

Machine-Checked Finite Quantum Event Algebras: Born Weights, Lüders Conditioning, Partition Pinching, and the Tsirelson Bound in Lean 4

Bernhard Mueller*

Pragma Research Inc., Washington, United States

bernhard@floatingpragma.ai

July 15, 2026

Paper release: r1542 **Released:** July 15, 2026

Abstract

We present a machine-checked development of the elementary theory of finite-dimensional quantum event algebras: events as Hermitian idempotents in the algebra of $n \times n$ complex matrices, states as positive-semidefinite matrices of unit trace, Born weights as the trace pairing, Lüders conditioning with its repeatability and fixed-point structure, the trace-preserving pinching onto the block-diagonal commutant of a projective partition of unity, the positive expectation functional of a state, and the Tsirelson bound for the CHSH combination. The development comprises 64 lemmas and 8 definitions in Lean 4 over the Mathlib library, with no proof placeholders. An axiom audit embedded in the source and re-executed on every build records the axiom dependencies of each result; no theorem depends on axioms outside Lean and Mathlib’s standard classical base [`propext`, `Classical.choice`, `Quot.sound`], and the core CHSH ring identity is choice-free. Every lemma is tagged as *algebra-only* (provable from the $\ast$ -algebra axioms alone) or *trace-dependent* (passing through the trace pairing), so the boundary between algebraic and probabilistic content is machine-visible. The Tsirelson bound is proved in operator-norm form in an abstract unital C^* -setting that assumes neither completeness nor an order structure, then instantiated to finite matrix algebras; it complements the order-form inequality already present in Mathlib and reuses Mathlib’s CHSH hypothesis interface for interoperability. We describe the formalization module by module, document the Mathlib gaps and workarounds encountered, and situate the development relative to the existing Lean and Isabelle/HOL formalizations of projective measurements and the CHSH inequality.

Keywords: quantum probability, event algebras, Born rule, Lüders conditioning, conditional expectation, Tsirelson bound, formalized mathematics, Lean 4, Mathlib

Mathematics Subject Classification (2020): Primary 81P15, 68V20; Secondary 81P40, 46L53, 81P10.

*Corresponding author: bernhard@floatingpragma.ai

1 Introduction

Finite-dimensional event algebras support a complete and elementary account of quantum probability. Over the matrix algebra $M_n(\mathbb{C})$, events are orthogonal projections, states are density matrices, the Born weight of an event under a state is a trace pairing, conditioning is the Lüders update [8, 21], classicality is captured by commutative subalgebras and coarse-graining onto a measurement context by a trace-preserving pinching, and the characteristic nonclassical phenomenon, the violation of the CHSH inequality up to Tsirelson’s bound [5, 9, 10], is already visible for 4×4 matrices. Every object in this list is finite-dimensional linear algebra; no unbounded operators, no measure theory, and no functional-analytic subtleties intervene. This elementary core is also where much of the current foundational discussion lives: the status and derivation of the Born rule [1, 26], the relation between quantum phenomena and Kolmogorov probability [25], and the derivability of the von Neumann–Lüders measurement postulate [18] remain active topics, including in this journal.

Published treatments of this core, from von Neumann [28] onward, leave the bookkeeping informal. Which statements require positivity of the state, which require only hermiticity, and which are identities of the trace calculus valid for arbitrary matrices? What exactly happens when one conditions on an event of weight zero? Is orthogonality of projections a symmetric hypothesis or a symmetric conclusion? Where, precisely, does the trace enter, and where does pure $*$ -algebra suffice? In informal mathematics these questions are resolved silently, case by case, by the reader. A proof assistant forces every such choice into the open and then certifies the result.

To the author’s knowledge (search current as of July 2026), this is the first Lean 4 development to package this specific combination: finite projective-event closure lemmas, a Lüders fixed-point / one-step-retraction characterization, the partition-pinching expectation with its Hilbert–Schmidt characterization, and a norm-form CHSH/Tsirelson bound, under a per-declaration axiom audit. Concretely, the development delivers:

- (a) 64 lemmas and 8 definitions, in Lean 4 [11] over Mathlib [27], with no `sorry` placeholders and no axioms beyond Lean’s standard classical base;
- (b) an axiom audit embedded in the source: every lemma is followed by a `#print axioms` command, so a complete per-declaration audit re-executes on every build; no theorem depends on axioms outside Lean and Mathlib’s standard classical base [`propext`, `Classical.choice`, `Quot.sound`], and individual dependency sets are subsets of this base (for example, the CHSH ring identity `chsh_mul_self` is choice-free, reporting [`propext`, `Quot.sound`]);
- (c) a tagging discipline under which every lemma is classified in its documentation string as *algebra-only* (content provable in the $*$ -algebra, with norm structure where relevant) or *trace-dependent* (content passing through the trace pairing $(\rho, P) \mapsto \text{Tr}(\rho P)$); trace-dependent statements divide further according to whether they additionally require positivity or state-normalization, which is visible in each statement’s hypotheses;
- (d) the Lüders update packaged with repeatability, idempotence, composition on commuting events, the classical restriction, and a fixed-point characterization identifying the fixed states of conditioning with the states certain of the event;
- (e) the trace-preserving pinching onto the block-diagonal commutant N of a projective partition of unity, with a Hilbert–Schmidt characterization: projector, state-preserving, trace-

selfadjoint, Pythagorean, contractive, trace-compatible, unique among commutant-valued trace-compatible maps, and compatible with Lüders conditioning on commutant events;

- (f) the Tsirelson bound in operator-norm form, proved in an abstract unital C^* -setting that assumes neither completeness nor any order structure, with the exact ring identity underlying it isolated as a standalone lemma, an instantiation to finite matrix algebras, and an interface lemma reusing Mathlib’s existing CHSH hypothesis bundle.

The development builds with the toolchain `leanprover/lean4:v4.29.1` and the Mathlib version pinned to match; it imports Mathlib and nothing else.

Three scope limitations should be stated at the outset. First, the novelty claim above is deliberately narrow and organizational. Formalizations of projective measurement, the Born rule for projective effects, and the CHSH/Tsirelson inequality already exist, in the Isabelle/HOL ecosystem (in particular the work of Echenim and Mhalla [13, 15, 16]) and in the Lean ecosystem (Mathlib’s order-form CHSH bound and the Physlib/QuantumInfo project [24]); we make no claim of first formalization for any of those individual results. The contribution here is the organization of the finite core around a projective partition with its trace-preserving pinching onto the block-diagonal commutant, the lemma-by-lemma separation of algebra-only from trace-dependent content, the abstract C^* -setting of the norm-form bound, and the prover ecosystem. Second, no attainment or tightness statement for the Tsirelson bound is formalized: the development ships the inequality only. Third, everything is finite-dimensional; extensions toward infinite dimensions are discussed as outlook in Section 11.

Finite event algebras arise, for example, as the measurement cores of finite models of quantum theory. Beyond that single remark, this paper is self-contained: Section 2 fixes the mathematical setting, Section 3 describes the architecture and conventions of the formalization, Sections 4 through 8 walk through the five modules with the main statements as numbered propositions carrying their Lean names, Section 9 documents Mathlib gaps and proof engineering, Section 10 discusses related work, and Section 11 concludes.

