Superdeterminism and Quantum Mechanics: A Compatibility Audit

Forked from the Born-counting program: theorems, prices, and open entailments

Research summary

19 August 2026

Abstract

This document organizes the results of the parent Born-counting program (whose cited artifacts are copied verbatim into `parent/`; the cross-project linkage is recorded in the phase notebook) and its compatibility fork (twenty-five further machine-verified phases, notebook `models/NOTES.md`) around the founding question: *is quantum mechanics, together with its contemporary entailments, fully compatible with superdeterminism?* The answer splits by what "superdeterminism" is allowed to be. Forward-causal smooth models: *no*, exactly and provably, already for one singlet (a C^1 rare-sector obstruction, a finite-affinity no-go, and a measure-class no-go force two-boundary structure that forward dynamics cannot generate). Finite or profinite counting ontologies: *no at every resolution* — and quantitatively, the deviation–spread law gives a resolution-independent floor $(\sigma/2 - 1/6)/(\frac{8}{3}\sigma)$ above spread fraction $\frac{1}{3}$: refining the ontic grid is not a route to the Born statistics. Two-boundary models: *yes, constructively and at quantified cost*, on a sector that the fork's phases extended from the qubit/singlet base to all pure two-qubit states, entanglement swapping, GHZ, unsharp qubit pairs, and all finite-dimensional pure projective measurements with sequences — via one recurring construction, the record-conditioned steering chain, whose telescoping realizes measure noncontextuality and additivity by certified algebra (the Busch–Gleason frame) while keeping outcome sets contextual as Kochen–Specker requires. Structure theorems price the compliance: equivariant models exist iff the state is unentangled (entanglement is exactly the property that forces the settings into the parity-odd sectors of the carrier measure); the unique one-wing singlet model carries a forced 0.2506-bit information premium over the two-setting minimum; the sharp-measurement limit uniquely forces the cusped basin, with finite affinity $\log \frac{1+3\eta}{1-\eta}$ sufficing below it. The conspiracy objection resolves for the entire verified sector: every exact no-signalling identity is the fixed locus of an identified structure (source parity, branch normalization, record-affinity) — none is a tuning. The hard implementation-conjugacy item is closed: record-affinity is *derived* — the law of total probability for the canonical single-law microdynamics, whose measurement map is a measure-preserving record toggle reading an ontic apparatus coordinate; implementation conjugacy is derived on the full rotation orbit by canonical symplectic automorphisms; and across chain skeletons the rank-chart intertwiner supplies a *unique* canonical transport between implementations of the same effect, its uniqueness a kernel-checked rigidity theorem and its data the skeleton's own filtration order. What remains open is only secondary: the residual deviation bracket.

1 The question and its grades

The founding question of this program is whether quantum mechanics—not only its statics but everything it now entails operationally—is fully compatible with a superdeterministic ontology. The parent Born-counting archive pursued the stronger question of *deriving* the Born rule as a natural counting measure; in the course of that pursuit it proved a set of results that answer large parts of the compatibility question, without ever assembling them under that heading. This fork performs the assembly and states what remains.

Compatibility must be graded, because the theorems below separate the grades.

Definition 1.1 (Compatibility grades)**.** A deterministic model is *L1-compatible* (law-level exact) if it reproduces the quantum probability law exactly for every state and every measurement context in the relevant sector, including the counterfactual continuum of settings. It is *L2-compatible* (finite precision) if its predictions deviate from the quantum law by an amount bounded by a resolution parameter, uniformly in the sector. It is *L3-compatible* (history-level adequate) if it reproduces the finitely many outcome frequencies of the experiments actually performed in one history.

L3 is nearly vacuous—any finite dataset has rational frequencies and admits finite deterministic models—but purchasing exactness at L3 while failing L1 requires the model to encode which experiments occur, which is the conspiracy objection at full strength (Section 6). The scientifically contentful grades are L1 and L2.

Definition 1.2 (Model classes)**.** *SD-F* (forward superdeterminism): deterministic invertible dynamics on a connected phase space; hidden-variable densities smooth and strictly positive on their supports; setting correlations established through the common past; conditioning only on past data. *SD-C* (counting superdeterminism): finite or profinite ontic state space carrying (near-)uniform counting measure; the natural reading of digital ontologies such as 't Hooft's cellular automaton interpretation [28] and causal sets. *SD-T* (tuned superdeterminism): finite or continuum ontology with setting-dependent hidden-variable weights inserted to match the quantum statistics, as in Brans [7], Hall [16], and the permissive reading of [29]. *SD-B* (boundary superdeterminism): deterministic invertible dynamics together with a final or two-time boundary condition on admissible histories; the history measure fixed by conditioning the invariant measure on the boundary event. This is the class long advocated by the Lagrangian-schema and locally-mediated programs [30, 31, 33].

The classes are not rhetorical: the theorems below give SD-F an exact impossibility result, SD-C a resolution-independent floor, SD-T an exactness-for-naturalness trade, and SD-B a verified compatible sector with a quantified price.

