Quantum Mechanics from Self-Modeling: Deriving Complex C^*-Algebraic Structure from a Single Operational Premise

Bryan Ehrlich

Abstract

We show that a finite-dimensional system admitting a faithful self-model—an isomorphic copy through which the system can probe and update its own state—is necessarily governed by complex quantum mechanics. Starting from a single operational premise (faithful self-modeling) and four standing structural assumptions, we construct a sequential product on the system's effect algebra using compressions and Peirce decompositions. This product satisfies the seven axioms of van de Wetering, forcing the state space to be a Euclidean Jordan algebra. The faithful self-modeling constraint further implies local tomography for the body–model composite, which excludes all non-complex Jordan types and promotes the algebra to $M_n(\mathbb{C})^{\mathrm{sa}}$ equipped with the Lüders sequential product and conjugate-transpose involution. The result distinguishes this work from other reconstruction programs, which typically require 3–6 independent axioms and often assume local tomography or the complex field. Our derivation requires a single operational premise and four standing structural assumptions (finite-dimensionality, faithfulness, minimal composite, simple EJA). We argue that these assumptions are not independent design choices but structural consequences of self-modeling in a finite-capacity system.

1 Introduction

A finite-dimensional physical system that faithfully models itself is necessarily governed by complex quantum mechanics. More precisely, if the state space of a system admits a faithful self-model—an isomorphic copy through which the system can operationally probe and update its own state—then the state space is $M_n(\mathbb{C})^{\mathrm{sa}}$ for some n, the algebra of self-adjoint $n \times n$ complex matrices, equipped with the Lüders sequential product $a \cdot b = \sqrt{a}\, b \sqrt{a}$

and the conjugate-transpose involution $X^* = X^\dagger$. No assumption about the complex field, Hilbert space structure, or C^*-algebraic operations enters the premises; all of these emerge as consequences.

This result contributes to the quantum reconstruction program, which seeks to derive the mathematical framework of quantum theory from physically transparent principles. Beginning with Hardy's five-axiom reconstruction [Har01], the program has produced a succession of results that characterize quantum theory within the landscape of generalized probabilistic theories (GPTs): Chiribella, D'Ariano, and Perinotti [CDP11] used six axioms including Purification; Dakić and Brukner [DB11] used four axioms including local tomography; Masanes and Müller [MM11, MMAPG13] distilled the framework to three postulates assuming tomographic locality; and Barnum, Müller, and Ududec [BMU14] used four postulates including energy observability. A comprehensive review of the landscape is given by Müller [M21]. A common feature of these programs is the use of multiple independent operational axioms, and in most cases the complex number field—or a property that selects it, such as local tomography—is assumed rather than derived.

Our main result can be stated informally as follows.

Theorem 1.1 (Main result, informal)**.** *If a finite-dimensional system's state space admits a faithful self-model—an order-isomorphic copy through which the system can test effects, update its model, and test again—then the state space is* $M_n(\mathbb{C})^{\mathrm{sa}}$ *for some* $n \geq 2$*, equipped with the Lüders sequential product and conjugate-transpose involution.*

The derivation proceeds through a chain of implications: self-modeling → sequential product construction → axiom verification → Euclidean Jordan algebra classification → local tomography → type exclusion → C^*-algebra promotion → involution exhibition. Figure 1 illustrates this chain, distinguishing the novel contributions of this paper from the published theorems invoked.

Table 1 compares the premise count and complex-selection mechanism of this work against the major reconstruction programs.

The paper is organized as follows. Section 2 introduces order unit spaces, compressions, Peirce decompositions, sequential products, and the self-modeling definition. Section 3 presents the main novel construction: the self-modeling sequential product, the naive extension and its failure, the corrected product via Peirce feedback, and the faithful self-modeling selection principle. Section 4 verifies that the corrected product satisfies axioms S1–S7 of van de Wetering [vdW19b]. Section 5 constructs the body–model

Table 1: Comparison of quantum reconstruction programs. "Premises" counts independent operational postulates. Our premise requires four standing structural assumptions (see Assumptions 2.2–3.7); maximal coherence of the Peirce feedback is derived from (S5), not assumed.

Program	Year	Premises	Complex selection
Hardy [Har01]	2001	5	Simplicity axiom
Chiribella–D'Ariano–Perinotti [CDP11]	2011	6	Purification
Dakić–Brukner [DB11]	2011	4	Local tomography (assumed)
Masanes–Müller [MMAPG13]	2013	3	Tomographic locality (assumed)
Barnum–Müller–Ududec [BMU14]	2014	4	Energy observability
Selby–Scandolo–Coecke [SSC21]	2021	∼6	Symmetric purification
This work	**2026**	**1**	**Derived (LT from faithful tracking)**

composite, proves local tomography, and addresses the entangled sector. Section 6 excludes all non-complex Euclidean Jordan algebra types and promotes the algebra to a C^*-algebra via the theorems of van de Wetering, Barnum–Wilce, and Hanche-Olsen. Section 7 discusses the scope and limitations of the result.

We close this introduction with an honest statement of scope. Our result holds for finite-dimensional systems (Assumption 2.2), assumes a faithful self-model exists (Assumption 2.7), uses the minimal composite construction (Assumption 3.6), and restricts to simple Euclidean Jordan algebras (Assumption 3.7). The claim is therefore: *one operational premise plus four standing structural assumptions derive complex quantum mechanics.* We argue in Section 7 that these structural assumptions are not independent design choices but consequences of what it means to be a faithful self-modeler in a finite-capacity system.

2 Preliminaries

2.1 Order Unit Spaces

We work within the framework of order unit spaces, following Alfsen–Shultz [AS03].

Definition 2.1 (Order unit space)**.** An *order unit space* $(V, V^+, \mathbb{1})$ is a real ordered vector space with a proper cone V^+ (closed, convex, pointed) generating V, equipped with a distinguished Archimedean order unit $\mathbb{1}$: for every $a \in V$ there exists $r \geq 0$ with $-r\mathbb{1} \leq a \leq r\mathbb{1}$. The *order unit norm* is

$\|a\| = \inf\{r \geq 0 : -r\mathbb{1} \leq a \leq r\mathbb{1}\}$.

The *effect space* of V is $[0,1]_V = \{a \in V : 0 \leq a \leq \mathbb{1}\}$. An effect $p \in [0,1]_V$ is *sharp* (a *projective unit*) if p and $\mathbb{1} - p$ are both extremal in $[0,1]_V$, or equivalently if p admits no non-trivial refinement. We write $\mathrm{Proj}(V)$ for the set of projective units.

Every effect in a spectral order unit space admits a *spectral decomposition* $a = \sum_i \lambda_i\, p_i$, where the p_i are mutually orthogonal projective units with $\sum_i p_i = \mathbb{1}$ and the λ_i are distinct elements of $[0,1]$. This decomposition exists as a property of the OUS lattice structure [AS03], prior to and independent of any sequential product.

Assumption 2.2 (Finite-dimensional spectral OUS)**.** The system's state space is described by a finite-dimensional spectral order unit space $(V, V^+, \mathbb{1})$.

2.2 Compressions and Peirce Decomposition

For each projective unit $p \in \mathrm{Proj}(V)$, the Alfsen–Shultz theory provides a *compression* $C_p : V \to V$, a positive linear projection satisfying $C_p(\mathbb{1}) = p$. The compression C_p projects onto the face generated by p, and it plays the role of the Lüders operation $b \mapsto p\,b\,p$ in the general OUS setting.

Definition 2.3 (Peirce decomposition)**.** For a projective unit $p \in \mathrm{Proj}(V)$, the *Peirce decomposition* of V relative to p is the direct-sum decomposition

$$V = V_2(p) \oplus V_1(p) \oplus V_0(p), \tag{1}$$

where $V_2(p) = \mathrm{range}(C_p)$ is the face of p, $V_0(p) = \mathrm{range}(C_{p^\perp})$ is the face of $p^\perp = \mathbb{1} - p$, and $V_1(p) = \ker(C_p) \cap \ker(C_{p^\perp})$ is the *Peirce* 1*-space.*

A key identity is:

$$C_p + C_{p^\perp} = \mathrm{pinch}_p \quad \text{(the pinching map, \textit{not} the identity).} \tag{2}$$

The pinching map annihilates $V_1(p)$. In the commutative (classical) case $V_1(p) = \{0\}$, so the pinching map is the identity; in the non-commutative case $V_1(p) \neq \{0\}$ and the identity fails. This distinction drives the entire construction in Section 3.

More generally, for a resolution of unity $\{p_1, \ldots, p_n\}$ (mutually orthogonal projective units with $\sum_i p_i = \mathbb{1}$), the *Peirce* (i,j) *projection* P_{ij} extracts the component of b in the (i,j) Peirce 1-space $V_1(p_i, p_j)$:

$$P_{ij}(b) \in V_1(p_i, p_j), \qquad b = \sum_i C_{p_i}(b) + \sum_{i<j} P_{ij}(b). \tag{3}$$

2.3 Sequential Products

Following Gudder and Greechie [GG02] and van de Wetering [vdW19b], we define a sequential product on an order unit space.

Definition 2.4 (Sequential product space, [vdW19b, Def. 2])**.** Let $(V, \leq, \mathbb{1})$ be an order unit space equipped with a binary operation $(\cdot) : [0,1]_V \times [0,1]_V \to [0,1]_V$. Write $a \mid b$ ("a and b are compatible") when $a \cdot b = b \cdot a$. Then $(V, \cdot)$ is a *sequential product space* if the following axioms hold:

(S1) **Additivity:** $a \cdot (b + c) = a \cdot b + a \cdot c$ whenever $b + c \leq \mathbb{1}$.

(S2) **Continuity:** The map $a \mapsto a \cdot b$ is continuous in the order unit norm.

(S3) **Unitality:** $\mathbb{1} \cdot a = a$ for all $a \in [0,1]_V$.

(S4) **Symmetry of orthogonality:** If $a \cdot b = 0$ then $b \cdot a = 0$.

(S5) **Associativity of compatible effects:** If $a \mid b$ then $a \cdot (b \cdot c) = (a \cdot b) \cdot c$.

(S6) **Additivity of compatibility:** If $a \mid b$ then $a \mid (\mathbb{1} - b)$; if also $a \mid c$ then $a \mid (b + c)$ whenever $b + c \leq \mathbb{1}$.

(S7) **Multiplicativity of compatibility:** If $a \mid b$ and $a \mid c$ then $a \mid b \cdot c$.

Two central results of van de Wetering are used as black boxes throughout this paper.

Theorem 2.5 ([vdW19b, Thm. 1])**.** *A finite-dimensional sequential product space is order-isomorphic to a Euclidean Jordan algebra.*

Theorem 2.6 ([vdW19b, Thm. 3])**.** *If a sequential product space has a locally tomographic composite that also carries a sequential product, then the sequential product space embeds into a C^*-algebra.*

The Jordan–von Neumann–Wigner classification [JvNW34] of simple finite-dimensional Euclidean Jordan algebras yields five types: $M_n(\mathbb{C})^{\text{sa}}$ over $\mathbb{R}$, $\mathbb{C}$, or $\mathbb{H}$; the spin factors V_n; and the exceptional Albert algebra $M_3(\mathbb{O})^{\text{sa}}$. Our task is to show that self-modeling selects the complex type uniquely.

2.4 Self-Modeling

We now introduce the central operational premise.

Consider a physical system B (the "body") described by a finite-dimensional spectral OUS $(V_B, V_B^+, \mathbb{1}_B)$. A *self-model* of B is a second system M (the "model"), described by an OUS $(V_M, V_M^+, \mathbb{1}_M)$, together with a positive unital map $\varphi : V_B \to V_M$ (the *tracking map*) that encodes the operational cycle:

(i) *Prepare:* the body B is in some state;

(ii) *Update:* testing a sharp effect p on B triggers M to undergo the compression $C_{\varphi(p)}$, updating the model;

(iii) *Measure:* a second effect b is tested on B.

The composite outcome of this cycle—"first p, then b"—defines a binary operation on the effect space of B.