2 Mathematical setting

Throughout, n and k denote natural numbers and $M_n(\mathbb{C})$ denotes the $*$ -algebra of $n \times n$ complex matrices, with conjugate transpose $X \mapsto X^*$, trace Tr , identity $\mathbf{1}$, and zero 0 . All statements below are formalized for every n , including the degenerate case $n = 0$; only the matrix instantiation of the Tsirelson bound (Proposition 20) assumes $n \neq 0$.

The partial order on $\mathbb{C}$ is the one in which $z \leq w$ exactly when $\text{Re } z \leq \text{Re } w$ and $\text{Im } z = \text{Im } w$; in particular $0 \leq z$ asserts that z is real and nonnegative. Nonnegativity and boundedness of Born weights are proved below in this partial order, which is strictly stronger than the corresponding assertions about real parts; real-part corollaries are provided alongside each such statement.

Definition 1 (Event). *An event of $M_n(\mathbb{C})$ is a Hermitian idempotent: a matrix $P \in M_n(\mathbb{C})$ with $P^* = P$ and $P^2 = P$. Lean name: `IsEvent`.*

Events are exactly the orthogonal projections, the sharp yes/no tests of the finite quantum formalism [28]. The *subevent* relation is witnessed algebraically: P is a subevent of Q when $PQ = P$.

Definition 2 (State). *A state on $M_n(\mathbb{C})$ is a density matrix: a positive-semidefinite $\rho \in M_n(\mathbb{C})$ with $\text{Tr } \rho = 1$. Lean name: `IsState`.*

Definition 3 (Born weight). *The Born weight of P under ρ is the trace pairing*

$$\mu_\rho(P) := \text{Tr}(\rho P) \in \mathbb{C}.$$

Lean name: `bornWeight`.

The weight is defined as a complex number for arbitrary matrix arguments; reality under hermiticity is a theorem (Proposition 3), and the probabilistic bounds hold under the appropriate positivity hypotheses (Proposition 4).

Definition 4 (Lüders update). *The Lüders update of ρ by P is the compressed and renormalized matrix*

$$\mathcal{L}_P(\rho) := (\mu_\rho(P))^{-1} P \rho P.$$

Lean name: `luedersUpdate`.

The scalar inverse is taken in $\mathbb{C}$ with Lean’s convention $0^{-1} = 0$, so conditioning on an event of vanishing weight degenerates to the zero matrix; every substantive lemma about the update carries the explicit guard $\mu_\rho(P) \neq 0$. The rule goes back to Lüders [21]; see Busch and Lahti [8] for its place in the measurement literature.

Definition 5 (Projective partition). *A projective partition of $M_n(\mathbb{C})$ is a finite family $\pi = (P_i)_{i < k}$ of events that are pairwise orthogonal, $P_i P_j = 0$ for $i \neq j$, and complete, $\sum_{i < k} P_i = \mathbf{1}$. Lean name: `ProjectivePartition`.*

The members of a projective partition commute pairwise (Proposition 11), so they generate a commutative unital $*$ -subalgebra

$$D := \text{span}\{P_i : i < k\} \subseteq M_n(\mathbb{C}).$$

The associated *commutant* is the block-diagonal algebra

$$N := \pi' := \{X \in M_n(\mathbb{C}) : X P_i = P_i X \text{ for all } i < k\} \cong \bigoplus_{i < k} M_{\text{rank } P_i}(\mathbb{C}),$$

which is in general noncommutative. The commutative algebra D is exactly the center of the commutant, $D = Z(N)$: the block-scalar matrices. This is the one precise sense in which a “center” figures here. As the next definition records, the pinching lands in N , not in D .

Definition 6 (Partition pinching). *The partition pinching attached to a projective partition π is*

$$\mathcal{E}_\pi(X) := \sum_{i < k} P_i X P_i.$$

Lean name: `pinchingExpectation`.

Section 6 verifies that $\mathcal{E}_\pi$ is the trace-preserving conditional expectation onto the commutant π' , with the Hilbert–Schmidt characterization: it is a projector onto π' , maps states to states, is selfadjoint for the trace inner product $\langle X, Y \rangle = \text{Tr}(X^* Y)$, satisfies the Pythagorean identity, contracts

the squared trace norm, is uniquely determined by trace compatibility, and commutes with Lüders conditioning on commutant events. When the partition is rank one, so that every P_i has rank one, the commutant π' coincides with D and $\mathcal{E}_\pi$ is the classical conditional expectation onto D .

Finally, the CHSH combination. In a unital C^* -algebra A , take selfadjoint involutions a_0, a_1, b_0, b_1 (each equal to its own star, each squaring to $\mathbf{1}$) such that every a_i commutes with every b_j , and form

$$S := a_0b_0 + a_0b_1 + a_1b_0 - a_1b_1.$$

Classically, dichotomic ± 1 observables obey $|\langle S \rangle| \leq 2$ [5, 10]; quantum mechanically the sharp bound is Tsirelson's $\|S\| \leq 2\sqrt{2}$ [9]. The bridge from the event layer to these hypotheses is the dichotomic observable $\mathbf{1} - 2P$ of an event P , which is a selfadjoint involution (Proposition 1(ix)).

3 The formalization: architecture and conventions

3.1 Modules

The development consists of five modules under a single namespace, importing Mathlib and nothing else:

Module	Content	Definitions	Lemmas
Basic	events, states, Born weights	IsEvent , IsState , bornWeight	25
Lueders	conditioning, certainty sets, fixed points	luedersUpdate , certainStates	12
CenterExpectation	partitions, pinching onto the commutant	ProjectivePartition , pinchingExpectation	16
StateFromTrace	the expectation functional of a state	expectation	5
Tsirelson	CHSH ring identity, norm-form Tsirelson bound	(none)	6
Total		8	64

Table 1: Module layout of the development.

The event and state layers are formalized over the concrete carrier **Matrix (Fin n) (Fin n) ℂ**. This choice trades generality for the density of Mathlib's finite matrix API (positive semidefiniteness, trace calculus, spectral theorem), which carries most of the proofs. The Tsirelson module, whose content is norm-theoretic rather than matrix-theoretic, is instead stated over an abstract typeclass context and instantiated to matrices only at the end (Section 8).

3.2 Scoped instances

Three pieces of structure that the development relies on are *scoped* in Mathlib rather than globally active, and must be opened explicitly:

- **ComplexOrder**: the partial order on $\mathbb{C}$ of Section 2, together with the ordered-ring structure that makes products of nonnegative complex numbers nonnegative;

- **MatrixOrder**: the Loewner order on matrices;
- **Matrix.Norms.L2Operator**: the L^2 operator norm on $M_n(\mathbb{C})$ together with its normed-ring and C^* -identity instances. There is no globally installed norm on the matrix type, so the matrix instantiation of the Tsirelson bound lives behind this scope.

Section 9 discusses the practical consequences of this scoping.