2 Exact incompatibility of forward superdeterminism

The three theorems of this section are the sharpest results of the program. Together they show that L1 compatibility already fails for SD-F at the level of one singlet, for structural rather than quantitative reasons. All three are machine-checked.

Theorem 2.1 (C^1 rare-sector obstruction; moving-lune trace)**.** *Let a be fixed, let b_t make angle t with a, and let $G_-(t)$ be the hemispheric disagreement lune $\{n : \operatorname{sign}(a \cdot n) = \operatorname{sign}(b_t \cdot n)\}^c$. For every jointly continuous nonnegative density $\varrho(v, n)$ on settings $\times$ sphere, the disagreement mass satisfies the trace formula*

$$\lim_{t \to 0^+} \frac{P_\varrho(t)}{t} = \frac{1}{4\pi} \int_0^{2\pi} \varrho(0, n(\phi, 0))\, |\cos \phi|\, d\phi.$$

If $P_\varrho(t) = \sin^2(t/2)$, the trace must vanish, forcing ϱ to vanish on the response equator; if ϱ is jointly C^1, every first derivative also vanishes there and $P_\varrho(t) = o(t^2)$, contradicting $\sin^2(t/2) = t^2/4 + O(t^4)$. The exact coefficients mark the dichotomy: a uniform positive trace gives $P(t) \sim t/\pi$; the cusp density $f = 2|z|$ gives $t^2/4$; the smooth nodal density $f = 3z^2$ gives $2t^3/(3\pi)$. The Born coefficient sits exactly at the non-C^1 cusp.

Provenance: Parent Phase 39, the one-sided Liouville singlet; `parent/S4OneSidedLiouvilleSinglet.lean`, `parent/s4_one_sided_liouv` trace integrals verified independently in Mathematica.

Theorem 2.2 (Finite terminal affinity fails exactly)**.** *In the common-future compiler realization of the singlet, give accepted histories weight one and rejected histories relative weight $\delta \in [0,1]$. The resulting correlation obeys*

$$E_\delta(\theta) + \cos\theta = \frac{4\delta\,(\cos\theta - T_\theta)}{1+3\delta}, \qquad T_\theta = 1 - \tfrac{2\theta}{\pi},$$

with small-angle slope $m_\delta = 8q_{\mathrm{rej}}/(3\pi)$, $q_{\mathrm{rej}} = 3\delta/(1+3\delta)$. Exact singlet behavior forces $\delta = 0$: every strictly positive likelihood floor, hence every finite Gibbs affinity $\delta = e^{-\kappa}$, fails exactly. The hard boundary is a finite-information point ($D(Q_\delta\|\mu) \to \log 4$) reachable only with infinite likelihood contrast and exact loss of support.

Provenance: Phase 47; `parent/S4SoftTerminalAffinity.lean`, `parent/s4_soft_terminal_affinity.py`; Z3 positive-leakage no-go.

Theorem 2.3 (Measure-class no-go)**.** *Let λ be Liouville measure, ρ_0 an initial density with $0 < \rho_0 < \infty$ a.e., Φ_T a finite-time Hamiltonian flow, and ℓ a finite positive terminal likelihood. Then the tilted final law $\rho_{T,\ell} \propto \ell \cdot (\rho_0 \circ \Phi_T^{-1})$ is equivalent to λ: every rejected region of positive Liouville volume retains positive probability. Adjoining a finite-dimensional full-support apparatus does not evade the result, and the basin of attraction of any closed invariant null set is itself null. Consequently the exact support zero required by Theorems 2.1 and 2.2 cannot be generated by finite-time forward Hamiltonian dynamics from any full-support preparation; it must be imposed as admissible-state structure or as a boundary condition.*

Provenance: Phase 48; `parent/S4MeasureClassNoGo.lean`, `parent/s4_measure_class_no_go.py`.

Corollary 2.4 (Verdict on SD-F)**.** *No SD-F model is L1-compatible with quantum mechanics, already for the singlet's projective statistics. SD-F models remain L2-compatible with quantified deviation: the finite-resolution boundary kernel $K_\gamma(p \,|\, [a]) \propto e^{\gamma(a\cdot p)^2}$ yields a smooth strictly positive source density whose equatorial value is $r_\gamma \sim 2/\sqrt{\pi\gamma}$ and whose correlation retains the linear cusp*

$$E_\gamma(a, b_\theta) = -1 + \tfrac{2r_\gamma}{\pi}\,\theta + o(\theta),$$

recovering the singlet only as $\gamma \to \infty$. The small-angle repeatability slope is therefore the universal experimental signature of the entire SD-F class.

Provenance: Phase 42; `parent/S4BoundaryApparatus.lean`, `parent/s4_boundary_apparatus.py`.