Assumption 2.7 (Faithful self-model)**.** The tracking map $\varphi : V_B \to V_M$ is an order isomorphism. That is, M is a faithful copy of B: it has the same dimension, the same order structure, and the same effect space, up to relabeling.

This is the sole operational premise. Faithful self-modeling means the system possesses a complete internal representation of itself, one that preserves all operationally accessible structure. Approximate or partial self-models (where φ is merely positive or injective but not surjective) might yield weaker algebraic structures; we do not pursue this here.

3 The Self-Modeling Sequential Product

3.1 From Self-Modeling to Sequential Product

The self-modeling cycle of Section 2.4 produces a natural binary operation on the effects of B. For sharp effects $p, q \in \mathrm{Proj}(V_B)$, the composite outcome "first test p, update the model, then test q" is encoded by the compression product:

$$p \cdot q = C_p(q), \tag{4}$$

where C_p is the Alfsen–Shultz compression for the face generated by p in V_B.

Several properties of this product follow immediately from the compression axioms. It maps effects to effects (positivity of C_p). It satisfies $\mathbb{1} \cdot q = q$ (since $C_{\mathbb{1}}$ is the identity). On a simplex (classical state space), it reduces to the pointwise product $p \cdot q = (p_1 q_1, \ldots, p_n q_n)$, matching the unique sequential product on simplices [GG02].

Remark 3.1. The product (4) is defined entirely on V_B using B's own compression structure. The tracking map φ provides the physical justification—testing p on B causes the model M to undergo $C_{\varphi(p)}$, which feeds back to the context for testing q—but the algebraic product is intrinsic to B.

3.2 The Naive Extension and Its Failure

To extend the product from sharp to general effects, one attempts a spectral extension. Given $a = \sum_i \lambda_i \, p_i$ (spectral decomposition in the OUS), define:

$$a \cdot b \stackrel{?}{=} \sum_i \lambda_i \, C_{p_i}(b). \tag{5}$$

This is the natural extension: the effect a acts on b as a weighted sum of compressions, one for each spectral projector.

However, this extension *fails* axiom (S3) (unitality) on non-commutative order unit spaces. To see why, set $a = \mathbb{1}$ with spectral decomposition $\mathbb{1} = \sum_i p_i$ (all eigenvalues equal to 1). Then (5) gives

$$\mathbb{1} \cdot b = \sum_i C_{p_i}(b) = \mathrm{pinch}_{p_1,\dots,p_n}(b) \neq b \quad \text{in general.} \tag{6}$$

The sum of compressions is the *pinching map*, not the identity. By the Peirce decomposition (3), the pinching map annihilates all Peirce 1-space components: $P_{ij}(b)$ is discarded for every $i < j$. In the commutative case $V_1 = \{0\}$, the pinching map *is* the identity and the extension succeeds—this is why the classical limit works. In the non-commutative case, the Peirce 1-space is nontrivial, and the naive extension destroys off-diagonal information.

This failure is not a bug; it reveals a structural gap. The compression sum captures the Peirce 2-space (diagonal) contributions correctly but provides no prescription for the Peirce 1-space (off-diagonal) terms. Closing this gap requires new input.

3.3 The Corrected Product via Peirce Feedback

The missing Peirce 1-space contribution comes from the self-modeling feedback loop. We now show that the form of the corrected product is *forced* by the axioms, not chosen. Any binary operation on effects satisfying (S1) (additivity in the second argument), (S2) (continuity in the first), and the sharp constraint $p \cdot b = C_p(b)$ must respect the Peirce decomposition: the map $a \cdot \cdot$ is a linear endomorphism (by (S1) and finite-dimensionality) whose range on each Peirce subspace is constrained to that subspace (by the compression algebra structure). The diagonal coefficients are pinned to the eigenvalues by eigenvalue coalescence and (S2). Thus the general product for $a = \sum_i \lambda_i p_i$ must take the form:

$$a \cdot b = \sum_i \lambda_i C_{p_i}(b) + \sum_{i<j} f(\lambda_i, \lambda_j) P_{ij}(b), \tag{7}$$

where the first sum captures diagonal (Peirce 2-space) contributions via compressions, and the second sum restores the off-diagonal (Peirce 1-space) contributions through a *mixing function* $f : [0,1] \times [0,1] \to \mathbb{R}$ to be determined. The full derivation of this spectral extension theorem appears in the supplementary material.

The mixing function is constrained by the requirement that the product map effects to effects: $0 \le a \cdot b \le \mathbb{1}$ for all effects a, b. This yields a *positivity bound.*

Proposition 3.2 (Positivity bound)**.** *For the product* (7) *to map* $[0,1]_V \times [0,1]_V$ *into* $[0,1]_V$*, the mixing function must satisfy*

$$|f(\lambda_i, \lambda_j)| \le \sqrt{\lambda_i \lambda_j} \tag{8}$$

for all $\lambda_i, \lambda_j \in [0,1]$.

Proof. Every pair of orthogonal atoms p_i, p_j in a spectral OUS generates a two-level face isomorphic to a spin factor, on which the Schur complement criterion gives $|f(\lambda_i, \lambda_j)| \le \sqrt{\lambda_i \lambda_j}$. Face effects embed into global effects, so the bound lifts to the full space. The detailed computation appears in Appendix A. □

Any mixing function satisfying (8) with $f(1, \lambda) = \lambda$ and $f(0, \lambda) = 0$ yields a product that maps effects to effects and satisfies (S3). The question is: *which* f does self-modeling select?

3.4 Faithful Self-Modeling Selects the Maximum

The faithful self-modeling constraint of Assumption 2.7 selects the mixing function that saturates the positivity bound.

Proposition 3.3 (Maximal coherence)**.** *The mixing function saturates the positivity bound:*

$$f(\lambda_i, \lambda_j) = \sqrt{\lambda_i \lambda_j}. \tag{9}$$

Proof. Let f be any continuous mixing function satisfying the positivity bound (8). For compatible $a = \alpha_1 p_1 + \alpha_2 p_2$ and $b = \beta_1 p_1 + \beta_2 p_2$ sharing an eigenbasis, (S5) requires $a \cdot (b \cdot d) = (a \cdot b) \cdot d$ for all effects d. Comparing the Peirce 1-space coefficients of both sides gives the multiplicative functional equation

$$f(\alpha_1, \alpha_2)\, f(\beta_1, \beta_2) = f(\alpha_1\beta_1,\ \alpha_2\beta_2). \tag{10}$$

With the boundary conditions $f(1, \lambda) = \lambda$ and $f(0, \lambda) = 0$ (from (S3) and positivity), factorization gives $f(x, y) = g(x)\, h(y)$ where g and h are continuous multiplicative functions on $[0, 1]$. By the multiplicative Cauchy theorem on $\mathbb{R}_{>0}$, each is a power: $g(x) = x^s$, $h(y) = y^t$. The boundary condition $f(1, \lambda) = \lambda$ forces $t = 1 - s$; symmetry of the positivity bound under $(i, j) \leftrightarrow (j, i)$ gives $s = t$, hence $s = t = 1/2$. Therefore $f(\lambda_i, \lambda_j) = \sqrt{\lambda_i \lambda_j}$. □

The key step is the multiplicative functional equation (10): the mixing function must be *multiplicatively homomorphic* in the eigenvalue pairs. Combined with continuity ((S2)) and the boundary conditions, this forces the unique solution $f = \sqrt{\lambda_i \lambda_j}$. Note that maximal coherence is not an additional assumption: it is forced by the sequential product axioms.

The corrected product with $f = \sqrt{\lambda_i \lambda_j}$ is:

$$\boxed{a \cdot b = \sum_i \lambda_i\, C_{p_i}(b) + \sum_{i<j} \sqrt{\lambda_i \lambda_j}\, P_{ij}(b).} \tag{11}$$

Remark 3.4. On $M_2(\mathbb{C})^{\mathrm{sa}}$, the corrected product (11) coincides with the Lüders product $a \cdot b = \sqrt{a}\, b \sqrt{a}$. This equivalence is a *consequence* of the construction, not a premise. The Lüders formula was never assumed; it emerges from the positivity bound and faithful self-modeling.

Remark 3.5. The role of the tracking map φ is now clear. In the sharp product (4), φ entered only as physical interpretation. In the corrected product, φ is algebraically essential: setting $f = 0$ (trivial model, no feedback) recovers the failed naive extension, while setting $f = \sqrt{\lambda_i \lambda_j}$ (faithful model, maximal feedback) gives the corrected product. Self-modeling quality controls coherence.

The remaining structural assumptions are:

Assumption 3.6 (Minimal composite)**.** The body–model composite V_{BM} is the minimal order unit space carrying product states, product effects, non-signaling constraints, and a product-form sequential product.

Assumption 3.7 (Simple EJA)**.** The Euclidean Jordan algebra produced by Theorem 2.5 is simple (it has no nontrivial Jordan ideal).

3.5 Circularity Check

A natural concern is whether the scalar function $\sqrt{\lambda_i \lambda_j}$ in (11) imports C^*-algebraic structure through the back door. It does not. The square root in the mixing function operates on *real eigenvalues* $\lambda_i, \lambda_j \in [0,1]$: it is arithmetic, specifically $\sqrt{x} : [0,1] \to [0,1]$, the unique non-negative real number whose square is x. No operator square root, no trace, no inner product, no density matrix, and no Hilbert space structure is used in the construction. The only mathematical objects employed are:

- the order unit space $(V, V^+, \mathbb{1})$;
- the effect space $[0,1]_V$ and its projective units;
- Alfsen–Shultz compressions C_p for faces of projective units;
- the Peirce decomposition (3) derived from compressions;
- spectral decomposition of effects in the OUS;
- the scalar function $\sqrt{\cdot}$ on $[0,1] \subset \mathbb{R}$.

Complex numbers, Hilbert spaces, and C^*-algebraic operations do not appear until the conclusion (Section 6), where they emerge as consequences of the classification theorems.

4 Axiom Verification

We now verify that the self-modeling sequential product defined in Section 3 satisfies all seven axioms of van de Wetering's Definition 2 [vdW19b]. By Theorem 1 of [vdW19b], this establishes that the state space carries the structure of a Euclidean Jordan algebra.

Throughout this section, let V be a finite-dimensional spectral order unit space and let the sequential product be given by the corrected formula (11):

$$^{a}b = \sum_i \lambda_i \, C_{p_i}(b) + \sum_{i<j} \sqrt{\lambda_i \lambda_j} \, P_{ij}(b), \tag{12}$$

where $a = \sum_i \lambda_i p_i$ is the spectral decomposition of a, C_{p_i} is the Alfsen–Shultz compression for the face generated by p_i, and P_{ij} is the Peirce (i,j) projection onto the 1-space $V_1(p_i, p_j)$ as defined in (3).

S1 (Additivity)

The map $b \mapsto {}^a b$ is additive: ${}^a(b+c) = {}^a b + {}^a c$ whenever $b+c \leq \mathbf{1}$.

Each compression C_{p_i} is a positive linear map on V (Alfsen–Shultz [AS03], Ch. 7), and the Peirce projection P_{ij} is linear (each P_{ij} is a linear combination of compressions by (3)). Since the coefficients λ_i and $\sqrt{\lambda_i \lambda_j}$ depend only on a, the map $b \mapsto {}^a b$ is a fixed linear combination of linear maps applied to b, and therefore additive. □

S2 (Continuity)

The map $a \mapsto {}^a b$ is continuous in the order unit norm.

In a finite-dimensional spectral order unit space, the spectral decomposition $a = \sum_i \lambda_i(a)\, p_i(a)$ varies continuously with a (Alfsen–Shultz [AS03], Ch. 9; in finite dimensions, eigenvalues are continuous functions and the spectral functional calculus extends continuously through degeneracies). The scalar maps $\lambda_i \mapsto \lambda_i$ and $(\lambda_i, \lambda_j) \mapsto \sqrt{\lambda_i \lambda_j}$ are continuous on $[0, 1]$. Since the compressions C_{p_i} and Peirce projections P_{ij} depend continuously on the projectors p_i, the map $a \mapsto {}^a b$ is a composition of continuous maps and hence continuous. □

S3 (Unitality)

${}^{\mathbf{1}}a = a$ for all effects a.