3.3 The algebra-only/trace-dependent split

Every lemma in the development carries one of two tags in its documentation string:

- **algebra-only**: the content is pure $*$ -algebra (ring and star structure; in the Tsirelson module, also norm structure), consuming no state and no trace functional;
- **trace-dependent**: the content passes through the trace pairing $(\rho, P) \mapsto \text{Tr}(\rho P)$; trace-dependent statements divide further by whether they also require positivity or state-normalization, which is visible in each statement's hypotheses.

The split is a statement-classification discipline rather than an import boundary: it makes visible, lemma by lemma, exactly where probabilistic content enters the theory. All event-closure lemmas, the entire partition algebra, the projector laws of the pinching map, and the entire Tsirelson module are algebra-only; every statement about Born weights, updates of states, and trace inner products consumes the trace. In the propositions below, this classification is reproduced: statements in Sections 4 through 7 are tagged individually, and everything in Section 8 is algebra-only.

3.4 Axiom audit and build discipline

Each of the 64 lemmas is followed in the source by a `#print axioms` command, so the axiom audit is part of the source text and re-executes on every build. No result depends on any axiom outside Lean's standard classical base

`[propext, Classical.choice, Quot.sound]`

(propositional extensionality, choice, and quotient soundness); individual dependency sets are subsets of this base. The ring identity `chsh_mul_self` is choice-free, reporting only `[propext, Quot.sound]`. The complete declaration-by-declaration audit accompanies the artifact as `AXIOM_AUDIT.txt`. The development contains no `sorry` placeholders, no custom axioms, and no unfinished obligations; the toolchain is `leanprover/lean4:v4.29.1` with Mathlib pinned at the matching version, and the project builds cleanly under `lake` (8254 jobs).

Proofs in this paper are given as sketches only; the complete formal proofs constitute the accompanying development, and each proposition cites the corresponding Lean names in `typewriter` font.

4 Events and Born weights

The first module establishes closure of the event set under the natural logical operations and the probabilistic behavior of the Born weight.

Proposition 1 (Event closure; algebra-only). *Let $P, Q \in M_n(\mathbb{C})$ be events. Then:*

- (i) 0 and 1 are events (*isEvent_zero, isEvent_one*);
- (ii) the complement $1 - P$ is an event (*IsEvent.compl*);
- (iii) orthogonality is symmetric: if $PQ = 0$ then $QP = 0$ (*IsEvent.orthogonal_symm*);
- (iv) if $PQ = 0$ then $P + Q$ is an event (*IsEvent.add*);
- (v) if $PQ = QP$ then PQ is an event (*IsEvent.mul_of_commute*);
- (vi) subevent absorption: if $PQ = P$ then $QP = P$ (*IsEvent.absorb_of_le*);
- (vii) if $PQ = P$ then $Q - P$ is an event (*IsEvent.sub_of_le*);
- (viii) every event is positive semidefinite (*IsEvent.posSemidef*);
- (ix) $1 - 2P$ is a selfadjoint involution: $(1 - 2P)^* = 1 - 2P$ and $(1 - 2P)^2 = 1$ (*IsEvent.one_sub_two_smul_involution*).

Proof sketch. All items are short computations in the $*$ -algebra. Symmetry of orthogonality (iii) and absorption (vi) take the conjugate transpose of the hypothesis and use hermiticity of both matrices; positive semidefiniteness (viii) exhibits $P = P^*P$ as a Gram matrix. $\square$

Item (ix) is the bridge from events to the CHSH hypotheses of Section 8: each event answers a yes/no question, and $1 - 2P$ is the associated ± 1 -valued observable.

Proposition 2 (Trace-pairing calculus; trace-dependent). *For all $\rho, P, Q \in M_n(\mathbb{C})$, all finite families $(P_i)_{i \in s}$, and all $c \in \mathbb{C}$:*

- (i) $\mu_\rho(P + Q) = \mu_\rho(P) + \mu_\rho(Q)$ and $\mu_\rho(P - Q) = \mu_\rho(P) - \mu_\rho(Q)$ (*bornWeight_add, bornWeight_sub*);
- (ii) $\mu_\rho(\sum_{i \in s} P_i) = \sum_{i \in s} \mu_\rho(P_i)$ (*bornWeight_sum*);
- (iii) $\mu_{c\rho}(P) = c \mu_\rho(P)$ (*bornWeight_smul*);
- (iv) if $P^2 = P$ then $\text{Tr}(P\rho P) = \mu_\rho(P)$ (*trace_sandwich*).

Item (iv), the sandwich identity, requires only idempotence of P ; hermiticity is never used. It is the workhorse that lets positivity arguments compress the state by the event and is applied throughout the conditioning module.

Proposition 3 (Reality; trace-dependent). *If ρ and P are Hermitian (neither positivity nor normalization is assumed), then $\overline{\mu_\rho(P)} = \mu_\rho(P)$; equivalently, the weight equals its own real part (*star_bornWeight, bornWeight_eq_re*).*

Proposition 4 (Probabilistic bounds; trace-dependent). *Let P, Q be events. Then, in the partial order of $\mathbb{C}$:*

- (i) if ρ is positive semidefinite, then $0 \leq \mu_\rho(P)$ (*bornWeight_nonneg*), with real-part form *bornWeight_re_nonneg*;

- (ii) if ρ is a state, then $\mu_\rho(\mathbf{1}) = 1$ (*bornWeight_one*);
- (iii) if $PQ = 0$, then $\mu_\rho(P + Q) = \mu_\rho(P) + \mu_\rho(Q)$ (*bornWeight_add_of_orthogonal*);
- (iv) if ρ is a state, then $\mu_\rho(P) \leq 1$ (*bornWeight_le_one*), with real-part form *bornWeight_re_le_one*;
- (v) if ρ is positive semidefinite and $PQ = P$, then $\mu_\rho(P) \leq \mu_\rho(Q)$ (*bornWeight_mono*);
- (vi) if ρ is positive semidefinite, then $\mu_\rho(P) \neq 0$ exactly when $0 < \text{Re } \mu_\rho(P)$ (*bornWeight_ne_zero_iff_re_pos*).

Proof sketch. For (i), compress: $\mu_\rho(P) = \text{Tr}(P\rho P)$ by Proposition 2(iv), and $P\rho P = P\rho P^*$ is a positive-semidefinite sandwich, whose trace is nonnegative. The upper bound (iv) applies (i) to the complement $\mathbf{1} - P$; monotonicity (v) applies (i) to the difference event $Q - P$ of Proposition 1(vii). $\square$

Together, Propositions 1 and 4 verify that $P \mapsto \mu_\rho(P)$ is a finitely additive, normalized, monotone probability assignment on events, with reality and the bounds stated in the strong ($\mathbb{C}$ -order) form. Item (vi) is the exact interface later needed by the conditioning guards.

5 Lüders conditioning

The second module packages the Lüders update as a state transformer and proves its characteristic laws. A remark on hypothesis minimality is in order: because the formal statements register exactly what each proof uses, several conditioning identities below hold for an *arbitrary* matrix ρ , with state hypotheses appearing only where positivity or normalization genuinely enters.