Remark 2.5 (The survival filter: a canonical approximant with a derived scale). The parent program contains a second, dynamical route to the boundary density, which its exact constructions superseded and this audit reinstates as the canonical SD-F approximant. Conditioning a reversible tape dilation on past survival in the preparation window (a quasi-stationary filter) selects the first Dirichlet mode: the exact cosine profile on the circle, $(p \cdot n)_+/\pi$ on the hemisphere, and the full projective Born overlap $(1 + p \cdot a)/2$ in the infinite-conditioning limit — inside a time-reversible, $SO(3)$-covariant dilation. At finite conditioning depth the deviation is controlled by the spectral gap of the survival kernel: a *derived* deviation scale, unlike the free parameter γ of the soft-boundary kernel above. Exactness is an infinite-duration conditional limit — a zero-measure event in the unconditioned ensemble — consistent with Theorem 2.3: the filter approaches, and never dynamically reaches, the boundary condition.

Provenance: Parent Phase 7; `parent/s4_reversible_survival.py`, `parent/s4_spherical_survival.py`, `parent/S4Survival.lean`.

Remark 2.6. Theorem 2.3 is the fork in the road for the whole program: exactness pushes every deterministic completion out of the initial-value class. What survives is SD-B—the final condition read either as postselection (one-boundary language) or as a law-like two-time boundary (two-boundary language)—and the maximum-caliber theorem of Section 4 shows the measure is then unique given the boundary. The two-time formulation is exactly the Lagrangian schema Wharton has long advocated [30, 31]; the theorem upgrades that schema from an available alternative to a forced one, and the locally-mediated framework surveyed by Wharton and Argaman [33] is the model class in which Section 4's constructions live. The residual physical question is the status of the boundary itself (Section 6).

3 Finite ontologies are finite-precision

The results of this section bound SD-C away from L1 at *every* resolution, by arithmetic, cardinality, and topology rather than dynamics.

Proposition 3.1 (Cardinality floor)**.** *Let Ω be a finite ontic space with a setting-independent probability weight and deterministic responses $A_a : \Omega \to \{\pm 1\}$. As a ranges over a continuum of settings, the outcome probability $P(+\,|\,a)$ takes at most $2^{|\Omega|}$ distinct values and therefore cannot equal $\cos^2(\theta/2)$ on any interval. Exactness across a settings continuum requires infinitely many ontic states or a setting-dependent measure.*

Proof. $P(+\,|\,a)$ is the weight of $A_a^{-1}(+1) \subseteq \Omega$, and Ω has $2^{|\Omega|}$ subsets. ☐

Theorem 3.2 (Niven floor)**.** *Under uniform counting measure on a finite ontic space, every outcome probability is rational. Since $\cos^2(\theta/2)$ is rational at rational angles only for $\theta \in \{0°, 60°, 90°, 120°, 180°\}$ [23], a counting model can be exactly Born at most on this five-angle set, at every resolution K.*

Provenance: Parent archive, opening of the S4 analysis; parent Phase 1 setup (linkage in `models/NOTES.md`).

Even on the five-angle set the discrete structure is severely constrained:

Theorem 3.3 (Trine frustration and the parity no-go)**.** *(i) For rotation-covariant antipodal cell models on $\mathbb{Z}_K$, the shift-$K/3$ three-cycle bound gives $\gamma(K/3) \geq -K/3$, while the Born value at $120°$ demands $\gamma(K/3) = -K/2$: impossible at every K. The $90°$ value is forced for free by antiperiodicity, and $60°$ is equivalent to $120°$. (ii) Without any covariance assumption, no three fixed ± 1 outcome cells with uniform conditional measures can be pairwise anticorrelated at $-\frac{1}{2}$ on any nonempty space: $\sum_x (g_1 + g_2 + g_3)^2$ would be 0 yet each summand is ≥ 1. Hence the outcome function of any exact finite counting model must depend on apparatus hidden variables; hidden-measurement structure is forced, not chosen.*

Provenance: Phase 1; Lean theorem `no_exact_trine_cells` (kernel-checked, project `borncount`); `parent/niven_qubit.py` (exhaustive $K = 24$, Z3 $K = 48$).

Theorem 3.4 (One-third law)**.** *In the outcome-function class $j = f(\lambda - a, u)$ with antipodality, repeatability, and exact Born values at the Niven angles: the maximum preparation spread is exactly $K/3$ (proved at $K = 24$ and $K = 48$, apparatus-independently; lower bound for all $K = 24m$ by cell splitting), while the classical linear profile $P(\Delta) = 1 - \Delta/12$ reaches full spread $K/2$. In the continuum the bound is sharp: a Born-exact preparation ensemble occupies at most $1/3$ of the circle. Carrying Born rather than Bell's linear law costs exactly one third of the epistemicity budget.*

Provenance: Phases 2–4; `parent/outcome_model.py`, `parent/unsat_t17.py`, `parent/continuum_one_third.py`; Lean theorems `witness_certified` and `six_cycle_bound` (`parent/SixCycle.lean`).

Theorem 3.5 (Profinite setting obstruction)**.** *Every continuous map from a connected space to a totally disconnected space is constant; hence every continuous action of the setting circle or sphere on a finite or profinite ontic space is trivial. Exact rotation covariance is impossible for SD-C; covariance can only be emergent and approximate.*

Provenance: Phase 5; `parent/S4Topology.lean`; the parent archive's profinite-action lemma.