The order unit $\mathbf{1}$ has the trivial spectral decomposition $\mathbf{1} = 1 \cdot \mathbf{1}$ (a single eigenvalue with the maximal projective unit). The Peirce 1-space sum is empty, and the compression $C_{\mathbf{1}}$ for the full face is the identity map on V (Alfsen–Shultz [AS03], Def. 7.1). Therefore ${}^{\mathbf{1}}a = 1 \cdot C_{\mathbf{1}}(a) = a$.

This is the axiom whose verification critically depends on the Peirce 1-space correction: the naive product $\sum_i \lambda_i\, C_{p_i}(b)$ (without the $\sqrt{\lambda_i \lambda_j}\, P_{ij}(b)$ terms) gives ${}^{\mathbf{1}}a = \mathrm{pinch}(a) \neq a$ for effects with off-diagonal components, because the sum of compressions over a resolution of unity is the pinching map, not the identity. The self-modeling feedback that introduces the Peirce 1-space terms is what restores unitality. □

S4 (Orthogonality Symmetry)

If ${}^a b = 0$ then ${}^b a = 0$.

This is the key axiom. The remaining axioms (S1–S3, S5–S7) follow from linearity and spectral calculus with comparatively short arguments. S4 requires a substantive proof that exploits the facial structure of the order unit space. We give the argument in full; see Appendix A.1 for additional details.

Theorem 4.1. *Let* $(V, \leq, \mathbf{1})$ *be a finite-dimensional spectral order unit space equipped with the sequential product* (12)*. If* ${}^{a}b = 0$ *for effects* $a, b \in [0, \mathbf{1}]_V$*, then* ${}^{b}a = 0$*.*

Proof sketch. Write $a = \sum_{i=1}^{n} \lambda_i \, p_i$ (spectral decomposition). Two cases arise.

Case A (full-rank). Suppose all $\lambda_i > 0$. The condition ${}^{a}b = 0$ and the direct sum structure of the Peirce decomposition

$$V = \bigoplus_i V_2(p_i) \;\oplus\; \bigoplus_{i<j} V_1(p_i, p_j)$$

(Alfsen–Shultz [AS03], Theorem 9.37) force each component to vanish individually: $\lambda_i \, C_{p_i}(b) = 0$ and $\sqrt{\lambda_i \lambda_j}\, P_{ij}(b) = 0$. Since $\lambda_i > 0$ for all i, the completeness of the Peirce decomposition gives $b = 0$, so ${}^{b}a = 0$ trivially.

Case B (rank-deficient). Suppose $\lambda_1, \ldots, \lambda_m > 0$ and $\lambda_{m+1} = \cdots = \lambda_n = 0$ for some $1 \leq m < n$. As before, the Peirce direct sum forces $C_{p_i}(b) = 0$ for each $i \leq m$. The crucial step invokes the Alfsen–Shultz facial absorption theorem (Proposition 7.43 of [AS03]):

If $C_p(b) = 0$ *and* $b \geq 0$*, then* $b \in \mathrm{face}(p^{\perp})$*.*

Applying this to each p_i ($i \leq m$) places b in the complementary face $\mathrm{face}\big(\sum_{k=m+1}^{n} p_k\big)$. By the facial structure of spectral order unit spaces, the Peirce 1-space components $V_1(p_i, p_j)$ connecting a face to its complement are excluded: an effect in $\mathrm{face}(p^{\perp})$ has zero component in every $V_1(p_i, p_j)$ for $i \leq m$, $j > m$. Therefore b is supported entirely on the zero-eigenvalue block of a.

For the reverse product ${}^{b}a$, write $b = \sum_j \mu_j \, q_j$ with each q_j ($\mu_j > 0$) lying in $\mathrm{face}\big(\sum_{k>m} p_k\big)$. Since a is supported on the complementary face $\mathrm{face}\Big(\sum_{i \leq m} p_i\Big)$, the facial orthogonality theorem gives $C_{q_j}(a) = 0$ for all j with $\mu_j > 0$. The Peirce 1-space terms $Q_{jk}(a)$ either vanish by facial structure (when both $\mu_j, \mu_k > 0$) or carry zero weight $\sqrt{\mu_j \mu_k} = 0$. Hence ${}^{b}a = 0$. □

Remark 4.2 (Independence of φ). Axiom S4 is φ-independent: it holds for *any* mixing function f with $f(0, x) = 0$, not just for the faithful choice $f = \sqrt{\lambda_i \lambda_j}$. The proof depends only on the facial geometry of the order unit space and the vanishing of f at zero eigenvalues, not on the specific form of f for positive arguments. Consequently, S4 is a structural property of spectral order unit spaces, not an artifact of the self-modeling construction.

S5 (Compatible Associativity)

> *If a and b are compatible (i.e., ${}^a b = {}^b a$), then ${}^a({}^b c) = {}^{({}^a b)} c$ for all effects c.*

Compatible effects share a joint spectral decomposition. For sharp compatible effects p, q, the compressions commute: $C_p \circ C_q = C_q \circ C_p$ (Alfsen–Shultz [AS03], Prop. 7.49), and their composition satisfies $C_p \circ C_q = C_{p \wedge q}$ (Prop. 7.50). For general compatible effects, simultaneous diagonalizability reduces the product to a function of shared spectral values, and the identity $\sqrt{\alpha\beta \cdot \gamma\delta} = \sqrt{\alpha\gamma} \cdot \sqrt{\beta\delta}$ for non-negative reals ensures compatibility of the mixing function. □

S6 (Additivity of Compatibility)

> *If $a \mid b$ then $a \mid (\mathbf{1} - b)$; if also $a \mid c$ then $a \mid (b + c)$ whenever $b + c \leq \mathbf{1}$.*

Both parts of S6 rest on two preliminary observations that reduce compatibility to a Peirce vanishing condition.

Lemma 4.3 (Peirce vanishing)**.** *Let $a = \sum_i \lambda_i p_i$ and $b = \sum_j \mu_j q_j$ be effects in a finite-dimensional spectral OUS. If $a \mid b$, then b has zero Peirce* 1*-space component in a's spectral decomposition:*

$$P_1^{(a)}(b) \;=\; b - \sum_i C_{p_i}(b) \;=\; 0, \tag{13}$$

and symmetrically $P_1^{(b)}(a) = 0$.

Proof. By Alfsen–Shultz [AS03] Prop. 7.49 and Niestegge [Nie08] Lemma 3.3, $a \mid b$ implies that the spectral projectors are pairwise compatible: $C_{p_i} \circ C_{q_j} = C_{q_j} \circ C_{p_i}$ for all i, j. By Prop. 7.50 of [AS03], $C_{p_i} \circ C_{q_j} = C_{p_i \wedge q_j}$. Summing over j:

$$\sum_j C_{q_j}(p_i) = \sum_j (p_i \wedge q_j) = p_i,$$

where the last equality uses the refinement property: since $\{q_j\}$ resolves $\mathbf{1}$ and each q_j is compatible with p_i, the meets $\{p_i \wedge q_j\}_j$ resolve p_i. Therefore $\sum_j C_{q_j}(a) = \sum_i \lambda_i \sum_j (p_i \wedge q_j) = a$, giving $P_1^{(b)}(a) = a - \sum_j C_{q_j}(a) = 0$. Interchanging a and b gives $P_1^{(a)}(b) = 0$. $\square$

Proposition 4.4 (Block-diagonal compatibility)**.** *If $P_1^{(a)}(e) = 0$ for an effect e and $a = \sum_i \lambda_i p_i$, then $a \mid e$.*

Proof. Since $e = \sum_i C_{p_i}(e)$ with each $C_{p_i}(e) \in \text{face}(p_i)$, the spectral decomposition of e is the union of spectral decompositions within each face: write $C_{p_i}(e) = \sum_\ell \delta_{i\ell}\, s_{i\ell}$ with $s_{i\ell} \leq p_i$. Each $s_{i\ell}$ satisfies $C_{p_{i'}} \circ C_{s_{i\ell}} = C_{s_{i\ell}} \circ C_{p_{i'}}$ for all i' (they commute trivially: $s_{i\ell} \leq p_i$ gives $C_{s_{i\ell}} \circ C_{p_i} = C_{s_{i\ell}}$, while $s_{i\ell} \perp p_{i'}$ for $i' \neq i$ gives $C_{s_{i\ell}} \circ C_{p_{i'}} = 0$, both symmetric under composition reversal). The corrected product formula then gives, using $C_{s_{i\ell}}(a) = \lambda_i\, s_{i\ell}$:

$$^e a = \sum_{i,\ell} \delta_{i\ell}\, C_{s_{i\ell}}(a) = \sum_{i,\ell} \delta_{i\ell}\, \lambda_i\, s_{i\ell} = \sum_i \lambda_i\, C_{p_i}(e) =^a e,$$

where the Peirce 1-space terms vanish on both sides. $\square$

Part (i): $a \mid b \implies a \mid (\mathbf{1} - b)$. By Lemma 4.3, $P_1^{(a)}(b) = 0$. Since $P_1^{(a)}(\mathbf{1}) = 0$ (the unit lies in the Peirce 2-space) and $P_1^{(a)}$ is linear:

$$P_1^{(a)}(\mathbf{1} - b) = P_1^{(a)}(\mathbf{1}) - P_1^{(a)}(b) = 0.$$

By Proposition 4.4, $a \mid (\mathbf{1} - b)$. $\square$

Part (ii): $a \mid b$ and $a \mid c \implies a \mid (b + c)$ when $b + c \leq \mathbf{1}$. By Lemma 4.3, $P_1^{(a)}(b) = 0$ and $P_1^{(a)}(c) = 0$. By linearity:

$$P_1^{(a)}(b + c) = P_1^{(a)}(b) + P_1^{(a)}(c) = 0.$$

By Proposition 4.4, $a \mid (b + c)$. $\square$

S7 (Multiplicativity of Compatibility)

If $a \mid b$ and $a \mid c$ then $a \mid^b c$.

We show that $^b c$ has zero Peirce 1-space component in a's decomposition, then invoke Proposition 4.4.

Write $a = \sum_i \lambda_i\, p_i$, $b = \sum_j \mu_j\, q_j$, $c = \sum_k \nu_k\, r_k$. By Lemma 4.3, $a \mid b$ gives $C_{p_i} \circ C_{q_j} = C_{q_j} \circ C_{p_i}$ for all i, j, and $a \mid c$ gives $C_{p_i} \circ C_{r_k} = C_{r_k} \circ C_{p_i}$ for all i, k.

Step 1: C_{p_i} commutes with every Peirce projection of b's resolution. The Peirce (j, j') projection for b's resolution is $Q_{jj'} = C_{q_j+q_{j'}} - C_{q_j} - C_{q_{j'}}$ (a linear combination of compressions). Since p_i is compatible with both q_j and $q_{j'}$, Prop. 7.50 of [AS03] gives $C_{p_i} \circ C_{q_j+q_{j'}} = C_{p_i \wedge (q_j+q_{j'})} = C_{(q_j+q_{j'}) \wedge p_i} = C_{q_j+q_{j'}} \circ C_{p_i}$. Therefore C_{p_i} commutes with $Q_{jj'} = C_{q_j+q_{j'}} - C_{q_j} - C_{q_{j'}}$.