Proposition 5 (State property; trace-dependent). *Let P be an event.*

- (i) If ρ is positive semidefinite, then $\mathcal{L}_P(\rho)$ is positive semidefinite, with no guard: in the degenerate case the update is the zero matrix (*luedersUpdate_posSemidef*).
- (ii) If $\mu_\rho(P) \neq 0$, then $\text{Tr } \mathcal{L}_P(\rho) = 1$ for arbitrary ρ (*trace_luedersUpdate*).
- (iii) If ρ is a state and $\mu_\rho(P) \neq 0$, then $\mathcal{L}_P(\rho)$ is a state (*luedersUpdate_isState*).

By Proposition 4(vi), the guard $\mu_\rho(P) \neq 0$ is equivalent, for positive-semidefinite ρ , to the weight having strictly positive real part.

Proposition 6 (Repeatability; trace-dependent). *Let P be an event and $\mu_\rho(P) \neq 0$. Then*

$$\mu_{\mathcal{L}_P(\rho)}(P) = 1 \quad (\text{bornWeight_luedersUpdate_self}).$$

Neither positivity nor normalization of ρ is used: repeatability is an identity of the trace calculus, valid for every matrix ρ subject to the guard.

Proposition 7 (Idempotence; trace-dependent). *Let P be an event and $\mu_\rho(P) \neq 0$. Then*

$$\mathcal{L}_P(\mathcal{L}_P(\rho)) = \mathcal{L}_P(\rho) \quad (\text{luedersUpdate_idem}).$$

Proposition 8 (Commuting composition and order exchange; trace-dependent). *Let P, Q be events with $PQ = QP$.*

- (i) *If $\mu_\rho(P) \neq 0$, then $\mathcal{L}_Q(\mathcal{L}_P(\rho)) = \mathcal{L}_{PQ}(\rho)$ (*luedersUpdate_luedersUpdate_of_commute*); the product PQ is an event by Proposition 1(v). Only the first guard is needed: if the joint weight vanishes, both sides degenerate to zero together.*
- (ii) *If moreover $\mu_\rho(Q) \neq 0$, then conditioning is order-exchangeable: $\mathcal{L}_Q(\mathcal{L}_P(\rho)) = \mathcal{L}_P(\mathcal{L}_Q(\rho))$ (*luedersUpdate_comm*).*

Proposition 9 (Classical restriction; trace-dependent). *Let P be an event and suppose the state matrix commutes with it, $\rho P = P\rho$. Then*

$$\mathcal{L}_P(\rho) = (\mu_\rho(P))^{-1} \rho P \quad (\textit{luedersUpdate_of_commute}),$$

with no nonvanishing guard: the identity holds for arbitrary ρ under the stated commutation.

In the commuting case the Lüders update is the normalized restriction, which is the formula of classical conditional probability; quantum conditioning strictly extends the classical rule.

5.1 Conditioning as a one-step retraction onto its fixed points

Repeatability and idempotence assemble into a fixed-point description of measurement update.

Definition 7 (Certainty set). *The certainty set of an event P is*

$$\mathcal{C}_P := \{ \sigma \in M_n(\mathbb{C}) : \sigma \text{ is a state and } \mu_\sigma(P) = 1 \}.$$

Lean name: `certainStates`.

Proposition 10 (Fixed-point characterization; trace-dependent). *Let P be an event.*

- (i) *If ρ is a state with $\mu_\rho(P) \neq 0$, then $\mathcal{L}_P(\rho) \in \mathcal{C}_P$: conditioning lands in the certainty set in a single step (*luedersUpdate_mem_certainStates*).*
- (ii) *Support: if $\sigma \in \mathcal{C}_P$, then $\sigma P = \sigma = P\sigma$ (*mul_eq_self_of_bornWeight_one*).*
- (iii) *Every $\sigma \in \mathcal{C}_P$ is a fixed point: $\mathcal{L}_P(\sigma) = \sigma$ (*luedersUpdate_eq_self_of_mem_certainStates*).*
- (iv) *Conversely, for a state σ with $\mu_\sigma(P) \neq 0$,*

$$\mathcal{L}_P(\sigma) = \sigma \iff \sigma \in \mathcal{C}_P \quad (\textit{luedersUpdate_eq_self_iff}).$$

Proof sketch. The substantive item is (ii). The complement $Q = \mathbf{1} - P$ has weight zero under σ , so the compression $Q\sigma Q$ is positive semidefinite with zero trace and hence vanishes. The Cauchy–Schwarz inequality for the quadratic form of σ (in the form: if $\langle x, \sigma x \rangle = 0$ for a positive-semidefinite σ then $\sigma x = 0$) then kills σQ itself, column by column. Items (iii) and (iv) unfold the update using (ii). $\square$

Remark 1. Among the states assigning P nonzero weight, conditioning on P is a retraction onto its own fixed-point set $\mathcal{C}_P$: measurement update reaches a fixed point in one step. The formalization thus supports the reading of the Lüders rule as a projection onto certainty rather than a dynamical process with transient behavior. For a state σ the guard in the equivalence is automatic in one direction, since a zero-weight update is the zero matrix, which is never a state.

6 Trace-preserving pinching onto the block-diagonal commutant

The third module is the organizational core of the development. A projective partition $\pi = (P_i)_{i < k}$ (Definition 5) plays the role of a measurement context; the pinching $\mathcal{E}_\pi$ (Definition 6) is verified to be *the* trace-preserving conditional expectation onto the block-diagonal commutant π' , in the precise sense of the uniqueness statement below. The range π' is in general noncommutative, so $\mathcal{E}_\pi$ preserves quantum coherence within a degenerate block; it reduces to the classical conditional expectation onto $D = \text{span}\{P_i\}$ exactly when the partition is rank one.

Proposition 11 (Partition algebra; algebra-only). *For a center partition π and $i, j < k$:*

- (i) $P_i P_j = \delta_{ij} P_i$ (*ProjectivePartition.proj_mul_proj*);
- (ii) $P_i P_j = P_j P_i$: *the partition generates a commutative subalgebra* (*ProjectivePartition.proj_commute*).

Proposition 12 (Projector laws; algebra-only). *For every $X \in M_n(\mathbb{C})$ and $i < k$:*

- (i) *absorption: $\mathcal{E}_\pi(X) P_i = P_i X P_i = P_i \mathcal{E}_\pi(X)$* (*pinchingExpectation_mul_proj, proj_mul_pinchingExpectation*);
- (ii) $\mathcal{E}_\pi(X) \in \pi'$: *the pinching is commutant-valued* (*proj_commute_pinchingExpectation*);
- (iii) $\mathcal{E}_\pi$ *fixes π' pointwise: if $X P_i = P_i X$ for all i , then $\mathcal{E}_\pi(X) = X$* (*pinchingExpectation_fixes*);
- (iv) $\mathcal{E}_\pi \circ \mathcal{E}_\pi = \mathcal{E}_\pi$: *the pinching is a projector* (*pinchingExpectation_idem*);
- (v) *star equivariance: $\mathcal{E}_\pi(X)^* = \mathcal{E}_\pi(X^*)$* (*conjTranspose_pinchingExpectation*).