Two imported theorems close the circle from the operational side. Hardy's excess-baggage theorem shows that reproducing the quantum statistics of even one qubit for all preparations and measurements requires infinitely many ontic states [19]; Montina's dimension bounds sharpen this under Markovian dynamics [22]; and Pfister–Wehner show a discrete operational state space obeying a nondisturbance postulate is classical [25].

Corollary 3.6 (Verdict on SD-C)**.** *No SD-C model is L1-compatible with quantum mechanics, at any resolution. SD-C is an intrinsically L2 class, with deviation floors given by Theorems 3.2–3.5. In particular, any framing in which the Planck scale is read* ontically*—finitely many ontic microstates per bounded region, counting measure—is a finite-precision model of quantum mechanics. Read merely* operationally *(minimal length, discrete spectra, finite horizon entropy), the Planck scale forces nothing of the kind: a finite-dimensional Hilbert space has the connected continuum $\mathbb{CP}^{N-1}$ behind it, and the compact Liouville models of Section 4 have finite total phase volume with continuum-valued basin ratios.*

Remark 3.7 (The 't Hooft trichotomy). The cellular automaton interpretation [28, 29] therefore splits three ways. With counting weights it is SD-C: inexact at every resolution (Proposition 3.1, Theorems 3.2, 3.5), consistent with 't Hooft's own expectation that quantum mechanics is emergent and that large-scale quantum computation may fail. With tuned ontic weights it is SD-T: exactness is possible, the weights are stationary under permutation dynamics (ontology conservation), but nothing selects them—the Born structure is inserted at $t = -\infty$, the same boundary-condition move as SD-B without the conditioning rationale, and squarely exposed to the fine-tuning objection. Only the continuum SD-B sector below achieves exactness with a *derived* measure.

4 The verified compatible sector

Under the two-boundary repair, exact L1 compatibility is machine-verified for the following sector. Everything in this section is constructive, reversible, and checked in Lean, Z3/exact arithmetic, and Mathematica.

Theorem 4.1 (Pure-qubit basin)**.** *On T^*S^2 with Liouville measure, the covariant dipole Hamiltonian*

$$H_p(n,\xi) = \frac{|\xi|^2}{2m} + U_0 - \kappa\, p \cdot n$$

has invariant sublevel $\Gamma_p = \{H_p \le U_0\}$ with normalized base marginal exactly $(p \cdot n)_+/\pi$; a measurement along a toggles a record bit by an involution with outcome volume exactly $(1 + p \cdot a)/2$. The law is independent of m, κ, U_0; a minimal kinetic cap gives a reference filter of success probability $1/4$ whose conditional state is exactly uniform on Γ_p; cotangent-lifted rotations transport it to every preparation and setting. With purity and repeatability the sequential law

$$P(r, s \,|\, p, a, b) = \frac{1 + r\, p \cdot a}{2} \cdot \frac{1 + rs\, a \cdot b}{2}$$

follows at all depths, with an invertible stationary tape dilation.

Provenance: Phases 8–9; `parent/s4_hamiltonian_basin.py`, `parent/s4_sequential_tape.py`; `parent/S4HamiltonianRecord.lean`, `parent/S4Sequential.lean`.

Theorem 4.2 (Singlet from basins; the cusp derived)**.** *The biased-source singlet density of Degorre–Laplante–Roland [12], $f_a(n) = 2|a \cdot n|$, is exactly the equal mixture of the two dipole preparation basins: $f_a\, d\mu = \frac{1}{2}\rho_a\, d\Omega + \frac{1}{2}\rho_{-a}\, d\Omega$ with $\rho_p = (p \cdot n)_+/\pi$. Its disagreement mass is $\sin^2(\theta/2)$: the linear-to-quadratic rare-sector suppression is derived from the closing momentum-fiber area, not imposed. One compact connected autonomous Liouville system on $S^2_p \times T^*S^2_n$ (total volume $8\pi^3 m\kappa$) realizes the full channel with consistent $\mathbb{Z}_2$ projective holonomy, and explains the cusp: the momentum disk collapses at $p \cdot n = 0$, so the projection is critical exactly where Theorem 2.1 requires.*

Provenance: Phases 39–40; `parent/S4OneSidedLiouvilleSinglet.lean`, `parent/S4EndogenousSettingDilation.lean` and paired scripts.

Theorem 4.3 (Two-wing Lüders law from local interactions)**.** *Let the source emit an antipodal pair $(n, -n)$ and a common-past selector bit J; let each wing carry only the local endpoint dipole constraint of Theorem 4.1, active in the sector its label selects; and condition histories on the joint terminal event. The conditioned source density is $f_{ab}(n) = |a \cdot n| + |b \cdot n|$, the singlet joint law $P(\alpha, \beta) = (1 - \alpha\beta\, a \cdot b)/4$ follows with zero marginals and operational no-signalling, and for arbitrary finite local projective sequences the full two-qubit Lüders law*

$$P(\alpha_1 \ldots \alpha_r, \beta_1 \ldots \beta_s) = \frac{1 - \alpha_1\beta_1\, a_1 \cdot b_1}{4} \prod_{i=2}^{r} \frac{1 + \alpha_{i-1}\alpha_i\, a_{i-1} \cdot a_i}{2} \prod_{j=2}^{s} \frac{1 + \beta_{j-1}\beta_j\, b_{j-1} \cdot b_j}{2}$$

holds independently of the spacelike interleaving. No post-emission cross-wing interaction is used.

