monoidal* iff \mathcal{C} is equipped with a symmetric monoidal structure \otimes such that $(f \otimes g)^\dagger = f^\dagger \otimes g^\dagger$, and such that the canonical isomorphisms $\alpha_{A,B,C}, \sigma_{A,B}, \lambda_A$ and ρ_A are all unitary.

A dual for an object A in a symmetric monoidal category \mathcal{C} is a structure (A', η, ϵ) where $A' \in \mathcal{C}$ and $\eta : I \rightarrow A \otimes A'$ and $\epsilon : A' \otimes A \rightarrow I$, such that:

$$(\text{id}_A \otimes \epsilon) \circ (\eta \otimes \text{id}_A) = \text{id}_A \quad \text{and} \quad (\epsilon \otimes \text{id}_{A'}) \circ (\text{id}_{A'} \otimes \eta) = \text{id}_{A'}$$

up to the natural associator and unit isomorphisms. If \mathcal{C} is \dagger -monoidal and $\epsilon = \sigma_{A,A'} \circ \eta_{A'}^\dagger$, then (A', η, ϵ) is a dagger-dual. A category in which every object A has a specified dual (A', η_A, ϵ_A) is compact closed, and a dagger-monoidal category in which every object has a given dagger-dual is dagger-compact. See [18,30] for details.

An important example of all this is the category $\mathbf{FdHilb}_{\mathbb{R}}$ of finite-dimensional real Hilbert spaces and linear mappings. If \mathcal{H} and \mathcal{K} are two such spaces and $\phi : \mathcal{H} \rightarrow \mathcal{K}$, let ϕ^\dagger be the usual adjoint of ϕ with respect to the given inner products. Letting $\mathcal{H} \otimes \mathcal{K}$ be the usual tensor product of \mathcal{H} and \mathcal{K} (in particular, with $\langle x \otimes y, u \otimes v \rangle = \langle x, u \rangle \langle y, v \rangle$ for $x, u \in \mathcal{H}$ and $y, v \in \mathcal{K}$), $\mathbf{FdHilb}_{\mathbb{R}}$ is a dagger-monoidal category with \mathbb{R} as the monoidal unit.

Since any $\mathcal{H} \in \mathbf{FdHilb}_{\mathbb{R}}$ is canonically isomorphic to its dual space, we have also a canonical isomorphism $\mathcal{H} \otimes \mathcal{H} \simeq \mathcal{H}^* \otimes \mathcal{H} = \mathcal{L}(\mathcal{H})$ and a canonical trace functional $\text{Tr}_{\mathcal{H}} : \mathcal{H} \otimes \mathcal{H} \rightarrow \mathbb{R}$, uniquely defined by $\text{Tr}_{\mathcal{H}}(x \otimes y) = \langle x, y \rangle$ for all $x, y \in \mathcal{H}$. Taking $\mathcal{H}' = \mathcal{H}$, let $\eta_{\mathcal{H}} \in \mathcal{H} \otimes \mathcal{H}$ be given by $\eta_{\mathcal{H}} = \sum_i x_i \otimes x_i$, where the sum is taken over any orthonormal basis $\{x_i\}$ for \mathcal{H} ; then, for any $a \in \mathcal{H} \otimes \mathcal{H}$, $\langle \eta_{\mathcal{H}}, a \rangle = \text{Tr}(a)$. It is routine to show that $\text{Tr}_{\mathcal{H}} = \sigma_{\mathcal{H},\mathcal{H}} \circ \eta_{\mathcal{H}}^\dagger$, so that $\eta_{\mathcal{H}}$ and $\text{Tr}_{\mathcal{H}}$ make \mathcal{H} its own dagger-dual.

In any compact closed symmetric monoidal category \mathcal{C} , every morphism $\phi : A \rightarrow B$ yields a dual morphism $\phi' : B' \rightarrow A'$ defined by:

$$\phi' = (\text{id}_{A'} \otimes \epsilon_B) \circ (\text{id}_{A'} \otimes \phi \otimes \text{id}_{B'}) \circ (\eta_A \otimes \text{id}_{B'}).$$

(again, suppressing associators and left and right units). For $\phi : \mathcal{H} \rightarrow \mathcal{K}$ in $\mathbf{FdHilb}_{\mathbb{R}}$, one has, for any $v \in A$,

$$\phi'(v) = \sum_{x \in M} \langle v, f(x) \rangle x = \sum_{x \in M} \langle f^\dagger(v), x \rangle x = f^\dagger(v),$$

i.e., $\phi' = \phi^\dagger$.