Proposition 13 (State preservation; trace-dependent). *If ρ is a state, then $\mathcal{E}_\pi(\rho)$ is a state* (*pinchingExpectation_isState*).

Proposition 14 (Trace-inner-product geometry; trace-dependent). *Write $\langle X, Y \rangle = \text{Tr}(X^* Y)$. For all $X, Y \in M_n(\mathbb{C})$:*

- (i) *selfadjointness: $\langle \mathcal{E}_\pi(X), Y \rangle = \langle X, \mathcal{E}_\pi(Y) \rangle$* (*trace_conjTranspose_pinchingExpectation_mul*);
- (ii) *Pythagoras:*

$$\langle X, X \rangle = \langle \mathcal{E}_\pi(X), \mathcal{E}_\pi(X) \rangle + \langle X - \mathcal{E}_\pi(X), X - \mathcal{E}_\pi(X) \rangle$$
(trace_pinchingExpectation_pythagoras);
- (iii) *squared- L^2 contraction, in the partial order of $\mathbb{C}$: $\langle \mathcal{E}_\pi(X), \mathcal{E}_\pi(X) \rangle \leq \langle X, X \rangle$* (*trace_pinchingExpectation_contraction*).

This is the quantum counterpart of the classical L^2 energy identity for conditional expectations: the pinching is the orthogonal projection, for the trace inner product, onto the commutant of the context.

Proposition 15 (Trace compatibility and uniqueness; trace-dependent). *(i) For every $X \in M_n(\mathbb{C})$ and every $C \in \pi'$,*

$$\text{Tr}(\mathcal{E}_\pi(X) C) = \text{Tr}(XC) \quad (\text{trace_pinchingExpectation_mul_central}).$$

(ii) Conversely, if $Y \in \pi'$ satisfies $\text{Tr}(YC) = \text{Tr}(XC)$ for every $C \in \pi'$, then $Y = \mathcal{E}_\pi(X)$ (`pinchingExpectation_unique`).

Proof sketch. Item (i) is a cyclic-trace computation using absorption. For (ii), set $D = Y - \mathcal{E}_\pi(X)$, observe $D \in \pi'$, test the hypothesis and item (i) against $C = D^*$, and conclude $\text{Tr}(DD^*) = 0$; faithfulness of the trace forces $D = 0$. $\square$

Item (i) is the defining property of a conditional expectation, and (ii) says the pinching is the only commutant-valued map with that property: existence and uniqueness are both machine-checked.

Proposition 16 (Compatibility with conditioning; trace-dependent). *Let $P \in \pi'$ (an event commuting with every P_i). For every $\rho \in M_n(\mathbb{C})$:*

(i) Born-weight invariance: $\mu_{\mathcal{E}_\pi(\rho)}(P) = \mu_\rho(P)$ (`bornWeight_pinchingExpectation`);

(ii) conditioning commutes with the center expectation:

$$\mathcal{E}_\pi(\mathcal{L}_P(\rho)) = \mathcal{L}_P(\mathcal{E}_\pi(\rho)) \quad (\text{pinchingExpectation_luedersUpdate}).$$

Item (i) guarantees that the normalizing scalars on both sides of (ii) agree, so the exchange law holds at the level of matrices, degenerate cases included. Conditioning on commutant events therefore commutes with the pinching onto π' . In the rank-one case, where $\pi' = D$ is commutative, this is exactly the classical statement that conditioning interacts with a sub- σ -algebra through the conditional expectation; in general it is the block-diagonal refinement of that statement.

7 The expectation functional

The fourth module attaches to each state its expectation functional on observables and proves that it is a positive normalized linear functional, strengthening Born-weight positivity from events to arbitrary positive-semidefinite observables.

Definition 8 (Expectation functional). *For $\rho, M \in M_n(\mathbb{C})$, write $\langle M \rangle_\rho := \text{Tr}(\rho M)$. Lean name: `expectation`.*

Proposition 17 (Positive normalized functional; trace-dependent). *For all $\rho, M, N \in M_n(\mathbb{C})$ and $c \in \mathbb{C}$:*

(i) the Born weight is the restriction of the functional to events: $\mu_\rho(P) = \langle P \rangle_\rho$ (`bornWeight_eq_expectation`);

(ii) additivity and homogeneity: $\langle M + N \rangle_\rho = \langle M \rangle_\rho + \langle N \rangle_\rho$ and $\langle cM \rangle_\rho = c\langle M \rangle_\rho$ (`expectation_add`, `expectation_smul`);

(iii) normalization: $\langle 1 \rangle_\rho = 1$ for a state ρ (`expectation_one`);

(iv) *positivity*: if ρ and M are both positive semidefinite, then $0 \leq \langle M \rangle_\rho$ in the partial order of $\mathbb{C}$ (*expectation_nonneg*).

Proof sketch. Positivity no longer has an idempotent to compress with, so the proof decomposes M spectrally: $M = VDV^*$ with V unitary and D a nonnegative diagonal, cycles the trace to $\text{Tr}((V^*\rho V)D)$, and sums termwise products of nonnegative diagonal entries of the sandwiched state against nonnegative eigenvalues. Section 9 explains why this elementary route was chosen over the abstract Gram-factorization route. $\square$

8 The Tsirelson bound

The fifth module is entirely algebra-only: no state and no trace appear anywhere in it. Its content splits into an exact ring identity, two C*-norm lemmas, the abstract bound, and interoperability statements.

8.1 The ring identity

Proposition 18 (CHSH square identity; algebra-only). *Let R be any ring and $a_0, a_1, b_0, b_1 \in R$ with $a_i^2 = b_j^2 = 1$ and $a_i b_j = b_j a_i$ for all $i, j \in \{0, 1\}$. With $S = a_0 b_0 + a_0 b_1 + a_1 b_0 - a_1 b_1$,*

$$S^2 = 4 \cdot 1 - (a_0 a_1 - a_1 a_0)(b_0 b_1 - b_1 b_0) \quad (\text{chsh_mul_self}).$$

The identity involves no star, no order, and no norm: it is the exact algebraic content of the Tsirelson bound, stated in a bare ring. The sign convention is that of commutators $[x, y] = xy - yx$, with the minus sign in front of the product of commutators. As noted in Section 3, this lemma's axiom audit reports only `[propext, Quot.sound]`. Its proof engineering, a hand-rolled confluent rewrite system standing in for a missing noncommutative tactic, is described in Section 9.

8.2 The norm bound

The analytic half runs in an abstract typeclass context: a normed ring carrying a star-ring structure satisfying the C*-identity $\|x^*x\| = \|x\|^2$, and nontrivial (so that $\|\mathbf{1}\| = 1$). In Lean the hypotheses are the typeclasses

$$[\text{NormedRing } A] [\text{StarRing } A] [\text{CStarRing } A] [\text{Nontrivial } A].$$

Neither completeness nor any order structure on A is assumed; in particular no ambient `CStarAlgebra` bundle is required.