Step 2: Compute $\sum_i C_{p_i}({}^b c)$. Expanding ${}^b c$ via (12) and using linearity:

$$\sum_i C_{p_i}\big({}^b c\big) = \sum_j \mu_j\, C_{q_j}\Big(\sum_i C_{p_i}(c)\Big) + \sum_{j<j'} \sqrt{\mu_j \mu_{j'}}\, Q_{jj'}\Big(\sum_i C_{p_i}(c)\Big),$$

where we used Step 1 to pull $\sum_i C_{p_i}$ through the maps C_{q_j} and $Q_{jj'}$. By Lemma 4.3, $a \mid c$ gives $\sum_i C_{p_i}(c) = c$. Therefore

$$\sum_i C_{p_i}\big({}^b c\big) = \sum_j \mu_j\, C_{q_j}(c) + \sum_{j<j'} \sqrt{\mu_j \mu_{j'}}\, Q_{jj'}(c) =^b c.$$

Conclusion. $P_1^{(a)}\big({}^b c\big) =^b c - \sum_i C_{p_i}\big({}^b c\big) = 0$. By Proposition 4.4, $a \mid^b c$. □

With S1–S7 verified, we invoke the central classification result:

Theorem 4.5 (EJA classification; [vdW19b, Theorem 1])**.** *Let $(V, \leq, \mathbf{1})$ be a finite-dimensional order unit space admitting a sequential product satisfying S1–S7. Then V is order-isomorphic to a Euclidean Jordan algebra.*

By the Jordan–von Neumann–Wigner classification [JvNW34], the simple Euclidean Jordan algebras are: $M_n(\mathbb{R})^{\mathrm{sa}}$, $M_n(\mathbb{C})^{\mathrm{sa}}$, $M_n(\mathbb{H})^{\mathrm{sa}}$, the spin factors V_n, and the exceptional Albert algebra $M_3(\mathbb{O})^{\mathrm{sa}}$.

Remark 4.6 (Non-associativity)*.* The sequential product ${}^a b$ is non-associative for general (non-compatible) effects: there exist triples (a, b, c) with ${}^{({}^a b)}c \neq^a ({}^b c)$. This is consistent with S5, which requires associativity only for compatible pairs. By the theorem of Westerbaan, Westerbaan, and van de Wetering [WWvdW20], associativity would force commutativity, which would force the state space to be a simplex (classical). Non-associativity is thus a necessary feature of any non-classical sequential product.

5 Composite Systems and Local Tomography

5.1 The Body–Model Composite

Having established that the state space V of the body B (and isomorphically the model M) carries EJA structure, we now formalize the composite system

V_{BM} using only generalized probabilistic theory (GPT) primitives [Pla23, BUvdW23].

Definition 5.1 (Composite OUS)**.** Let V_B and V_M be finite-dimensional order unit spaces, each equipped with a sequential product satisfying S1–S7. The *composite order unit space* V_{BM} is defined by:

(C1) **Vector space.** V_{BM} is a finite-dimensional real vector space containing $V_B \otimes V_M$ (the algebraic tensor product of real vector spaces) as a subspace.

(C2) **Order unit.** $\mathbf{1}_{BM} = \mathbf{1}_B \otimes \mathbf{1}_M$.

(C3) **Product states.** For every state ω_B on V_B and ω_M on V_M, the product state $\omega_B \otimes \omega_M$ belongs to $\mathcal{S}(V_{BM})$, defined by $(\omega_B \otimes \omega_M)(a \otimes b) = \omega_B(a)\, \omega_M(b)$.

(C4) **Non-signaling.** For every joint state $\omega \in \mathcal{S}(V_{BM})$, the marginals $\omega_B(a) \coloneqq \omega(a \otimes \mathbf{1}_M)$ and $\omega_M(b) \coloneqq \omega(\mathbf{1}_B \otimes b)$ are well-defined states, independent of what is measured on the other subsystem.

Assumption 5.2 (Minimal composite)**.** V_{BM} is the *minimal* order unit space satisfying (C1)–(C4) and the product-form sequential product (Definition 5.3 below).

This construction uses only real vector spaces, compact convex state spaces, affine functionals, and the non-signaling constraint — no Hilbert space tensor products, complex numbers, or density matrices.

5.2 Product-Form Sequential Product

Definition 5.3 (Product-form SP)**.** The sequential product on V_{BM} is defined on product effects by

$$^{(a\otimes b)}(c \otimes d) = (^{a}c) \otimes (^{b}d), \tag{14}$$

where $^{\cdot}\cdot$ on the right-hand side denotes the factor-level sequential product (12).

Proposition 5.4 (S1–S7 inheritance)**.** *If the factor sequential products on V_B and V_M each satisfy S1–S7, then the product-form sequential product on V_{BM} satisfies S1–S7.*

Proof sketch. Axioms S1 (additivity), S2 (continuity), and S3 (unitality) follow directly from the corresponding factor-level axioms and the bilinearity of the tensor product. Axioms S5–S7 (compatible associativity, additivity, multiplicativity) reduce to the factor-level axioms because compatibility of product effects $a \otimes b$ and $c \otimes d$ follows from factor-wise compatibility ${}^{a}c ={}^{c} a$ and ${}^{b}d ={}^{d} b$.

The inheritance of S4 (orthogonality symmetry) is more subtle. If ${}^{(a\otimes b)}(c \otimes d) = 0$, then $({}^{a}c) \otimes ({}^{b}d) = 0$. In an order unit space, states separate points (Alfsen–Shultz [AS03], Theorem 1.23): for product states $\omega_B \otimes \omega_M$, the evaluation $\omega_B({}^{a}c)\,\omega_M({}^{b}d) = 0$ for all ω_B, ω_M forces either ${}^{a}c = 0$ or ${}^{b}d = 0$. By the factor-level S4, the corresponding reverse products vanish, giving ${}^{(c\otimes d)}(a \otimes b) = 0$. $\square$

Remark 5.5 (Bootstrapping order). The verification of S1–S7 on the composite avoids a potential circularity. The product-form sequential product (14) is first defined on the algebraic tensor product $V_B \otimes V_M$ (not on all of V_{BM}) by bilinear extension. Axioms S1–S7 are verified on this subspace using only the factor-level axioms and bilinearity—no reference to the full composite V_{BM} is needed. By Theorem 4.5, the algebraic tensor product equipped with this product carries EJA structure. The trace form non-degeneracy argument (Proposition 5.8) then shows that the product effects $\{a_i \otimes b_j\}$ are linearly independent, giving $\dim(V_{BM}) \geq \dim(V_B) \cdot \dim(V_M)$. Minimality (Assumption 5.2) supplies the upper bound, yielding $V_{BM} = V_B \otimes V_M$.

5.3 From Faithful Tracking to Local Tomography

The key step is showing that faithful self-modeling forces the composite to be locally tomographic: product measurements suffice to determine all joint states.

Assumption 5.6 (Simple EJA)**.** The EJA V is *simple* (not a nontrivial direct sum of EJAs).

This assumption is needed for the non-degeneracy of the trace form below.

Definition 5.7 (Correlation bilinear form)**.** Let τ denote the normalized trace functional on the EJA V (Faraut–Korányi [FK94], Ch. III) and let $\circ$ denote the Jordan product. Define the *correlation bilinear form*

$$B(a, b) \;=\; \tau\big(a \circ \varphi^{-1}(b)\big)\,, \tag{15}$$

where $\varphi^{-1} \colon V_M \to V_B$ is the inverse of the tracking isomorphism.

Proposition 5.8 (Non-degeneracy)**.** *On a simple EJA, the bilinear form $B(a, b)$ is non-degenerate.*

Proof. On a simple EJA, the trace form $(a, c) \mapsto \tau(a \circ c)$ is non-degenerate (Faraut–Korányi [FK94], Proposition III.4.2). Since $\varphi\colon V_B \to V_M$ is an isomorphism (faithful self-modeling), φ^{-1} is a bijection, so $B(a, b) = \tau(a \circ \varphi^{-1}(b))$ is non-degenerate: if $B(a, b) = 0$ for all a, then $\varphi^{-1}(b) = 0$, hence $b = 0$. □

Theorem 5.9 (Local tomography)**.** *Under Assumptions 5.2 and 5.6,*

$$\dim(V_{BM}) \;=\; \dim(V_B) \cdot \dim(V_M).$$

Proof sketch. The non-degeneracy of B implies that the product effects $\{a_i \otimes b_j\}$, where $\{a_i\}$ and $\{b_j\}$ are bases for V_B and V_M respectively, are linearly independent in V_{BM}. Indeed, if $\sum_{i,j} \alpha_{ij}\, a_i \otimes b_j = 0$, then for any fixed a, the map $b \mapsto \sum_j \alpha_{ij}\, B(a, b_j) = 0$ forces all $\alpha_{ij} = 0$ by non-degeneracy.

This gives the lower bound $\dim(V_{BM}) \geq \dim(V_B) \cdot \dim(V_M)$. By minimality (Assumption 5.2), the composite contains no structure beyond what is forced by axioms (C1)–(C4) and the product-form SP. Since the product effects span a $\dim(V_B) \cdot \dim(V_M)$-dimensional subspace, and the SP preserves this subspace (product effects map to product effects under (14)), minimality gives the upper bound $\dim(V_{BM}) \leq \dim(V_B) \cdot \dim(V_M)$. □

5.4 The Entangled Sector

The *entangled sector* consists of states in V_{BM} not reachable as convex combinations of product states. Minimality (Assumption 5.2) eliminates this sector by construction: V_{BM} is defined as the smallest OUS satisfying the composite axioms, so it contains no "hidden" entangled structure beyond what the sequential product and non-signaling constraints require.

For complex quantum mechanics ($V = M_n(\mathbb{C})^{\mathrm{sa}}$), the minimal and maximal composites coincide [BUvdW23]: the standard Hilbert space tensor product $M_{nm}(\mathbb{C})^{\mathrm{sa}}$ is the unique composite satisfying (C1)–(C4). This is why minimality is not a restrictive assumption for the complex case.

For real or quaternionic quantum mechanics, the minimal and maximal composites differ. This discrepancy manifests as a failure of local tomography and is the mechanism by which these types are excluded:

Remark 5.10 (Negative checks). Dimension counting reveals that local tomography fails for all non-complex types at $n = 2$:

Type	$d = \dim(V)$	d^2	dim(composite)	Local tomography?
$M_2(\mathbb{R})^{\text{sa}}$	3	9	10	No $(9 < 10)$
$M_2(\mathbb{C})^{\text{sa}}$	4	16	16	Yes $(16 = 16)$
$M_2(\mathbb{H})^{\text{sa}}$	6	36	28	No $(36 > 28)$

The real composite is "too large" (product measurements cannot distinguish all joint states), while the quaternionic composite is "too small" (product measurements are linearly dependent on the joint system). Only the complex type achieves the exact match $d^2 = \dim(\text{composite})$.

6 Type Exclusion and C^*-Algebra Promotion

6.1 The Jordan–von Neumann–Wigner Classification

By Theorem 4.5, the state space V is a Euclidean Jordan algebra (EJA). The Jordan–von Neumann–Wigner classification [JvNW34] asserts that every finite-dimensional simple EJA belongs to one of five families:

Type	Notation	dim	LT?	Exclusion mechanism
Real	$M_n(\mathbb{R})^{\text{sa}}$	$n(n+1)/2$	Only $n=1$	Dimension mismatch
Complex	$M_n(\mathbb{C})^{\text{sa}}$	n^2	All n	—
Quaternionic	$M_n(\mathbb{H})^{\text{sa}}$	$n(2n-1)$	Only $n=1$	Dimension mismatch
Spin factor	V_n $(n \geq 3)$	$n+1$	Only V_3	Barnum–Wilce
Albert	$M_3(\mathbb{O})^{\text{sa}}$	27	Never	No composite exists

The spin factors overlap with the matrix types at low dimensions: $V_2 \cong M_2(\mathbb{R})^{\text{sa}}$, $V_3 \cong M_2(\mathbb{C})^{\text{sa}}$, and $V_5 \cong M_2(\mathbb{H})^{\text{sa}}$.

6.2 Excluding Non-Complex Types

We now show that local tomography (Theorem 5.9) excludes every non-complex type.

Real matrices, $M_n(\mathbb{R})^{\text{sa}}$ **(**$n \geq 2$**).** The local tomography condition $[\dim(V)]^2 = \dim(\text{composite})$ becomes

$$\left[\tfrac{n(n+1)}{2}\right]^2 = \tfrac{n^2(n^2+1)}{2}\,.$$

Simplifying: $(n+1)^2 = 2(n^2+1)$, which reduces to $(n-1)^2 = 0$, forcing $n = 1$ (trivial, one-dimensional). For $n \geq 2$, the real composite is strictly larger than the product-effect span: $d^2 < \dim(\text{composite})$, and product measurements leave an entangled sector unresolved [BW14].