Now, let \mathcal{C} be a monoidal probabilistic theory; that is, a category of probabilistic models and processes, with a symmetric monoidal structure $A, B \mapsto AB$, where AB is a (non-signaling) composite in the sense discussed in Section 6. Let \mathbb{C} is multiplicative, so that for $A, B \in \mathcal{C}$, we have $\pi_{AB} : X(A) \times X(B) \rightarrow X(AB)$. Henceforward, I will write $x \otimes y$ for $\pi(x, y)$ where $x \in X(A)$ and $y \in X(B)$. I will further assume that \mathcal{C} 's tensor unit is $I = \mathbb{R}$, and that:

- (a) Every $A \in \mathcal{C}$ has a conjugate, $\bar{A} \in \mathcal{C}$, with $\bar{\bar{A}} = A$;
- (b) For all $A, B \in \mathcal{C}$ and $\phi \in \mathcal{C}(A, B)$, $\bar{\phi} \in \mathcal{C}(\bar{A}, \bar{B})$;
- (c) $\bar{\bar{A}} = A$, with $\eta_{\bar{A}}(\bar{a}, b) := \eta_A(a, \bar{b})$.

Remark A1. (1) The chosen conjugate \bar{A} for $A \in \mathcal{C}$ required by Condition (a) is equipped with a canonical isomorphism $\gamma_A : A \simeq \bar{A}$, with $\bar{x} = \gamma(x)$ for every $x \in X(A)$. As discussed in Section 4, this extends to an order-isomorphism $\mathbf{E}(A) \simeq \mathbf{E}(\bar{A})$, which we again write as $\gamma_A(a) = \bar{a}$ for $a \in \mathbf{E}(A)$. Notice, however, that γ_A is not assumed to be a morphism in \mathcal{C} .

(2) In spite of this, Condition (b) requires that $\bar{\phi} = \gamma_B \circ \phi \circ \gamma_A^{-1}$ does belong to $\mathcal{C}(\bar{A}, \bar{B})$ for $\phi \in \mathcal{C}(A, B)$. Notice here that $\phi \mapsto \bar{\phi}$ is functorial.

(3) The second part of Condition (c) is redundant if every model A in \mathbb{C} is sharp (since in this case, there is at most one correlator between \bar{A} and A). Notice, too, that Condition (c) implies that:

$$\langle \bar{x}, \bar{y} \rangle = \eta_{\bar{A}}(\bar{x}, y) = \eta_A(x, \bar{y}) = \langle \hat{x}, \hat{y} \rangle$$

for all $x, y \in \mathbf{E}(A)$.

We are now ready to prove Theorem 4. We continue to assume that \mathcal{C} is a locally-tomographic, multiplicative monoidal probabilistic theory, satisfying Conditions (a), (b) and (c) above. We wish to show that if every $A \in \mathcal{C}$ is sharp and spectral, then \mathcal{C} has a canonical dagger, with respect to which it is dagger-compact.

Before proceeding, it will be convenient to dualize our representation of morphisms, so that $\phi \in \mathcal{C}(A, B)$ means that ϕ is a positive linear mapping $\mathbf{E}(B) \rightarrow \mathbf{E}(A)$ (thus, our co-unit $\eta \in \mathcal{C}(I, A \otimes A')$ becomes a positive linear mapping $\eta_A : \mathbf{E}(A \otimes A') \rightarrow \mathbb{R}$, and similarly, a unit $\epsilon_A \in \mathcal{C}(A' \otimes A, I)$ becomes a positive linear mapping $\mathbb{R} \rightarrow \mathbf{E}(A' \otimes A)$, i.e, an element of $\mathbf{E}(A \otimes A')$). By Lemma 1, for every $A \in \mathcal{C}$, the space $\mathbf{E}(A)$ carries a canonical self-dualizing inner product $\langle \cdot, \cdot \rangle_A$, with respect to which $\mathbf{E}(A) \simeq \mathbf{V}(A)$.