Provenance: Phase 44; `parent/S4CommonSelectorBoundary.lean`, `parent/s4_common_selector_boundary.py` (includes a five-record direct two-qubit Lüders check).

Theorem 4.4 (The architecture is forced)**.** *Within the four-mode local gate class $s_{ab}(n) = u_0 + u_A|a \cdot n| + u_B|b \cdot n| + u_2|a \cdot n||b \cdot n|$, exact singlet behavior forces $u_0 = u_2 = 0$ (the small-angle linear term kills u_0; strict Cauchy–Schwarz on $M_\theta < \frac{1}{3}$ kills u_2): exactly one endpoint is active. Resolved outcome statistics cannot identify the wing weight λ, but the intrinsic source information $I_\theta(\lambda)$ is strictly convex with unique minimum at $\lambda = \frac{1}{2}$: maximum caliber, not data, selects the exchange-symmetric selector. The reversible compiler's ready bit must be exactly pure ($\varepsilon = 0$, negentropy $\log 2$).*

Provenance: Phases 45–46; `parent/S4ExclusiveSelectorUniqueness.lean`, `parent/S4CommonFutureCompiler.lean` and paired scripts.

Theorem 4.5 (Measure uniqueness given the boundary)**.** *On each oriented component of the endpoint cap, any smooth positive density invariant under every compactly supported Hamiltonian flow is constant, and projective symmetry equates the two constants: canonical naturality selects normalized Liouville measure as reference. Given the terminal compatibility event C, the minimum-relative-entropy (maximum-caliber) update is $\mu(\cdot \mid C)$, with $\mu_R(C_{[a]}) = \frac{1}{4}$ and cost $\log 4$ per measurement; conditioning composes associatively, tensorizes over independent apparatuses, and commutes with deterministic transport. N independent copies have compatible fraction 4^{-N}.*

Provenance: Phase 43; `parent/S4CanonicalHistoryMeasure.lean`, `parent/s4_canonical_history_measure.py`.

The price list

The measurement dependence this sector requires is not a free parameter; it is bounded and, at the boundaries, uniquely structured.

Theorem 4.6 (Measurement-dependence floors)**.** *(i) With local deterministic responses and hidden distributions ρ_{xy} of total-variation diameter D, $|S_{\mathrm{CHSH}}| \leq 2 + 6D$; the singlet value $2\sqrt{2}$ forces $D \geq (\sqrt{2} - 1)/3 = 0.138071\ldots$, and the bound is saturated by an explicit eight-table model with unbiased marginals. (ii) The same channel is the unique minimum-information setting-dependent channel at the CHSH point, with rate–distortion value $R(\delta_*) = 0.0320745\ldots$ nats $= 0.0462738\ldots$ bits, $\delta_* = (2 - \sqrt{2})/4$, realized by a Gibbs family with $\kappa_* = \log \frac{3+2\sqrt{2}}{3}$. (iii) Hall's continuous model [16] is its exact continuous extension by minimum relative entropy given the Born sector masses, costing $0.0280347\ldots$ bits on Haar-average; the volume-derived model of Theorem 4.2, which* derives *the sector masses, costs $\log 2 - \frac{1}{2}$ nats $= 0.278652\ldots$ bits. Whether the tenfold premium is the necessary price of derivation is open (Section 9).*

Provenance: Phases 36–39; `parent/S4FiniteSpeedBell.lean`, `parent/S4MinimumInformationBell.lean`, `parent/S4ContinuousHallBr` and paired scripts; exact LP optimum confirmed in Mathematica.

Theorem 4.7 (Intervention theorem)**.** *In the connected realization of Theorem 4.2, a modular post-emission overwrite $\mathrm{do}(L{=}[a])$ that leaves the source state unchanged yields the Bell-local triangular law $E_{\mathrm{do}}(\theta) = -1 + 2\theta/\pi$ with $|S| = 2$, maximal correlation gap $0.2105\ldots$ at $\theta_* = \arcsin(2/\pi)$. Restoring the singlet requires re-conditioning on the boundary (TV cost $\frac{1}{4}$, information cost $\log 2 - \frac{1}{2}$ nats), which changes the counted ensemble of complete histories. Settings are therefore necessarily endogenous: every physical mechanism that could implement a genuine intervention is, in an exact model, already enrolled in the boundary correlation.*

Provenance: Phase 41; `parent/S4InterventionalFork.lean`, `parent/s4_interventional_fork.py`.