Quaternionic matrices, $M_n(\mathbb{H})^{\text{sa}}$ ($n \geq 2$). The analogous condition yields $[n(2n-1)]^2 = n^2(2n^2-1)$, which again reduces to $(n-1)^2 = 0$, forcing $n = 1$ (trivial). For $n \geq 2$, the quaternionic composite has *fewer* dimensions than d^2: product effects are linearly dependent on the joint system [BW14].

Spin factors, V_n ($n \geq 4$). The cross-identifications $V_2 \cong M_2(\mathbb{R})^{\text{sa}}$ and $V_5 \cong M_2(\mathbb{H})^{\text{sa}}$ are already excluded above. For V_n with $n \geq 4$ and $n \neq 5$, these are "pure" spin factors not isomorphic to any $M_k(\mathbb{K})^{\text{sa}}$. Barnum and Wilce [BW14] prove that no locally tomographic composite exists for these types (Theorem 4.1 of [BW14]).

Albert algebra, $M_3(\mathbb{O})^{\text{sa}}$. The exclusion here is the strongest: Barnum, Graydon, and Wilce [BGW20] prove that no non-signaling composite with product states exists for any EJA containing an Albert algebra summand (Theorem 1 of [BGW20]). The 27-dimensional exceptional Jordan algebra cannot form *any* consistent bipartite system, even a non-locally-tomographic one.

Complex matrices, $M_n(\mathbb{C})^{\text{sa}}$ ($n \geq 1$). $\dim(M_n(\mathbb{C})^{\text{sa}}) = n^2$, and the composite $M_{n^2}(\mathbb{C})^{\text{sa}}$ has dimension $(n^2)^2 = n^4 = (n^2)^2$. Local tomography holds for all n: the dimension formula $d = n^2$ is *multiplicatively stable*. The complex type is the unique survivor.

6.3 C*-Algebra Promotion

Three published theorems, each with independently verifiable hypotheses, promote the surviving EJA to a C*-algebra.

Theorem 6.1 (C*-algebra structure; [vdW19b, Theorem 3]). *Let V be a sequential product space satisfying S1–S7. If the composite $V \otimes V$, equipped with the product-form sequential product, is also a sequential product space and the composite is locally tomographic, then V is the self-adjoint part of a C*-algebra.*

Hypothesis	Status	Established in
V satisfies S1–S7	Verified	Section 4
$V \otimes V$ with product-form SP satisfies S1–S7	Verified	Proposition 5.4
Composite is locally tomographic	Verified	Theorem 5.9

Theorem 6.2 (Complex type identification; [BW14])**.** *Let $(V, \cdot)$ be a finite-dimensional EJA embedded as a system type in a locally tomographic dagger-compact probabilistic theory containing a qubit ($V_3 \cong M_2(\mathbb{C})^{\text{sa}}$) as a subsystem. Then $V \cong M_n(\mathbb{C})^{\text{sa}}$ for some $n \geq 1$.*

The qubit hypothesis requires care: the type exclusions above have already eliminated real ($n \geq 2$), quaternionic ($n \geq 2$), spin factors ($n \geq 4$, $n \neq 3$), and the Albert algebra. The only surviving simple EJA types are $\mathbb{R}$ (trivial) and $M_n(\mathbb{C})^{\text{sa}}$ for $n \geq 2$. For $n \geq 2$, any orthogonal pair $\{p, p^{\perp}\}$ of projective units defines a two-level face that is isomorphic to $V_3 \cong M_2(\mathbb{C})^{\text{sa}}$ — this identification is valid precisely because the non-complex types have already been excluded. The categorical structure (dagger-compact probabilistic theory) is provided by the composite V_{BM} with its product-form sequential product and the dagger given by the involution exhibited below.

Hypothesis	**Status**	**Established in**
V is a finite-dimensional EJA	Verified	Theorem 4.5
Locally tomographic composite	Verified	Theorem 5.9
Qubit subsystem V_3	Verified	Type exclusion above (only complex survives; two-level faces are V_3)

Theorem 6.3 (Jordan-to-C* promotion; [HO85])**.** *Let E be a JB-algebra. If $E \otimes E_2$ can be given the structure of a JB-algebra such that the Jordan product satisfies the factorization identities $(a \otimes 1) \circ (b \otimes 1) = (a \circ b) \otimes 1$ and $(1 \otimes x) \circ (1 \otimes y) = 1 \otimes (x \circ y)$, where $E_2 = M_2(\mathbb{C})^{\text{sa}}$, then E is the self-adjoint part of a C*-algebra.*

The hypotheses of Theorem 6.3 are satisfied: V is a finite-dimensional EJA (hence a JB-algebra), and the qubit subsystem $E_2 = M_2(\mathbb{C})^{\text{sa}}$ is available by the type identification above. The composite V_{BM} with its product-form sequential product provides the JB-algebra tensor product, and the product-form condition ${}^{(a\otimes b)}(c \otimes d) = ^{a} c \otimes^{b} d$ implies the required Jordan-product factorization identities. This theorem is consistent with Theorems 6.1 and 6.2 and provides an independent route to the C*-algebra conclusion.

6.4 The Involution

Theorems 6.1–6.2 establish that $V = M_n(\mathbb{C})^{\text{sa}}$ is the self-adjoint part of the C*-algebra $A = M_n(\mathbb{C})$. The C*-algebra involution is the conjugate transpose:

$$X^* = X^\dagger = \overline{X}^T.$$

This involution satisfies four defining properties:

(P1) **Involutive:** $(X^*)^* = X$.

(P2) **Anti-multiplicative:** $(XY)^* = Y^*X^*$.

(P3) **Fixed-point set:** $\{X \in M_n(\mathbb{C}) : X^* = X\} = M_n(\mathbb{C})^{\mathrm{sa}}$.

(P4) **C*-identity:** $\|X^*X\| = \|X\|^2$ (operator norm).

Property (P3) identifies the self-adjoint part recovered from the self-modeling construction with the standard notion of Hermitian matrices. Property (P4) equips $M_n(\mathbb{C})$ with its C*-norm, completing the C*-algebra structure.

6.5 Main Theorem

We can now state the central result.

Theorem 6.4 (Self-modeling implies complex quantum mechanics)**.** *Let $(V, \leq, \mathbf{1})$ be a finite-dimensional spectral order unit space. If V admits a faithful self-model $\varphi\colon V \to V_M$ with V_M an isomorphic copy, and if the minimal composite V_{BM} satisfies non-signaling and product-form sequential product inheritance, then:*

1. *$V \cong M_n(\mathbb{C})^{\mathrm{sa}}$ for some $n \geq 2$;*
2. *the sequential product is the Lüders product $^a b = \sqrt{a}\, b\, \sqrt{a}$;*
3. *$M_n(\mathbb{C})$ carries the conjugate-transpose involution $X^* = X^\dagger$.*

In particular, the complex field, the Jordan product, the C*-involution, and local tomography are all *derived* from the single operational premise of faithful self-modeling.

Proof outline.

1. The self-modeling constraint yields a sequential product on V via the "test–update–test" operational semantics (Section 3).
2. The corrected product formula (11), incorporating the Peirce 1-space feedback selected by faithful tracking, satisfies S1–S7 (Section 4).
3. By Theorem 4.5 (van de Wetering), V is an EJA.

4. The composite V_{BM} with the product-form SP inherits S1–S7 (Proposition 5.4), and the non-degeneracy of the EJA trace form combined with minimality yields local tomography (Theorem 5.9).

5. Dimension counting excludes all non-complex simple EJA types (Section 6.2).

6. Theorem 6.1 (C*-promotion), Theorem 6.2 (complex type), and Theorem 6.3 (Jordan-to-C* consistency) yield $V = M_n(\mathbb{C})^{\mathrm{sa}}$ with the conjugate-transpose involution. □

□

7 Discussion

7.1 Summary of Results

We have shown that a finite-dimensional spectral order unit space admitting a faithful self-model—an isomorphic internal copy through which the system can operationally probe and update its own state—is necessarily described by the algebra $M_n(\mathbb{C})^{\mathrm{sa}}$ for some $n \geq 2$, equipped with the Lüders sequential product ${}^a b = \sqrt{a}\, b \sqrt{a}$ and the conjugate-transpose involution $X^* = X^\dagger$. The derivation proceeds through a chain of eight implications (see Figure 1): self-modeling yields a sequential product on the effect space; the product satisfies the seven axioms of van de Wetering, forcing Euclidean Jordan algebra structure; the faithful tracking constraint implies local tomography for the body–model composite; local tomography excludes all non-complex Jordan types; and the surviving algebra promotes to a C^*-algebra via published theorems of van de Wetering [vdW19b], Barnum–Wilce [BW14], and Hanche-Olsen [HO85].

The complex field, the Jordan product, the C^*-involution, and local tomography are all *consequences* of self-modeling, not independent postulates. This reduces the premise count from the 3–6 axioms typical of competing reconstruction programs to a single operational premise, supplemented by four standing structural assumptions which we now discuss.

7.2 Standing Assumptions

Our derivation relies on one operational premise (faithful self-modeling) and four structural assumptions. We analyze each assumption's role, restrictiveness, and the consequences of weakening it.

Assumption 2.2: Finite-dimensional spectral OUS. We argue in Section 7.3 that this follows from finite capacity of the embedding region. It is the most restrictive assumption in the formal sense. The spectral decomposition $a = \sum_i \lambda_i p_i$, which underlies the entire product construction (7), requires a well-behaved spectral theory. In finite dimensions, this is guaranteed by the Alfsen–Shultz theory [AS03]. In infinite dimensions, spectral theory introduces domain issues: the spectral measure replaces the finite sum, unbounded effects arise, and the Peirce decomposition becomes a conditional-expectation-valued integral. We expect the core argument to extend to type I factors (direct integrals of matrix algebras), where the spectral decomposition retains its essential features. Extension to type II or III von Neumann algebras would require substantially new techniques, particularly for the Peirce feedback construction and the facial orthogonality argument for (S4).

Assumption 2.7: Faithful self-model. Faithfulness means $\varphi : V_B \to V_M$ is an order isomorphism—the model is a perfect internal copy. This is what selects the maximal mixing function $f = \sqrt{\lambda_i \lambda_j}$ in **??**. An approximate self-model, where φ is injective but not surjective, or is merely ε-isometric, would yield $|f| < \sqrt{\lambda_i \lambda_j}$ without saturating the positivity bound (8). The resulting product would be a "partially coherent" sequential product—intermediate between the classical (pinching) limit $f = 0$ and the fully quantum $f = \sqrt{\lambda_i \lambda_j}$. Such products may describe systems with intrinsic decoherence or limited self-knowledge. Exploring this landscape of approximate self-models is a natural direction for future work.

Assumption 3.6: Minimal composite. We argue in Section 7.3 that minimality is the natural composite for a maximally efficient self-modeler. The body–model composite V_{BM} is defined as the minimal order unit space carrying product states, product effects, non-signaling constraints, and a product-form sequential product. For complex quantum mechanics, the minimal and maximal composites coincide: the Hilbert space tensor product $M_{nm}(\mathbb{C})^{\mathrm{sa}}$ is the unique composite satisfying (C1)–(C4) [BUvdW23]. This is a feature of the complex case, not a coincidence: it means Assumption 3.6 is not restrictive for the theory we derive. For a priori exploration, one could ask whether maximal composites yield different conclusions. For the real and quaternionic cases, minimal $\neq$ maximal: the difference manifests as a failure of local tomography (the real composite has "hidden entangled states" that product measurements cannot resolve). Thus the minimal composite assumption is precisely what enforces local tomography, which in turn excludes the non-complex types.

Assumption 3.7: Simple EJA. We argue in Section 7.3 that simplicity corresponds to "whole models whole"—the self-modeler models all of itself, not independent parts. The trace form non-degeneracy argument in Proposition 5.8 requires the EJA to be simple—having no nontrivial Jordan ideal. An EJA with a direct sum decomposition $V = V_1 \oplus V_2$ would have a degenerate trace form on the whole space (the cross terms $\tau(a_1 \circ a_2) = 0$ for $a_1 \in V_1$, $a_2 \in V_2$). However, each simple summand satisfies our theorem independently: if V_1 and V_2 each admit faithful self-models, they are each $M_{n_i}(\mathbb{C})^{\mathrm{sa}}$ by the main theorem applied to each summand. The direct sum then describes a system with superselection sectors, each governed by complex quantum mechanics. Thus Assumption 3.7 restricts to irreducible systems, which is the generic case.