Proposition 19 (Norm lemmas; algebra-only). *In the context above:*

- (i) *a selfadjoint involution has norm one: if $a^* = a$ and $a^2 = \mathbf{1}$, then $\|a\| = 1$ (*norm_eq_one_of_selfAdjoint_involution*);*
- (ii) *in any normed ring (star and C*-structure are dropped for this item), if $\|a_0\| = \|a_1\| = 1$, then $\|a_0 a_1 - a_1 a_0\| \leq 2$ (*norm_commutator_le_two*).*

Theorem 1 (Tsirelson bound, norm form; algebra-only). *Let A be as above and let $a_0, a_1, b_0, b_1 \in A$ be selfadjoint involutions such that every a_i commutes with every b_j . Then*

$$\|a_0b_0 + a_0b_1 + a_1b_0 - a_1b_1\| \leq 2\sqrt{2} \quad (\text{tsirelson_bound}).$$

Proof sketch. The classical square trick, run entirely on the three preceding statements. By Proposition 18, $S^2 = 4 \cdot \mathbf{1} - [a_0, a_1][b_0, b_1]$; by Proposition 19, each commutator has norm at most 2, so $\|S^2\| \leq 4 + 4 = 8$. The combination S is selfadjoint, so the C*-identity gives $\|S\|^2 = \|S^*S\| = \|S^2\| \leq 8$, and $\|S\| \leq \sqrt{8} = 2\sqrt{2}$. $\square$

No attainment or tightness statement is formalized anywhere in the module: the development ships the bound as an inequality, and nothing in it asserts that $2\sqrt{2}$ is achieved.

8.3 Relation to Mathlib’s CHSH file and the matrix instantiation

Mathlib already contains an order-theoretic form of the Tsirelson inequality: the file `Mathlib/Algebra/Star/CHSH` proves, in a star-ordered algebra over the reals, that the CHSH combination satisfies $S \leq 2\sqrt{2} \cdot \mathbf{1}$ in the algebra’s order, together with the hypothesis bundle `IsCHSHTuple` packaging selfadjointness, involutivity, and the cross commutations. The present development contributes the complementary *norm-form* statement in the C*-setting, with the underlying algebraic identity isolated as Proposition 18, and connects the two through the shared interface:

Proposition 20 (Interoperability and instantiation; algebra-only). *(i) The conclusion of Theorem 1 holds for any tuple satisfying Mathlib’s `IsCHSHTuple` hypothesis bundle (`tsirelson_bound_of_isCHSHTuple`).*

(ii) For every $n \neq 0$, the bound holds in $M_n(\mathbb{C})$ equipped with the L^2 operator norm, activated by the scoped instance set `Matrix.Norms.L2Operator` (`matrix_tsirelson_bound`).

Item (ii) returns the abstract statement to the ambient algebra of the event-algebra modules; combined with Proposition 1(ix), any four events with the appropriate commutation pattern yield dichotomic observables to which the bound applies. Because the abstract theorem is typeclass-generic, the matrix case is an instantiation of a few lines rather than a separate proof; the only genuine work is a hand-constructed nontriviality instance for the matrix algebra (Section 9).

9 Mathlib gaps and proof engineering

This section reports the friction points and affordances encountered against the pinned Mathlib version, as raw material for library builders and for others formalizing quantum-probabilistic content.

9.1 Scoped instances

Three layers of structure needed by this development are scoped. The partial order on $\mathbb{C}$ requires `open scoped ComplexOrder`; without it, goals mentioning $0 \leq z$ or `Matrix.PosSemidef` over $\mathbb{C}$

fail to elaborate, with instance errors that do not point to the missing scope. The Loewner order on matrices is likewise scoped (`MatrixOrder`), as is the ordered-ring structure on $\mathbb{C}$ that supplies `mul_nonneg` in the complex order. Most consequentially, there is no globally installed norm on the matrix type: the L^2 operator norm and its `NormedRing` and `CStarRing` instances activate only under `open scoped Matrix.Norms.L2Operator`, so the matrix instantiation of Theorem 1 must live behind that scope. In the same corner, no `Nontrivial` instance fired for the matrix algebra, and nontriviality (needed for $\|1\| = 1$) had to be constructed by hand, entrywise, from $n \neq 0$. The scoping policy is defensible library design, since competing orders and norms exist on the same types; downstream, it is a real usability tax, and the failure modes when a scope is missing are hard to diagnose.

9.2 Positivity of the trace pairing

Mathlib provides $\text{Tr } M \geq 0$ for a single positive-semidefinite matrix and covers sandwiches BAB^* of positive-semidefinite matrices, yet the bilinear statement, $\text{Tr}(AB) \geq 0$ for positive-semidefinite A and B , is absent. For Born weights this is harmless: an event is idempotent, so $\text{Tr}(\rho P) = \text{Tr}(P\rho P)$ by cycling the trace, and the single-matrix lemma applies. For the expectation functional (Proposition 17(iv)) there is no idempotent to compress with, and the natural abstract route failed in an instructive way. Mathlib’s Gram-factorization characterization of positivity, `CStarAlgebra.nonneg_iff_eq_star_mul_self` combined with the Loewner-order bridge `Matrix.PosSemidef.nonneg`, is gated behind the continuous functional calculus, and instance synthesis for two of its classes, `NonUnitalContinuousFunctionalCalculus` over $\mathbb{R}$ with the `IsSelfAdjoint` predicate and `NonnegSpectrumClass` over $\mathbb{R}$, did not resolve for the matrix algebra carrying the Π topology, even with the order scopes open. (Mathlib’s own matrix-order file uses the lemma successfully inside its own section context; the instance chain did not reproduce downstream for us.) The workaround is elementary and robust: decompose spectrally via `Matrix.IsHermitian.spectral_theorem`, extract nonnegative eigenvalues with `PosSemidef.eigenvalues_nonneg`, cycle the trace, and sum termwise; about a dozen lines.

9.3 No noncommutative linear_combination

The CHSH square identity (Proposition 18) is pure noncommutative-ring algebra with side relations: involutivity of the four letters and the four cross commutations. Lean’s `linear_combination` tactic is commutative-only, and `noncomm_ring` normalizes without consuming hypotheses, so neither closes the goal. The formal proof instead builds a small confluent rewrite system by hand: from each commutation hypothesis it derives the association-compatible, universally quantified form $b(ax) = a(bx)$, and from each involution the cancellation form $a(ax) = x$; running these as a `simp only` set over right-associated words moves every b past every a and cancels adjacent equal letters, and `abel` finishes the additive bookkeeping. The system terminates because each rewrite strictly decreases the number of letters out of normal order. A noncommutative analogue of `linear_combination`, even one restricted to rewriting modulo hypothesis-supplied commutation and involution relations, would compress this style of proof considerably.

9.4 Occurrence-sensitive rewriting

A recurring low-level friction: on goals containing several syntactic occurrences of the same subterm, `rw` rewrites the first occurrence, which in conditioning goals is typically the normalizing scalar $(\mu_\rho(P))^{-1}$ rather than the intended trace argument. The stable pattern was to isolate each trace identity as a standalone `have` with explicit arguments to `trace_mul_comm` or `trace_mul_cycle`, and only then rewrite. A minor lexical note in the same spirit: the script capital letter often used for partitions in informal notation is outside Lean’s legal identifier alphabet in this toolchain, so partition variables are plain ASCII.