Lemma A4. *For all models $A, B \in \mathcal{C}$, the inner product on $\mathbf{E}(AB)$ factors, in the sense that if $u, x \in \mathbf{E}(A)$ and $v, y \in \mathbf{E}(B)$, then $\langle u \otimes v, x \otimes y \rangle = \langle u, x \rangle \langle v, y \rangle$.*

Proof. This follows from the sharpness of A, B and AB . For $u \in X(A), v \in X(B)$, let δ_u, δ_v and $\delta_{u \otimes v}$ denote the unique states of A, B and AB such that $\delta_u(u) = \delta_v(v) = \delta_{u \otimes v}(u \otimes v) = 1$. Since $(\delta_u \otimes \delta_v)(u \otimes v)$ is also one, we conclude that $\delta_{u \otimes v} = \delta_u \otimes \delta_v$. However, we also have $\delta_u(x) = n \langle \hat{u}, \hat{x} \rangle$, $\delta_v(y) = m \langle \hat{v}, \hat{y} \rangle$ and $\delta_{u \otimes v}(x \otimes y) = nm \langle \hat{u} \otimes \hat{v}, \hat{x} \otimes \hat{y} \rangle$, where n, m and nm are the ranks, respectively, of A, B and $A \otimes B$. This establishes the claim. \square

It follows that \mathcal{C} is a monoidal subcategory of $\mathbf{FdHilb}_{\mathbb{R}}$. In effect, we are going to show that \mathcal{C} inherits a dagger-compact structure from $\mathbf{FdHilb}_{\mathbb{R}}$, with the minor twist that we will take \bar{A} , rather than A , as the dual for $A \in \mathcal{C}$. We define the dagger of $\phi \in \mathcal{C}(A, B)$ to be the Hermitian adjoint of $\phi : \mathbf{E}(A) \rightarrow \mathbf{E}(B)$ with respect to the canonical inner products on $\mathbf{E}(A)$ and $\mathbf{E}(B)$. At this point, it is not obvious that ϕ^\dagger belongs to \mathcal{C} . In order to show that it does, we first need to show that \mathcal{C} is compact closed. To define the unit, let $e_A \in \mathbf{E}(\bar{A}) \otimes \mathbf{E}(\bar{A}) = \mathbf{E}(\bar{A}A)$ (note the use of local tomography here) to be the vector with $\langle e_A, \cdot \rangle = \eta_A$, i.e., for all $a, b \in \mathbf{E}(A)$,

$$\langle e_A, \bar{a} \otimes b \rangle = \eta_A(a \otimes \bar{b}) = \langle a, b \rangle.$$

Since $\mathbf{E}(A\bar{A})$ is self-dual, $e_A \in \mathbf{E}(A\bar{A})_+$.

Lemma A5. *With η_A and e_A defined as above, \bar{A} is a dual for A for every $A \in \mathcal{C}$. In particular, \mathcal{C} is compact closed.*

Proof. Choose an orthonormal basis $M \subseteq \mathbf{E}(A)$. Local tomography and Lemma A4 tell us that $\bar{M} \otimes M = \{\bar{a} \otimes a | a \in M\}$ is then an orthonormal basis for $\mathbf{E}(\bar{A}A)$ (note here that $a, b \in M$ are not necessarily even positive, let alone in $X(A)$). If we expand e_A with respect to this basis, we have:

$$e_A = \sum_{a,b \in M} \langle e_A, \bar{a} \otimes b \rangle \bar{a} \otimes b$$

Since the basis is orthonormal, we have:

$$\langle e_A, \bar{a} \otimes a \rangle = \langle a, a \rangle = \|a\|^2 = 1$$

and for $a \neq b$, both in M ,

$$\langle e_A, \bar{a} \otimes b \rangle = \langle a, b \rangle = 0$$

Hence, $e_A = \sum_{a \in M} \bar{a} \otimes a$. Regarding e_A as a morphism $I \rightarrow \bar{A} \otimes A$, we now have, for any $v \in \mathbf{E}(A)$,

$$\begin{aligned} (\eta_A \otimes \text{id}_A) \circ (\text{id}_A \otimes e_A)(v) &= (\eta_A \otimes \text{id}_A) \left(\sum_{x \in M} v \otimes \bar{a} \otimes a \right) \\ &= \sum_{x \in M} \eta_A(v \otimes \bar{a})a \\ &= \sum_{x \in M} \langle v, a \rangle a = v. \end{aligned}$$