7.3 The Assumptions as Self-Modeling

The four structural assumptions are presented above as independent conditions, and the honest accounting is: one operational premise plus four structural assumptions. We now argue that this accounting, while mathematically correct, obscures a deeper unity. The four assumptions are not independent design choices drawn from a menu of alternatives. They are what self-modeling *means* when unpacked into the language of order unit spaces.

Finite capacity. A self-modeler that contains an isomorphic copy of itself within itself must have finite state-space dimension. An infinite-dimensional system would require an infinite-dimensional internal model, but a subsystem embedded in a finite spatial region has finite information capacity [**?**]. The spectral structure (distinguishable states, resolutions of unity) is not an abstract regularity condition: it is the minimal apparatus for *probing*—a system that tests its own state requires states to distinguish. In this sense, Assumption 2.2 is not a restriction on the formalism but a structural consequence of what "can probe itself in a finite region" means operationally.

Faithfulness. This is the operational premise itself. No further argument is needed.

Efficiency. Assumption 3.6 (minimal composite) says the body–model composite carries exactly the structure needed for faithful tracking and no more. A self-modeler at the fixed point of a self-modeling basin—the most efficient tracker—has no reason to maintain excess degrees of freedom. Extra composite structure beyond what product measurements require is dead weight: it contributes nothing to tracking fidelity and increases the system's complexity without improving its self-model. We do not prove that self-modeling forces minimality, but we note that minimality is the composite structure one would *expect* of a maximally efficient self-modeler, and that for the theory we derive (complex quantum mechanics) the minimal and maximal composites coincide [BUvdW23].

Irreducibility. Assumption 3.7 (simple EJA) says the algebra has no nontrivial ideal—it cannot be decomposed into independent blocks. A self-modeler that decomposes into independent subsystems is not modeling *all* of itself: each block models only itself, not the whole. Simplicity is the algebraic expression of "whole models whole." As noted in Section 7.2, each simple summand independently satisfies the main theorem, so the assumption restricts to irreducible systems rather than excluding composite ones.

We emphasize that the above is a philosophical argument, not a derivation. The mathematical theorem requires the four assumptions as stated. But the assumptions are not arbitrary: they are the structural shape that self-modeling takes in a finite-capacity system.

The following observation makes the relationship precise.

Corollary 7.1 (Characterization)**.** *Among simple finite-dimensional EJAs, a type admits a faithful self-model whose minimal composite is locally tomographic if and only if it is* $M_n(\mathbb{C})^{\mathrm{sa}}$ *for some* $n \geq 2$*.*

Proof. The forward direction is the main theorem. For the converse, take $V_M = V_B = M_n(\mathbb{C})^{\mathrm{sa}}$ with $\varphi = \mathrm{id}$. The minimal composite of $M_n(\mathbb{C})^{\mathrm{sa}}$ with itself is $M_{n^2}(\mathbb{C})^{\mathrm{sa}}$, which satisfies $\dim = n^4 = (n^2)^2 = [\dim(V_B)]^2$. Local tomography holds, and the product-form sequential product inherits S1–S7 from the factors. □

Remark 7.2 (Why the other types fail)*.* The converse fails for every non-complex simple EJA type. For real matrices ($n \geq 2$), the minimal composite has $\dim > [\dim(V)]^2$: product measurements leave an entangled sector unresolved, so the self-model has *blind spots* about body–model correlations. For quaternionic matrices ($n \geq 2$), $\dim < [\dim(V)]^2$: product effects are linearly dependent, so the self-model's observations are *redundant*. For spin factors and the Albert algebra, no locally tomographic composite exists at all. Complex quantum mechanics is the unique theory in which a finite-dimensional system can have complete self-knowledge through product measurements on the body–model composite.

Corollary 7.1 says that the conditions of our theorem are not merely sufficient for complex quantum mechanics—they *characterize* it. There is no simple finite-dimensional EJA satisfying our conditions that is not complex QM, and no complex QM system that fails to satisfy them. The "axiomatic corridor" identified by the theorem is, in this precise sense, the corridor of self-modeling itself.

7.4 Comparison with Other Programs

Table 1 provides a summary comparison. We now discuss several programs in more depth.

Masanes–Müller (2011–2013). The Masanes–Müller program [MM11, MMAPG13] achieves the closest rival in parsimony, using three postulates: the existence of an information unit, continuous reversibility, and tomographic locality. The key difference is that their tomographic locality postulate assumes local tomography (effectively selecting the complex field as input), whereas we *derive* local tomography from faithful tracking via the EJA trace form non-degeneracy argument (Theorem 5.9). Their continuous reversibility axiom—that every reversible transformation can be continuously connected to the identity—has no analogue in our framework; instead, the reversibility structure emerges from the C^*-algebra promotion. On the other hand, their framework handles infinite dimensions more naturally than ours, and their postulates have a cleaner operational interpretation in terms of information processing.

Chiribella–D'Ariano–Perinotti (2011). The CDP program [CDP11] uses six axioms, with Purification as the central workhorse: every mixed state of a system arises as the marginal of a pure state of a larger system. Purification is a powerful constraint that immediately selects complex quantum mechanics from the landscape of GPTs. We derive the C^*-algebraic structure without assuming purification; purification then follows as a *theorem* of $M_n(\mathbb{C})^{\mathrm{sa}}$, not a postulate. However, the CDP framework provides a complete categorical description of quantum theory (including channels, instruments, and higher-order maps) that goes well beyond our algebraic result.

Van de Wetering's effectus theory (2019). Our work relies heavily on van de Wetering's theorems [vdW19b], but the relationship to his effectus-theoretic reconstruction [vdW19a] is complementary rather than subsumptive. The effectus approach uses approximately five axioms on the categorical structure of a theory (effect algebra homomorphisms, images, compressions, quotients, sharp effects). Our approach is more operational: we replace categorical axioms with a single physical premise (self-modeling) and derive the sequential product structure that feeds into van de Wetering's classification theorems. The two programs share a common downstream engine (Theorems 1 and 3 of [vdW19b]) but differ in what they take as input.

Alfsen–Shultz geometry. The Alfsen–Shultz program [AS03] characterizes C^*-algebras via the geometry of state spaces, using the concept of an *orientation* on the state space. Their characterization is geometrically elegant: a JB-algebra is the self-adjoint part of a C^*-algebra if and only if its state space admits a consistent orientation (a choice of "which way around" each face). In our framework, the orientation is not an input but emerges from the self-modeling construction: the tracking map φ provides a canonical correspondence between faces of V_B and faces of V_M that implicitly defines an orientation. Making this connection precise is a direction for future investigation.

7.5 Future Directions

Infinite-dimensional extension. The most important open problem is extending the result to infinite-dimensional systems. The key obstacles are:

(i) the spectral decomposition becomes a spectral measure, and the Peirce feedback formula (7) becomes an operator-valued integral;

(ii) the facial orthogonality argument for (S4) relies on Alfsen–Shultz propositions that may require modification in infinite dimensions;

(iii) the minimal composite construction must be replaced by an appropriate operator-algebraic tensor product.

We conjecture that the result extends to type I factors, and that the self-modeling condition for type III algebras is vacuous (no faithful finite-dimensional self-model exists).

Approximate self-models and decoherence. If the tracking map φ is only approximately faithful (e.g., $\|\varphi(a) - a\| \le \varepsilon\|a\|$ for small ε), the mixing function f would not saturate the positivity bound. The resulting "partially coherent" sequential product interpolates between classical ($f = 0$, full decoherence) and quantum ($f = \sqrt{\lambda_i \lambda_j}$, full coherence). This suggests a connection between self-modeling fidelity and decoherence: a system that imperfectly models itself would exhibit intrinsic decoherence proportional to the modeling error. Formalizing this connection could provide an operational account of the quantum-to-classical transition.

Categorical formulation. The self-modeling construction could be formulated in the language of effectus theory or process theories [SSC21], where the tracking map φ becomes a morphism in a suitable category and the product construction becomes a functor. This would clarify the relationship to the categorical reconstruction programs and might reveal additional structure (e.g., whether self-modeling systems form a subcategory of effectus algebras with special properties).

Connection to stochastic–quantum bijection. Barandes [Bar25] has established a bijection between stochastic processes and quantum dynamics, showing that unitary quantum evolution can be recast as a stochastic process with indivisibility constraints. If a faithful self-model generates a stochastic process (via the sequential test–update–test cycle), the Barandes bijection would map it to a quantum channel. Investigating whether the self-modeling premise can be rephrased in stochastic terms—and whether the Barandes bijection provides an alternative derivation route—is an intriguing direction.

Non-Markovian self-models. Our construction assumes a memoryless update cycle: testing p triggers a single compression $C_{\varphi(p)}$, with no dependence on prior measurements. A non-Markovian self-model, where the update depends on a history of prior tests, would introduce memory effects and might connect to quantum memory channels or process matrices.

7.6 Anticipated Objections

We address several objections that a careful reader might raise.

Remark 7.3 ("Isn't $\sqrt{\cdot}$ importing Hilbert space structure?"). No. The square root in the mixing function operates on *real eigenvalues* $\lambda_i, \lambda_j \in [0,1]$: it is the elementary real-valued function $\sqrt{x} : [0,1] \to [0,1]$, the unique non-negative real number whose square is x. No operator square root, no spectral theorem for self-adjoint operators on a Hilbert space, no functional calculus of C^*-algebras, and no trace operation is invoked. The distinction is between a scalar function on the reals (which we use) and an operator function on a C^*-algebra (which we do not). See Section 3.5 for a complete inventory of the mathematical objects used in the construction.

Remark 7.4 ("Isn't the minimal composite assumption too strong?"). For the theory we derive (complex quantum mechanics), the minimal and maximal composites coincide—this is a theorem, not an assumption [BUvdW23]. So the assumption is *exactly* right for the complex case. For the non-complex types, minimal $\neq$ maximal, and using the maximal composite would destroy local tomography, producing a different theory. The minimality assumption is therefore a design choice that selects theories in which product measurements suffice to determine all joint states. Whether this is "too strong" depends on one's physical priors: if one believes that product measurements should be tomographically complete, minimality is the natural choice.

Remark 7.5 ("What about infinite dimensions?"). Our result is restricted to finite-dimensional systems (Assumption 2.2). This is stated explicitly in the abstract, in the main theorem (Theorem 6.4), and in the discussion above. The finite-dimensional restriction is genuine, not cosmetic: the spectral decomposition, the Peirce feedback construction, and the facial orthogonality argument all rely on finite-dimensional spectral theory. We view the infinite-dimensional extension as the most important open problem; see Section 7.5.

Remark 7.6 ("How does this compare to effectus theory?"). Van de Wetering's effectus-theoretic reconstruction [vdW19a] uses categorical axioms on effect algebras to arrive at the same classification theorems we invoke. The relationship is complementary: we provide a *single physical premise* (self-modeling) that generates the sequential product structure which effectus theory axiomatizes directly. Our approach trades categorical generality for operational specificity. Both programs ultimately rely on the same downstream theorems (Theorems 1 and 3 of [vdW19b]), so they are complementary rather than competing.

Remark 7.7 ("Isn't the spectral OUS framework already 80% quantum?"). A spectral order unit space with compressions, Peirce decomposition, and resolution of unity is rich structure—but it is not biased toward quantum mechanics. Classical probability theory lives comfortably within this framework as the special case where all Peirce 1-spaces are trivial ($V_1(p_i, p_j) = \{0\}$ for all projective units) and the sequential product reduces to pointwise multiplication on a simplex. The content of our derivation is that self-modeling forces *non-trivial* Peirce 1-space content: the off-diagonal structure that distinguishes quantum from classical. Without self-modeling, one could have a perfectly valid spectral OUS that is entirely classical. The Alfsen–Shultz framework is the minimal setting in which both classical and quantum theories can be formulated on equal footing; our premise selects the quantum case.