9.5 Affordances

Three parts of Mathlib made the development markedly shorter than a bespoke stack would be. First, the `Matrix.PosSemidef` API is definition-agnostic: sandwich closure in both directions (`PosSemidef.mul_mul_conjTranspose_same`, `PosSemidef.conjTranspose_mul_mul_same`), sums (`posSemidef_sum`), scalar multiples (`PosSemidef.smul`, which accepts a complex scalar that is nonnegative in the complex order, so the Lüders normalization is a single application), diagonal nonnegativity, trace nonnegativity, the zero-trace criterion, and the Gram construction `posSemidef_conjTranspose_mul_self` covered every need without ever unfolding the definition. Second, trace faithfulness is in the box: `Matrix.trace_mul_conjTranspose_self_eq_zero_iff` powers the uniqueness of the center expectation, and `PosSemidef.trace_eq_zero_iff` together with `PosSemidef.dotProduct_mulVec_zero_iff` powers the support lemma of Proposition 10(ii). Third, the C*-norm layer, `CStarRing.norm_star_mul_self`, `IsSelfAdjoint.norm_mul_self`, the norm-one instance for nontrivial C*-rings, and `norm_nsmul_le`, reduces the analytic half of Theorem 1 to roughly forty lines over the bare typeclass context, with no completeness assumption anywhere. Mathlib’s existing CHSH file was itself an affordance: its `IsCHSHTuple` bundle fixed the interface against which the norm-form result could be made interoperable from the start.

10 Related work

The Isabelle/HOL line. The nearest prior art is the sequence of developments by Echenim and Mhalla. Echenim’s Archive of Formal Proofs entry on quantum projective measurements and the CHSH inequality [13] formalizes projective measurements and the CHSH expectation in Isabelle/HOL; the journal treatment by Echenim and Mhalla [15] presents a formalization of the CHSH inequality and Tsirelson’s upper bound, with the corresponding AFP entry [16]; a supporting AFP entry formalizes simultaneous diagonalization of pairwise commuting Hermitian matrices [14]. Relative to that line, the present development differs on two axes. Organizationally, it is event-algebra-first: the theory is arranged around the closure algebra of Hermitian idempotents, the block-diagonal commutant with an explicit trace-preserving pinching carrying existence, uniqueness, and contractivity, and a lemma-by-lemma, machine-visible split between algebra-only and trace-dependent content; the Isabelle line is organized matrix-first over a concrete density-matrix and L^2 -norm stack, with the CHSH development as the target. On the prover axis, the ecosystems differ in a way that shows up in proof economy: in Lean/Mathlib the norm-form Tsirelson bound is proved abstractly for any nontrivial unital C*-ring, with the matrix case a short instantiation, and Born-weight positivity and boundedness live directly in the partial order of $\mathbb{C}$ with no bespoke real-part bookkeeping; conversely, the concrete Isabelle stack obtains from its own spectral toolbox some

facts (such as Gram factorization of positive matrices) that are currently gated behind unresolved instance chains in Mathlib, as detailed in Section 9. The CHSH arithmetic itself is common ground; the organization is the delta.

Formalized quantum mathematics more broadly. Bordg, Lachnitt, and He verified quantum algorithms and information-theoretic results in Isabelle/HOL [7], with the accompanying Isabelle Marries Dirac library [6]. In Lean/Mathlib, Dupuis, Lewis, and Macbeth developed formalized functional analysis with semilinear maps [12], infrastructure relevant to any future infinite-dimensional extension of the present work. Garrigue and Saikawa formalize typed compositional quantum computation with lenses [17]. Reports on individual Mathlib formalization efforts, such as Nash on Engel’s theorem [23] and Baanen and collaborators on Dedekind domains and class groups of global fields [3], document both the mathematics and the practice of folding formalization-driven contributions back into the library; the library itself is described in the Mathlib community paper [27], and the Lean 4 system in de Moura and Ullrich [11]. Within Lean specifically, Mathlib already provides an order-form CHSH/Tsirelson inequality and the `IsCHSHTuple` interface used in Section 8, and the Physlib/QuantumInfo project [24] formalizes finite mixed states as positive-semidefinite trace-one matrices, positive operator-valued measures, and a Born-rule measurement. Relative to these, the present development contributes the event-algebra organization, the commutant pinching with its Hilbert–Schmidt characterization, the Lüders fixed-point layer, and the norm-form CHSH bound, packaged together under a per-declaration axiom audit; the novelty claim of Section 1 is exactly this specific combination.

Foundations context. The mathematical foundations of quantum probability remain under active discussion in this journal’s scope. Reddiger and Poirier examine the one-body Born rule in curved spacetime [26]; Reddiger analyzes the applicability of Kolmogorov’s probability theory to quantum phenomena [25]; Aharonov and Shushi revisit the derivation of the Born rule [1], and Beck surveys the textbook orthodoxy on measurement and state update [4]. Arsiwalla, Chester, and Kauffman reconstruct expectation values from a pre-quantum operator algebra, treating classical and quantum wave functions as consequences of that operator layer [2]; Kumar and Pan study network nonlocality [20], and Held derives the Tsirelson bound from geometric-algebra assumptions [19]. Goldstein bootstraps infinite-dimensional state spaces from finite ones, deriving the von Neumann–Lüders postulate for projective observables [18]. The present paper contributes to that discussion at the level of certified bookkeeping: for the finite core, every hypothesis of every standard statement is now explicit and machine-audited, so debates about what follows from what can point to a checked artifact.

11 Conclusion and outlook

The finite event-algebra core of quantum probability admits a complete machine-checked treatment at modest cost: 64 lemmas and 8 definitions cover event closure, Born weights with reality and the probabilistic bounds in the strong complex-order form, repeatable Lüders conditioning with its fixed-point characterization, the trace-preserving pinching onto the block-diagonal commutant with its Hilbert–Schmidt characterization (projector, state-preserving, selfadjoint, Pythagorean, contractive, trace-compatible, unique among commutant-valued trace-compatible maps, and compatible with conditioning on commutant events), a positive normalized expectation functional, and the Tsirelson bound in an abstract C^* -setting. The axiom audit is part of the source and re-executes on every build; the lemma-level split between algebra-only and state-consuming content gives the

development a second, structural payload beyond the individual theorems: it is a machine-visible map of where probability enters the quantum formalism.

Three extensions are natural. First, unsharp measurements: effects ($0 \leq E \leq \mathbf{1}$ in the Loewner order) and instruments generalize events and the Lüders map, and the certainty-set and fixed-point interface of Section 5 suggests the shape the generalized statements should take. Second, attainment: complementing Theorem 1 with a formalized two-qubit witness achieving $2\sqrt{2}$ would close the gap between the inequality and the bound’s sharpness, and is well within reach of the matrix instantiation already in place. Third, infinite dimensions: the finite development deliberately avoids completeness, so its Tsirelson theorem already applies to any nontrivial unital C^* -ring, yet the event-and-state layer is genuinely finite. Goldstein’s bootstrap of infinite-dimensional state spaces from finite ones in this journal [18] marks exactly the frontier a formal development would next confront: which parts of the finite bookkeeping survive the passage verbatim, and which require new analysis (normality of states, σ -additivity, von Neumann algebras in place of matrix algebras). The conditional-expectation package of Section 6 is the component we expect to generalize most cleanly, since its statements are already phrased against a commutant rather than against matrix entries.