Similarly, for $\bar{v} \in \bar{A}$,

$$\begin{aligned} (\text{id}_{\bar{A}} \otimes \eta_A) \circ (e_{\bar{A}} \otimes \text{id}_A)(\bar{v}) &= (\text{id}_{\bar{A}} \otimes \eta_A) \left(\sum_{a \in M} \bar{a} \otimes a \otimes \bar{v} \right) \\ &= \sum_{x \in M} \bar{a} \eta_A(a, \bar{v}) = \sum_{a \in M} \langle a, \bar{v} \rangle \bar{a} \\ &= \sum_{a \in M} \langle \bar{v}, \bar{a} \rangle \bar{a} = \bar{v}. \end{aligned}$$

□

Lemma A6. *If $\phi : \mathbf{E}(A) \rightarrow \mathbf{E}(B)$ belongs to \mathcal{C} , then so does $\phi^\dagger : \mathbf{E}(B) \rightarrow \mathbf{E}(A)$.*

Proof. Using the compact structure on \mathcal{C} defined above, if $\phi : A \rightarrow B$, we construct the dual of $\bar{\phi}$,

$$\bar{\phi}' := (\eta_B \otimes \text{id}_A) \circ (\text{id}_B \otimes \bar{\phi} \otimes \text{id}_A) \circ (\text{id}_B \otimes e_A) : \mathbf{E}(B) \rightarrow \mathbf{E}(A).$$

Applying this mapping to $b \in \mathbf{E}(B)$, we have:

$$\begin{aligned} b \mapsto (\eta_B \otimes \text{id}_A) \left(\sum_{a \in M} b \otimes \bar{\phi}(\bar{a}) \otimes a \right) &= \sum_{a \in M} \eta_B(b, \bar{\phi}(\bar{a}))a \\ &= \sum_{a \in M} \langle b, \phi(a) \rangle a \\ &= \sum_{a \in M} \langle \phi^\dagger(b), a \rangle a = \phi^\dagger(b). \end{aligned}$$

Thus, $\phi^\dagger = \bar{\phi}'$, which is evidently a morphism in \mathcal{C} .

Thus, \mathcal{C} is a dagger-, as well as a monoidal, sub-category of $\mathbf{FdHilb}_{\mathbb{R}}$. Hence, the associator, swap and left- and right-unit morphisms associated with an object $A \in \mathcal{C}$ are all unitary (since they are unitary in $\mathbf{FdHilb}_{\mathbb{R}}$), whence \mathcal{C} is dagger-monoidal. To complete the proof of Theorem 4, we need to check that $\eta_A = e_A^\dagger \circ \sigma_{A, \bar{A}} : \mathbf{E}(A \bar{A}) \rightarrow \mathbb{R}$. In view of our local tomography assumption, it is enough to check this on pure tensors, where a routine computation gives us $e_A^\dagger(\sigma_{A, \bar{A}}(a \otimes \bar{b})) = \langle e_A^\dagger(\bar{b} \otimes a), 1 \rangle_1 = \langle \bar{b} \otimes a, e_A \rangle_{\bar{A}A} = \langle a, b \rangle = \eta_A(a \otimes \bar{b})$. □

Remark A2. *Given that \mathcal{C} is compact closed, with \bar{A} the dual of A , the functoriality of $\phi \mapsto \bar{\phi}$ makes \mathcal{C} strongly compact closed, in the sense of [18]. This is equivalent to dagger-compactness.*

References

1. Von Neumann, J. *Mathematical Foundations of Quantum Mechanics*; Princeton University Press: Princeton, NJ, USA, 1955.
2. Schwinger, J. The algebra of microscopic measurement. *Proc. Natl. Acad. Sci. USA* **1959**, *45*, 1542–1553.