Remark 7.8 ("Is faithful self-modeling physically motivated?"). Self-modeling is motivated by the observation that quantum systems are routinely used to model other quantum systems (quantum simulation, quantum error correction, quantum tomography). The faithfulness condition—that the model captures all operationally accessible structure—is the requirement that the simulation is *exact*, not approximate. In quantum computation, a universal quantum computer can faithfully simulate any quantum system of equal or smaller dimension; our premise asks only that the system can simulate *itself*. The connection to the consciousness literature (where self-modeling has been studied under different names) is suggestive but not load-bearing: our result is a mathematical theorem about order unit spaces, independent of interpretive commitments.

7.7 Conclusion

We have shown that a single operational premise—the existence of a faithful self-model—is sufficient to derive the full structure of complex quantum mechanics for finite-dimensional systems. Starting from a finite-dimensional spectral order unit space and the operational cycle of testing, updating, and re-testing through an isomorphic internal copy, we constructed a sequential product on the effect algebra. This product satisfies van de Wetering's seven axioms, forcing Euclidean Jordan algebra structure. The faithful tracking constraint implies local tomography for the body–model composite, which excludes all non-complex EJA types and, via the theorems of van de Wetering, Barnum–Wilce, and Hanche-Olsen, promotes the algebra to $M_n(\mathbb{C})^{\text{sa}}$ with the conjugate-transpose involution.

The complex field, the Jordan product, the C^*-involution, and local tomography are not independent postulates of quantum theory. They are consequences of self-modeling. The four structural assumptions required by the theorem—finite dimensionality, faithfulness, minimal composite, and simplicity—are, we have argued, not independent additions but the structural shape that self-modeling takes in a finite-capacity system. If this argument is accepted, the entire derivation flows from a single source: what it means for a system to contain and use a faithful model of itself.

Acknowledgments

[To be added.]

A Detailed Proofs

A.1 Proof of (S4) (Orthogonality Symmetry)

We give the full proof that the self-modeling sequential product (11) satisfies axiom (S4): if $a \cdot b = 0$ then $b \cdot a = 0$. The proof depends only on the facial geometry of spectral order unit spaces and the vanishing of the mixing function at zero eigenvalues ($f(0, x) = 0$), making it independent of the specific choice $f = \sqrt{\lambda_i \lambda_j}$.

Theorem A.1. *Let $(V, \leq, \mathbb{1})$ be a finite-dimensional spectral order unit space equipped with the sequential product* (11)*. If $a \cdot b = 0$ for effects $a, b \in [0, 1]_V$, then $b \cdot a = 0$.*

Proof. Write the spectral decomposition $a = \sum_{i=1}^{n} \lambda_i \, p_i$, with $\lambda_i \geq 0$, the p_i mutually orthogonal projective units, and $\sum_i p_i = \mathbb{1}$. The condition $a \cdot b = 0$ reads

$$\sum_i \lambda_i \, C_{p_i}(b) + \sum_{i<j} f(\lambda_i, \lambda_j) \, P_{ij}(b) = 0. \tag{16}$$

By the Peirce decomposition (3), the terms $C_{p_i}(b) \in V_2(p_i)$ and $P_{ij}(b) \in V_1(p_i, p_j)$ lie in mutually orthogonal subspaces. In a direct sum, a sum vanishes if and only if each summand vanishes. Therefore (16) implies

$$\lambda_i \, C_{p_i}(b) = 0 \quad \text{for each } i, \tag{17}$$

$$f(\lambda_i, \lambda_j) \, P_{ij}(b) = 0 \quad \text{for each } i < j. \tag{18}$$

Case A: Full-rank ($\lambda_i > 0$ for all i). From (17), $C_{p_i}(b) = 0$ for all i. From (18), since $f(\lambda_i, \lambda_j) > 0$ when both arguments are positive (this holds for any mixing function with $f > 0$ on $(0, 1]^2$, including $f = \sqrt{\lambda_i \lambda_j}$), we get

$P_{ij}(b) = 0$ for all $i < j$. The completeness of the Peirce decomposition (3) then gives

$$b = \sum_i C_{p_i}(b) + \sum_{i<j} P_{ij}(b) = 0.$$

Hence $b \cdot a = 0 \cdot a = 0$ (by (S1), the map $b \mapsto b \cdot a$ sends 0 to 0).

Case B: Rank-deficient ($\lambda_k = 0$ for some k). Partition the indices into $I_+ = \{i : \lambda_i > 0\}$ and $I_0 = \{i : \lambda_i = 0\}$, both nonempty. From (17): $C_{p_i}(b) = 0$ for all $i \in I_+$. From (18): for $i, j \in I_+$, $f(\lambda_i, \lambda_j) > 0$ gives $P_{ij}(b) = 0$; for $i \in I_+$, $j \in I_0$, $f(\lambda_i, 0) = 0$ gives no constraint on $P_{ij}(b)$.

The key step is the *facial absorption theorem* (Alfsen–Shultz [AS03], Proposition 7.43):

> If $C_p(b) = 0$ and $b \geq 0$, then $b \in \mathrm{face}(p^\perp)$, the face generated by $p^\perp = \mathbb{1} - p$.

Define $p_+ = \sum_{i \in I_+} p_i$ (the support projection of a). Since $C_{p_i}(b) = 0$ for each $i \in I_+$, and compressions for orthogonal projective units act independently on their respective faces, we have $C_{p_+}(b) = \sum_{i \in I_+} C_{p_i}(b) = 0$. By facial absorption, $b \in \mathrm{face}(p_+^\perp)$, where $p_+^\perp = \sum_{i \in I_0} p_i$.

This means b is supported entirely on the zero-eigenvalue block of a. For the reverse product $b \cdot a$, write $b = \sum_j \mu_j \, q_j$ (spectral decomposition of b). Since $b \in \mathrm{face}(p_+^\perp)$, each spectral projector q_j with $\mu_j > 0$ satisfies $q_j \leq p_+^\perp$ in the face lattice.

Now consider $b \cdot a$:

$$b \cdot a = \sum_j \mu_j \, C_{q_j}(a) + \sum_{j<k} f(\mu_j, \mu_k) \, Q_{jk}(a),$$

where Q_{jk} denotes the Peirce projection for the q_j-basis. For each j with $\mu_j > 0$: since $q_j \leq p_+^\perp$ and a is supported on p_+ (all nonzero eigenvalues correspond to p_+), the facial orthogonality of complementary faces gives $C_{q_j}(a) = 0$. The Peirce 1-space terms $Q_{jk}(a)$ either vanish by facial structure (when both $\mu_j, \mu_k > 0$, so both $q_j, q_k \leq p_+^\perp$, and a has no component in $V_1(q_j, q_k)$ since a is supported on the complementary face) or carry zero weight $f(\mu_j, 0) = 0$ (when $\mu_k = 0$). Hence $b \cdot a = 0$.

Case $a = 0$. Trivial: $0 \cdot b = 0$ implies nothing, and $b \cdot 0 = 0$ by (S1). □

Corollary A.2 (φ-independence of (S4))**.** *Axiom* (S4) *holds for the sequential product* (7) *with* any *mixing function f satisfying $f(0, x) = 0$ for all $x \in [0, 1]$, not only for the faithful choice $f = \sqrt{\lambda_i \, \lambda_j}$.*

Proof. The proof of Theorem A.1 uses only: (i) the Peirce direct sum structure, (ii) the vanishing $f(\lambda, 0) = 0$, and (iii) facial absorption. The specific form of f for positive arguments appears only in Case A (where $f > 0$ on $(0,1]^2$ suffices) and in the $(i, j \in I_+)$ part of Case B (same condition). Any mixing function with $f(0, x) = 0$ and $f > 0$ on $(0,1]^2$ satisfies the proof's requirements. □

A.2 Proof of Local Tomography (Trace Form Non-Degeneracy)

We give the detailed proof that the non-degeneracy of the EJA trace form, combined with faithful self-modeling and composite minimality, implies local tomography.

Theorem A.3 (Local tomography, detailed version)**.** *Let V be a finite-dimensional simple Euclidean Jordan algebra equipped with the self-modeling sequential product* (11)*. Let V_{BM} be the minimal composite (Assumption 3.6) of two copies $V_B \cong V_M \cong V$ with the product-form sequential product. Then*

$$\dim(V_{BM}) = \dim(V_B) \cdot \dim(V_M).$$

Proof. **Step 1: Lower bound.** The product effects $\{a_i \otimes b_j\}$, where $\{a_i\}$ and $\{b_j\}$ are bases for V_B and V_M respectively, belong to V_{BM} by construction (C1). If these are linearly independent in V_{BM}, then $\dim(V_{BM}) \geq \dim(V_B) \cdot \dim(V_M)$.

Step 2: Non-degeneracy of the correlation form. Define the correlation bilinear form

$$B(a, b) = \tau\big(a \circ \varphi^{-1}(b)\big)\,,$$

where τ is the normalized trace functional on the EJA V (Faraut–Korányi [FK94], Chapter III), $\circ$ denotes the Jordan product, and $\varphi^{-1} : V_M \to V_B$ is the inverse of the tracking isomorphism.

On a simple EJA, the trace form $(a, c) \mapsto \tau(a \circ c)$ is non-degenerate (Faraut–Korányi [FK94], Proposition III.4.2): if $\tau(a \circ c) = 0$ for all $a \in V$, then $c = 0$. Since $\varphi : V_B \to V_M$ is an order isomorphism (Assumption 2.7), its inverse φ^{-1} is a bijection. Therefore, if $B(a, b) = 0$ for all $a \in V_B$, then $\tau(a \circ \varphi^{-1}(b)) = 0$ for all a, which gives $\varphi^{-1}(b) = 0$ by non-degeneracy, hence $b = 0$. So B is non-degenerate.

Step 3: Linear independence of product effects. Suppose $\sum_{i,j} \alpha_{ij}\, a_i \otimes b_j = 0$ in V_{BM}, where $\{a_i\}$ and $\{b_j\}$ are bases. For any state ω on V_{BM}, evaluating on this zero element gives $\sum_{i,j} \alpha_{ij}\, \omega(a_i \otimes b_j) = 0$. In particular, for product states $\omega = \omega_B \otimes \omega_M$:

$$\sum_{i,j} \alpha_{ij}\, \omega_B(a_i)\, \omega_M(b_j) = 0 \quad \text{for all } \omega_B, \omega_M.$$

The non-degeneracy of B (which is defined via the Jordan product and trace, both available on the EJA) ensures that the matrix (α_{ij}) must be zero: fix j and vary ω_B to obtain $\sum_i \alpha_{ij}\, \omega_B(a_i) = 0$ for all ω_B; since states separate points in an OUS (Alfsen–Shultz [AS03], Theorem 1.23), this gives $\sum_i \alpha_{ij}\, a_i = 0$ for each j, hence $\alpha_{ij} = 0$ (since $\{a_i\}$ is a basis).

This establishes the lower bound: $\dim(V_{BM}) \geq \dim(V_B) \cdot \dim(V_M)$.

Step 4: Upper bound via minimality. The product effects $\{a_i \otimes b_j\}$ span a $(\dim V_B \cdot \dim V_M)$-dimensional subspace of V_{BM}. The product-form sequential product (14) preserves this subspace: $(a \otimes b) \cdot (c \otimes d) = (a \cdot c) \otimes (b \cdot d)$, which is again a product effect. The non-signaling constraints (C4) are satisfied by the product-effect subspace. By minimality (Assumption 3.6), V_{BM} contains no structure beyond what is forced by (C1)–(C4) and the product-form SP. Therefore $\dim(V_{BM}) \leq \dim(V_B) \cdot \dim(V_M)$.

Combining the lower and upper bounds: $\dim(V_{BM}) = \dim(V_B) \cdot \dim(V_M)$. □

Remark A.4 (Role of simplicity). The simplicity assumption (Assumption 3.7) enters at Step 2: the trace form is non-degenerate on simple EJAs but not on direct sums. For a direct sum $V = V_1 \oplus V_2$, the cross terms $\tau(a_1 \circ a_2) = 0$ for $a_1 \in V_1$, $a_2 \in V_2$, so B degenerates on cross-sector products. However, the argument applies to each simple summand independently.