Corollary 4.8 (Verdict on SD-B)**.** *SD-B is L1-compatible with quantum mechanics on the mined sector $\{$pure-qubit projective sequences$\} \cup \{$singlet with arbitrary finite local projective sequences$\}$, with operational no-signalling, purely local spacetime interactions, and the quantified prices of Theorems 4.5–4.7; the fork's phases extend this sector to all pure and mixed states, all POVMs, all finite dimensions, entangled measurements, networks, circuits, and the two-mode field core (Sections 5 and 7). This is compatibility, not derivation: the terminal support condition is an axiom* *that Theorem 2.3 proves cannot be dynamically generated — no finite-time forward Hamiltonian evolution from a full-support preparation produces exact support zeros, and the basin of any invariant null set is null — and Theorem 4.5 makes the measure unique only* given *that axiom.*

5 The audit: contemporary entailments

Full compatibility means every operational entailment of quantum mechanics. The current state of the audit:

Entailment	Status	Sharpest known obstacle / next step
Qubit projective sequences	**verified**	Theorem 4.1
Singlet projective sequences	**verified**	Theorems 4.2–4.4
Arbitrary bipartite pure states	**closed**	exact a.c. setting-conditioned model for every pure two-qubit state (steering dilation); equivariant models exist iff the state is a product (saturation collapse, Phases 3–5); the two-wing chain extends to every pure state of $\mathbb{C}^{d_A} \otimes \mathbb{C}^{d_B}$ with CGLMP carried at Schmidt rank 3 (Phase 18)
POVMs / unsharp measurement	**closed**	proven: strict unsharpness admits exact smooth full-support densities with finite affinity $\log \frac{1+3\eta}{1-\eta}$; the sharp limit forces the dipole cusp uniquely by mass saturation (notebook Phase 7); arbitrary POVMs — non-covariant, informationally complete, any outcome count — by the Naimark chain (Phase 21)
$d \geq 3$, Kochen–Specker, Gleason	closed	the chain-of-effective-qubits dilation realizes measure noncontextuality and additivity by certified algebra, with contextual outcome sets (the KS-compatible split; notebook Phase 14); record-affinity and rotation-orbit conjugacy derived from the canonical single-law microdynamics (Phase 15); the rank-chart intertwiner closes hard conjugacy across skeletons (Phase 16)
Entanglement swapping / event-ready Bell	**closed**	the chained steering dilation reproduces the swapped law; $\sum_\mu T(\mu) = 0$, so the conditioning *constitutes* the correlation (notebook Phase 6); the station itself realized as a carrier chain with universal basins — no black box, event-ready derived (Phase 17)
GHZ / multipartite / monogamy	**closed**	the steering chain carries the exact GHZ law incl. the Mermin paradox ($M = 4$ vs. LHV 2); pairwise sectors classical (monogamy as architecture); the induction certified on W and cluster instances, depth-two recursion exact, with the W pair CHSH-local unconditionally and maximal conditionally (Phase 19)
Device-independent certification	**closed** (quantified)	fails under SD at a sharp, kernel-certified price: faking CHSH S needs measurement dependence $D^*(S) = (S-2)/6$ — $(\sqrt{2}-1)/3 \approx 13.8\%$ total variation at the Tsirelson point, 1/3 at the algebraic maximum (Phase 22)
Quantum computation at scale	**closed**	circuits run on the chain (Grover exact at 25/32, 121/128 on carriers); the 4^{-N} fraction costs a *linear* bit budget (2 bits per measurement, additive); the exponential sits in boundary data, so sampling-hardness survives (Phase 24)

Entailment	Status	Sharpest known obstacle / next step *(continued)*
Relativistic QFT sector	opened (two-mode core)	the squeezed vacuum's entanglement, thermality/RS shadow, and Bell violation up to the Tsirelson limit all run on the chain with exponentially controlled truncation (Phase 25); algebraic QFT proper — type III factors, exact microcausality, genuine RS cyclicity — remains open

At the fork, two rows were singled out: the $d \geq 3$ row as where "fully" lives (by [8, 9] the entire Born content in every dimension reduces to measure-noncontextuality plus additivity, and Kochen–Specker [21] obstructs neither, constraining outcome *sets* rather than *measures*), and the entanglement-swapping row as the cheapest potential kill shot. Both predictions were borne out in Section 7: the swapping stress test closed positively (Phase 6), and the $d \geq 3$ row closed constructively by exactly the set/measure split anticipated here (Phase 14).

6 The conspiracy objection, recast

The objection that superdeterministic correlations are "conspiratorial" is imprecise as stated: in a deterministic universe experimenters are physical subsystems, and *exact* statistical independence between settings and anything else would itself be unexplained. Bell conceded that a theory in which the correlations arise inevitably might make them digestible [5]; the modern literature has made the underlying complaint quantitative [34, 26, 27, 18, 20]. The results above sharpen it into three structural features any exact model must have, and two decidable questions.

The features: (i) *mechanism-independence* — by Theorem 4.7, no physical process may constitute a genuine intervention on the setting, so every possible setting mechanism (generator, brain, quasar photon [6, 14]) is enrolled in the same exact functional correlation; (ii) *exact cancellation* — the hidden setting-dependence is strong enough to violate Bell inequalities yet cancels exactly in every operational marginal [34]; (iii) *non-degradation* — the correlation survives arbitrary mixing between the common past and the experiment, with the exponential selectivity 4^{-N} of Theorem 4.5, while never appearing as bias in any classical statistical practice.