When the bookkeeping of a foundational formalism is mechanical and audited, discussion can concentrate on the questions that are actually open. For finite quantum event algebras, that state has now been reached.

Declarations

Funding. No funding was received for this work.

Competing interests. The author declares no competing interests.

Author contributions. Bernhard Mueller is the sole author: he designed and wrote the Lean development, conducted the axiom audit, and wrote the manuscript.

Data availability. No data were generated or analyzed in this work.

Use of AI tools. The author used AI-assisted tools in drafting the Lean development and the manuscript. All formal statements and proofs were checked by the Lean 4 proof assistant, the axiom audits were machine-executed, and the author reviewed and takes full responsibility for the content.

Code availability. The complete Lean development [22] accompanies the submission as a standalone Lake project with a pinned external Mathlib dependency; it does not vendor Mathlib. It builds free of `sorry` placeholders with the toolchain `leanprover/lean4:v4.29.1` and the Mathlib version pinned in the project manifest, and the complete per-declaration axiom audit re-executes as part of the build. A public release with a persistent identifier and an open-source license will accompany publication.

References

- [1] Yakir Aharonov and Tomer Shushi. Improving the proof of the Born rule using a physical requirement on the dynamics of quantum particles. *Quantum Studies: Mathematics and Foundations*, 11:287–290, 2024. DOI: 10.1007/s40509-024-00322-x.

- [2] Xerxes D. Arsiwalla, David Chester, and Louis H. Kauffman. On the operator origins of classical and quantum wave functions. *Quantum Studies: Mathematics and Foundations*, 11: 193–215, 2024. DOI: 10.1007/s40509-023-00311-6.
- [3] Anne Baanen, Sander R. Dahmen, Ashvni Narayanan, and Filip A. E. Nuccio Mortarino Majno di Capriglio. A formalization of Dedekind domains and class groups of global fields. *Journal of Automated Reasoning*, 66:611–637, 2022.
- [4] Geoff Beck. How to be an orthodox quantum mechanic. *Quantum Studies: Mathematics and Foundations*, 13:Article 6, 2026. DOI: 10.1007/s40509-025-00378-3.
- [5] John S. Bell. On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika*, 1(3):195–200, 1964.
- [6] Anthony Bordg, Hanna Lachnitt, and Yijun He. Isabelle marries Dirac: a library for quantum computation and quantum information. *Archive of Formal Proofs*, 2020. Formal proof development, ISSN 2150-914x.
- [7] Anthony Bordg, Hanna Lachnitt, and Yijun He. Certified quantum computation in Isabelle/HOL. *Journal of Automated Reasoning*, 65:691–709, 2021.
- [8] Paul Busch and Pekka Lahti. Lüders rule. In Daniel Greenberger, Klaus Hentschel, and Friedel Weinert, editors, *Compendium of Quantum Physics*. Springer, Berlin, Heidelberg, 2009.
- [9] Boris S. Cirel’son. Quantum generalizations of Bell’s inequality. *Letters in Mathematical Physics*, 4(2):93–100, 1980.
- [10] John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt. Proposed experiment to test local hidden-variable theories. *Physical Review Letters*, 23(15):880–884, 1969.
- [11] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021.
- [12] Frédéric Dupuis, Robert Y. Lewis, and Heather Macbeth. Formalized functional analysis with semilinear maps. *Journal of Automated Reasoning*, 68(1):Article 4, 2024.
- [13] Mnacho Echenim. Quantum projective measurements and the CHSH inequality. *Archive of Formal Proofs*, 2021. Formal proof development, ISSN 2150-914x.
- [14] Mnacho Echenim. Simultaneous diagonalization of pairwise commuting hermitian matrices. *Archive of Formal Proofs*, 2022. Formal proof development, ISSN 2150-914x.
- [15] Mnacho Echenim and Mehdi Mhalla. A formalization of the CHSH inequality and Tsirelson’s upper-bound in Isabelle/HOL. *Journal of Automated Reasoning*, 68:Article 2, 2024. DOI: 10.1007/s10817-023-09689-9.
- [16] Mnacho Echenim, Mehdi Mhalla, and Coraline Mori. The CHSH inequality: Tsirelson’s upper-bound and other results. *Archive of Formal Proofs*, 2023. Formal proof development, ISSN 2150-914x.
- [17] Jacques Garrigue and Takafumi Saikawa. Typed compositional quantum computation with lenses. *Journal of Automated Reasoning*, 70:Article 10, 2026. DOI: 10.1007/s10817-026-09750-3.

- [18] Norman J. Goldstein. State spaces: Infinite dimensions from finite. *Quantum Studies: Mathematics and Foundations*, 13:Article 11, 2026. DOI: 10.1007/s40509-025-00381-8.
- [19] Carsten Held. Quantum-mechanical correlations and Tsirelson bound from geometric algebra, 2021. Preprint.
- [20] Rahul Kumar and A. K. Pan. Sharing nonlocality in a network using the quantum violation of chain network inequality. *Quantum Studies: Mathematics and Foundations*, 10:353–372, 2023. DOI: 10.1007/s40509-023-00300-9.
- [21] Gerhart Lüders. Über die zustandsänderung durch den meßprozeß. *Annalen der Physik*, 8: 322–328, 1950. English translation by K. A. Kirkpatrick, *Concerning the state-change due to the measurement process*, *Annalen der Physik* **15**(9), 663–670 (2006).
- [22] Bernhard Mueller. EventAlgebra: A Lean 4 formalization of finite quantum event algebras, 2026. Software artifact accompanying this article; built with toolchain `leanprover/lean4:v4.29.1` over Mathlib.
- [23] Oliver Nash. Engel’s theorem in mathlib, 2023. Preprint, arXiv.
- [24] Physlib/QuantumInfo contributors. QuantumInfo: A Lean 4 library for quantum information, 2026. Software project; finite mixed states as positive-semidefinite trace-one matrices, POVMs, and Born-rule measurement.
- [25] Maik Reddiger. On the applicability of Kolmogorov’s theory of probability to the description of quantum phenomena. Part I: Foundations. *Quantum Studies: Mathematics and Foundations*, 13:Article 1, 2026. DOI: 10.1007/s40509-025-00375-6.
- [26] Maik Reddiger and Bill Poirier. Towards a probabilistic foundation of relativistic quantum theory: the one-body Born rule in curved spacetime. *Quantum Studies: Mathematics and Foundations*, 12:Article 5, 2025. DOI: 10.1007/s40509-024-00349-0.
- [27] The mathlib Community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020)*, pages 367–381. ACM, 2020.
- [28] John von Neumann. *Mathematische Grundlagen der Quantenmechanik*. Springer, Berlin, 1932. English translation: *Mathematical Foundations of Quantum Mechanics*, Princeton University Press, Princeton, 1955.