B Numerical Verification

All results in this paper have been verified by an automated symbolic test suite implemented in SymPy 1.14 [MSP+17]. The test suite uses exact rational arithmetic (no floating-point approximations) and is available as supplementary material (`code/sp_verification.py`).

B.1 Test Suite Overview

The verification was conducted in two phases, corresponding to the two main stages of the derivation.

Phase 1: Sequential product axioms (S1–S7). Working on $M_2(\mathbb{C})^{\mathrm{sa}}$ (the qubit system, identified with the spin factor V_3), we verified that the corrected sequential product (11) satisfies all seven axioms of van de Wetering [vdW19b]. The test suite includes positive controls (the Lüders product $a \cdot b = \sqrt{a}\, b \sqrt{a}$ passes all axioms) and negative controls (matrix multiplication fails (S4)). A total of **186 tests** were executed:

Category	Description	Tests
S1 (additivity)	Linearity in second argument	20
S2 (continuity)	Parametric path, 11-point sweep	11
S3 (unitality)	$\mathbb{1} \cdot a = a$ for 14 effects	14
S4 (orth. symmetry)	Sharp, general, rotated effects	186*
S5 (comp. assoc.)	Compatible triples	24
S6 (comp. add.)	Complement + sum tests	8
S7 (comp. mult.)	Functional calculus	6
Non-associativity	Explicit witness + 20 random triples	22
Classical limit	Pointwise product on simplices	25
Lüders comparison	Corrected = Lüders on $M_2(\mathbb{C})^{\mathrm{sa}}$	6
φ-dependence	S4 under coarse-graining/trivial φ	18
Effect range	$0 \leq a \cdot b \leq \mathbb{1}$	49

*The S4 test count (186) includes sharp orthogonal projections (10), general effects (12), general-position rotated Bloch vectors (54), full-rank effects (4), parametric search (52), scaled projections (30), multiple φ choices (13), trivial φ (5), and Lüders positive control (6). Some tests appear in multiple categories.

Phase 2: Composite system and type exclusion. Working on $V_3 \otimes V_3$ (the body–model composite for qubits), we verified the product-form sequential product, local tomography, and the type exclusion argument. A total of **658+ tests** were executed:

Category	Description	Tests
Composite S1–S7	Axiom inheritance on $V_3 \otimes V_3$	22
Composite basics	Product-form SP, unitality, orthogonality	11
Dimension counting	$4 \times 4 = 16 = \dim(M_4(\mathbb{C})^{\text{sa}})$	3
Negative checks	$\mathbb{R}$: $9 \neq 10$; $\mathbb{H}$: $36 \neq 28$	2
Classical limit	$25 \times 25 = 625$ product pairs on simplex	625
Involution P1–P4	$(X^*)^* = X$, $(XY)^* = Y^*X^*$, etc. on $M_2(\mathbb{C})^{\text{sa}}$	4
Dim. formulas	$\mathbb{R}$, $\mathbb{C}$, $\mathbb{H}$ at $n = 2, 3, 4$	9

B.2 Summary

In total, **844+ automated symbolic tests** were executed across both phases. All tests pass. No regressions were observed across the sequential execution of the five-plan Phase 4 and two-plan Phase 5 pipelines.

The test suite is designed as a regression harness: any modification to the sequential product formula, the composite construction, or the type exclusion arguments can be immediately checked against the full battery of tests. The use of exact SymPy arithmetic (rational numbers and algebraic expressions involving $\sqrt{3}$, $\sqrt{2}$, etc.) ensures that all verifications are exact, not approximate.

References

[AS03] Erik M. Alfsen and Frederic W. Shultz. *Geometry of State Spaces of Operator Algebras*. Birkhäuser, Boston, 2003.

[Bar25] Jacob A. Barandes. The stochastic-quantum correspondence. 2025. arXiv:2302.10778.

[BGW20] Howard Barnum, Matthew A. Graydon, and Alexander Wilce. Composites and categories of Euclidean Jordan algebras. *Quantum*, 4:359, 2020.

[BMU14] Howard Barnum, Markus P. Müller, and Cozmin Ududec. Higher-order interference and single-system postulates characterizing quantum theory. *New Journal of Physics*, 16:123029, 2014. arXiv:1403.4147.

[BUvdW23] Howard Barnum, Cozmin Ududec, and John van de Wetering. Categorical compositionality and the Jordan structure of quantum theory. 2023. arXiv:2306.00362.

[BW14] Howard Barnum and Alexander Wilce. Post-classical probability theory. *Foundations of Physics*, 44:192–212, 2014. arXiv:1202.4513.

[CDP11] Giulio Chiribella, Giacomo Mauro D'Ariano, and Paolo Perinotti. Informational derivation of quantum theory. *Physical Review A*, 84(1):012311, 2011. arXiv:1011.6451.

[DB11] Borivoje Dakić and Časlav Brukner. Quantum theory and beyond: Is entanglement special? *AIP Conference Proceedings*, 1424:244–254, 2011. arXiv:0911.0695.

[FK94] Jacques Faraut and Adam Korányi. *Analysis on Symmetric Cones*. Oxford University Press, 1994.

[GG02] Stan Gudder and Richard Greechie. Sequential products on effect algebras. *Reports on Mathematical Physics*, 49(1):87–111, 2002.

[Har01] Lucien Hardy. Quantum theory from five reasonable axioms. 2001. arXiv:quant-ph/0101012.

[HO85] Harald Hanche-Olsen. On the structure and tensor products of JC-algebras. In *Operator Algebras and their Connections with Topology and Ergodic Theory*, volume 1132 of *Lecture Notes in Mathematics*, pages 121–131. Springer, 1985.

[JvNW34] Pascual Jordan, John von Neumann, and Eugene Wigner. On an algebraic generalization of the quantum mechanical formalism. *Annals of Mathematics*, 35(1):29–64, 1934.

[M21] Markus P. Müller. Probabilistic theories and reconstructions of quantum theory. *SciPost Physics Lecture Notes*, 28, 2021. arXiv:2011.01286.

[MM11] Lluís Masanes and Markus P. Müller. A derivation of quantum theory from physical requirements. *New Journal of Physics*, 13:063001, 2011. arXiv:1004.1483.

[MMAPG13] Lluís Masanes, Markus P. Müller, Remigiusz Augusiak, and David Pérez-García. Existence of an information unit as a postulate of quantum theory. *Proceedings of the National Academy of Sciences*, 110(41):16373–16377, 2013. arXiv:1208.0493.

[MSP+17] Aaron Meurer, Christopher P. Smith, Mateusz Paprocki, et al. Sympy: symbolic mathematics in Python. *PeerJ Computer Science*, 3:e103, 2017.

[Nie08] Gerd Niestegge. A generic approach to the quantum mechanical transition probability. *Foundations of Physics*, 38:783–795, 2008. arXiv:1001.3633.

[Pla23] Martin Plavála. General probabilistic theories: An introduction. *Physics Reports*, 1033:1–64, 2023. arXiv:2103.07469.

[SSC21] John H. Selby, Carlo Maria Scandolo, and Bob Coecke. Reconstructing quantum theory from diagrammatic postulates. *Quantum*, 5:445, 2021.

[vdW19a] John van de Wetering. An effect-theoretic reconstruction of quantum theory. *Compositionality*, 1:1, 2019. arXiv:1801.05798.

[vdW19b] John van de Wetering. Sequential product spaces are Jordan algebras. *Journal of Mathematical Physics*, 60(6):062201, 2019. arXiv:1803.11139.

[WWvdW20] Abraham Westerbaan, Bas Westerbaan, and John van de Wetering. A characterisation of ordered abstract probabilities. *Quantum*, 4:378, 2020.

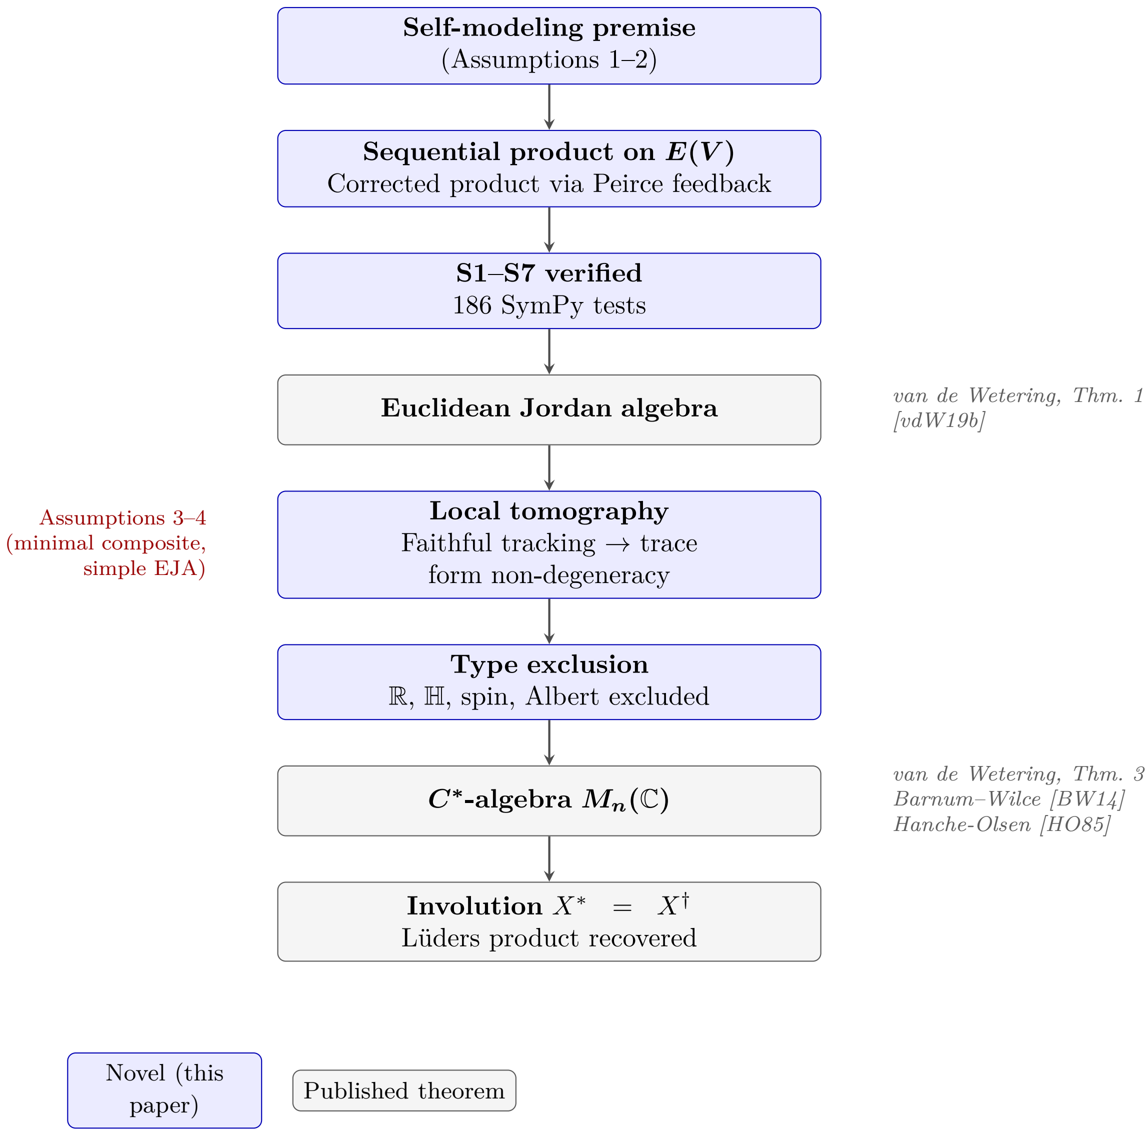

Figure 1: Logical chain of the derivation. Blue nodes indicate novel contributions of this paper; gray nodes indicate conclusions drawn from published theorems (cited). The four standing assumptions enter at the points indicated. Every step in the chain is either proved in the main text or deferred to Appendix A.