Question 6.1 (Symmetry protection)**.** Does there exist a microscopic symmetry whose Ward-type identity enforces exact no-signalling and operational measure-noncontextuality in a settings-dependent theory? On the quantum side the fine-tuning charge has been dissolved by generalizing Reichenbach's principle [1]; no analogous theorem exists on the superdeterministic side. If the symmetry exists, "fine-tuned" becomes "symmetry-protected" and the conspiracy objection dies; if it provably cannot, the exact program dies with it.

Question 6.2 (Law status of the boundary)**.** Theorem 2.3 shows the required support structure cannot be dynamical; its only non-conspiratorial home is law-like boundary status, where conditioning on a two-time boundary is no stranger than a variational endpoint or a microcanonical constraint [30, 33]. What must then be derived is implementation-independence: why every physical realization of the same setting induces the same boundary data. Theorems 4.3–4.4 are the first installments (proportional fiber profiles; forced exclusive gate; caliber-selected weights); the general statement is open.

On this recasting the original question "is it a conspiracy?" is replaced by Questions 6.1 and 6.2, both of which have mathematical content and either of which, resolved affirmatively, retires the objection.

7 The executed program

Between the fork and this revision, the standing agenda of Section 9 was executed as twenty-five machine-verified phases (notebook `models/NOTES.md`; every phase carries a Lean kernel job in the remote `compat` project, a passing `python3.12`/Z3 script, and an independent Mathematica check). The summary table; the paragraphs that follow are the per-phase records in chronological order, each closing an audit row of Section 5 or advancing a question of Sections 6–8.

Phase	Target	Result
1	Q-ward	parity Ward dichotomy: no-signalling $\Leftrightarrow$ parity fixed locus; Wood–Spekkens formalized and dissolved
2	bipartite row	partial entanglement forces two carriers; four-sector Ward decoupling; planar frustration
3	bipartite row	equivariant a.c. no-go with exact deficit $(s^2/c)\,\operatorname{artanh} c$
4	bipartite row	steering dilation: exact a.c. model for every pure two-qubit state; non-equivariance witnessed
5	bipartite row	equivariance collapse: equivariant models exist iff the state is a product
6	swapping row	chained steering through the common-future selector; $\sum_\mu T(\mu) = 0$: the conditioning constitutes the correlation
7	POVM row	smoothing theorem: strict unsharpness admits smooth full-support models at finite affinity; the cusp is the sharp limit
8	multipartite row	GHZ chain carries the Mermin paradox ($M = 4$ vs. 2); monogamy as architecture
9	Q-ward	record-affinity, the third protecting structure; the no-signalling fixed locus is a manifold
10	information premium	one-wing rigidity: the unique model carries a forced 0.2506-bit premium
11	Q-convergence	deviation–spread floor $(\sigma/2 - 1/6)/(\frac{8}{3}\sigma)$; resolution-independence
12	Q-convergence	six-cycle exhausted; floor doubly certified; arcs locally optimal
13	Q-convergence	twelve-cycle converges (90° dual weight vanishes); floor final for shift-orbit relaxations
14	$d \geq 3$ row	qudit chain dilation: Busch/CFMR requirements by certified algebra; KS set/measure split realized
15	hard conjugacy	record-affinity and rotation-orbit conjugacy derived from the canonical single-law microdynamics; open item narrowed to cross-skeleton intertwiners
16	hard conjugacy	the rank-chart intertwiner: unique canonical transport across skeletons (rigidity theorem); the item closes
17	swapping row	the Bell-basis chain: entangled measurement realized on carriers; universal station basins derive event-ready; no black box
18	bipartite row	the two-wing chain for every pure state of $\mathbb{C}^{d_A} \otimes \mathbb{C}^{d_B}$; unitarity protects; CGLMP 2.8729 vs kernel-checked LHV bound 2
19	multipartite row	the induction certified on W and cluster instances; the W pair CHSH-local unconditionally $(4\sqrt{2}/3)$, maximal conditionally $(2\sqrt{2})$
20	mixed states	purification marginals; the label is gauge (HJW); label-Bayes = Lüders; Werner $S = 2\sqrt{2}\,\eta$ with kernel-checked thresholds $1/3 < 1/\sqrt{2}$
21	POVM row	general POVMs by the Naimark chain; completeness = isometry; trine sequential statistics realized; SIC informational completeness as a kernel frame identity
22	DI row	the certification-defeat price: $S \leq 2+6D$ sharp, $D^*(S) = (S-2)/6$; Tsirelson costs $(\sqrt{2}-1)/3$, the maximum 1/3
23	normal form	chain forcing: exactness inverts the telescoping uniquely (triangular Jacobian = survivals); residual freedom is exactly the canonical gauge
24	computation row	Grover on carriers, exact rationals; caliber cost linear (2 bits per measurement); the exponential lives in boundary data — no simulation paradox
25	QFT row	the two-mode squeezed vacuum on the chain: geometric Schmidt tails, thermal full-rank reductions (RS shadow), pseudospin CHSH $2\sqrt{41}/5$ at $t = 1/4$, Tsirelson in the limit