3. Mackey, G.W. *Mathematical Foundations of Quantum Mechanics*; Dover Publications, Inc.: Mineola, NY, USA, 2004.
4. Ludwig, G. *Foundations of Quantum Mechanics I*; Springer: New York, NY, USA, 1983.
5. Piron, C. *Mathematical Foundations of Quantum Mechanics*; Academic Press: Cambridge, MA, USA, 1978.
6. Barnum, H.; Müller, M.; Ududec, C. Higher-order interference and single-system postulates characterizing quantum theory. *New J. Phys.* **2014**, *16*, 123029.
7. Hardy, L. Quantum theory from five reasonable axioms. *arXiv* **2001**, arXiv:quant-ph/0101012.
8. Chiribella, G.; D'Ariano, M.; Perinotti, P. Informational derivation of quantum theory. *Phys. Rev. A* **2011**, *84*, 012311.
9. Dakic, B.; Brukner, C. Quantum theory and beyond: Is entanglement special? *arXiv* **2009**, arXiv:0911.0695.
10. Masanes, L.; Müller, M. A derivation of quantum theory from physical requirements. *New J. Phys.* **2011**, *13*, 063001.
11. Baez, J. Division algebras and quantum theory. *Found. Phys.* **2012**, *42*, 819–855.
12. Janotta, P.; Lal, R. Generalized probabilistic theories without the no-restriction hypothesis. *Phys. Rev. A* **2013**, *87*, 052131.
13. Faraut, J.; Koranyi, A. *Analysis on Symmetric Cones*; Oxford University Press: London, UK, 1994.
14. Wilce, A. 4.5 axioms for finite-dimensional quantum probability. In *Probability in Physics*; Ben-Menahem, Y., Hemmo, M., Eds.; Springer: New York, NY, USA, 2012.
15. Wilce, A. Symmetry and composition in probabilistic theories. *Electron. Notes Theor. Comput. Sci.* **2011**, *270*, 191–207.
16. Wilce, A. Symmetry, self-duality and the Jordan structure of finite-dimensional quantum mechanics. *arXiv* **2011**, arxiv:1110.6607.
17. Wilce, A. Conjugates, Filters and Quantum Mechanics. *arXiv* **2012**, arxiv.org/pdf/1206.2897.
18. Abramsky, S.; Coecke, B. Abstract Physical Traces. *Theor. Appl. Categories* **2005**, *14*, 111–124.
19. Barnum, H.; Graydon, M.A.; Wilce, A. Some nearly quantum theories. *arXiv* **2015**, arXiv:1507.06278.
20. Jordan, P. Über ein Klasse nichtassoziativer hypercomplexe algebren. *Nachr. Akad. Wiss. Göttingen Math. Phys. Kl. I.* **1933**, *33*, 569–575. (In German)
21. Von Neumann, J. On an algebraic generalization of the quantum mechanical formalism (Part I). *Ann. Math.* **1936**, *1*, 415–484.
22. Aliprantis, C.D.; Toukey, R. *Cones and Duality*; American Mathematical Society: Providence, RI, USA, 2007.
23. Foulis, D.J.; Randall, C.H. Empirical logic and tensor products. In *Interpretations and Foundations of Quantum Theory*; Neumann, H., Ed.; Bibliographisches Inst.: Mannheim, Germany, 1981.
24. Barnum, H.; Wilce, A. Post-classical probability theory. In *Quantum Theory: Informational Foundations and Foils*; Chiribella, G., Spekkens, R., Eds.; Springer: Dordrecht, The Netherlands, 2016.
25. Popescu, S.; Rohrlich, D. Nonlocality as an axiom. *Found. Phys.* **1994**, *24*, 379–385.
26. Gunson, J. On the algebraic structure of quantum mechanics. *Commun. Math. Phys.* **1967**, *6*, 262–285.
27. Chiribella, G.; Scandolo, C.M. Operational axioms for state diagonalization. *arXiv* **2015**, arXiv:1506.00380.
28. Barnum, H.; Graydon, M.; Wilce, A. Composites and categories of Euclidean Jordan algebras. *arXiv* **2016**, arXiv:1606.09331.
29. Mac Lane, S. *Categories for the Working Mathematician*; Springer: New York, NY, USA, 1978.
30. Selinger, P. Dagger compact closed categories and completely positive maps. *Electron. Notes Theor. Comput. Sci.* **2007**, *170*, 139–163.