Progress on Question 6.1 (compat notebook, Phase 1). Within the exclusive local-gate class of Theorem 4.4 the protecting symmetry exists and is source parity. A parity-even boundary profile has identically zero signed marginals (Ward identity; the involution $n \mapsto -n$ flips both records and fixes every profile built from the projective setting data), while a linear odd term $\varepsilon\,(b \cdot n)$ produces the explicitly signalling marginal $\varepsilon\,(a\cdot b)/(2Z)$ — signalling is the order parameter of the broken involution. Conversely, unbiased marginals at all settings force the odd sector to vanish (Funk–Hecke spanning; exact low-degree checks), so in any finite parameterization the no-signalling set is Lebesgue-measure zero *and* equals the parity fixed locus: the Wood–Spekkens tuning claim is correct, and it is a symmetry criterion rather than a coincidence. The general criterion is equivariance — the normalized odd sector of the profile must be independent of the remote setting — which for biased (partially entangled) states becomes the structural ansatz $f_{ab} = \mathrm{even}(a, b, z) + \mathrm{odd}(z)$ and identifies Q-ward with the implementation-conjugacy condition of the $d \geq 3$ program restricted to remote contexts. Artifacts: `models/QWardParityDichotomy.lean`, `q_ward_parity_dichotomy.py`, with independent Mathematica confirmation.

Progress on the bipartite row and Question 6.1 (compat notebook, Phase 2). The equivariance ansatz was tested against $\cos\xi\,|00\rangle + \sin\xi\,|11\rangle$ (correlation matrix $\mathrm{diag}(s, -s, 1)$, marginals $c\,(a\cdot z)$ with $c = \cos 2\xi$, $s = \sin 2\xi$; verified directly and in Mathematica). Four results. (i) *Shared-carrier rigidity*: any locked direction — a matched setting where the two responses agree pointwise on the carrier, which occurs for every fixed intertwiner with a transverse eigenvector, including the reflection natural to the state's correlation matrix — forces matched-setting correlation one, while the quantum target $s + (1-s)(a\cdot z)^2$ is strictly below one off the Schmidt axis. The single-carrier architecture of Theorem 4.3 is therefore *exactly* the maximal-entanglement point of the family; partial entanglement forces a two-carrier ontology or apparatus hidden variables — the two-wing analogue of the hidden-measurement forcing in Theorem 3.3. (ii) *The Ward mechanism generalizes*: under the $\mathbb{Z}_2 \times \mathbb{Z}_2$ parity group on the two-carrier space, mass, the two marginals, and the correlation are carried by separate sectors; no-signalling is again a sector condition, and the quantum moments force the low-degree components uniquely ($2c(z\cdot n_A)$, $2c(z\cdot n_B)$, $4\,n_A\!\cdot\! D\, n_B$) — realizing $f_{ab} = \mathrm{even}(a, b, z) + \mathrm{odd}(z)$ exactly. (iii) *A perfect-correlation support law*: unit correlation in a direction forces carrier sign agreement almost surely; maximal states are perfectly correlated in every direction, forcing a deterministic carrier relation with product-null graph (an absolutely continuous two-carrier profile is impossible at maximal entanglement), while partial states need agreement only along the Schmidt axis — carrier-measure singularity is graded by entanglement. (iv) *A planar frustration*: exact rational targets at the five Niven settings are Z3-infeasible for circle carriers at $K = 12, 24, 36$, for maximal and partial states alike, and collectively — every tested proper subset of the constraints is feasible, and removing the single $(60^\circ, 120^\circ)$ pair restores feasibility — the two-wing analogue of the trine frustration, indicating the spherical fiber geometry is essential even at finite resolution. Artifacts: `models/QWardEquivariantBipartite.lean`, `q_ward_equivariant_bipartite.py`, with independent Mathematica confirmation.

Progress on the bipartite row: the a.c. carrier no-go and the exact entanglement deficit (compat notebook, Phase 3). Executing Phase 2's construction target closed it negatively, with a sharper theorem than the grading language of (iii) above suggested. On the two-carrier space, nonnegativity on each parity orbit is *exactly* the condition $f_{ee} \geq \Phi := \max(|x{+}y| - w,\ |x{-}y| + w)$, and pointwise — independently of the odd sectors — $\Phi \geq |w|$ with $w = 4\,n_A\!\cdot\! D\, n_B$ forced. Since odd sectors carry no mass, every equivariant profile obeys $1 = \int f_{ee} \geq 4\,\mathbb{E}\,|n_A\!\cdot\! D\, n_B| = \mu(s) = 1 + \frac{s^2}{c}\,\mathrm{artanh}\, c$, which exceeds one for every $s > 0$: *no absolutely continuous equivariant two-carrier profile exists at any nonzero entanglement.* Carrier singularity is not graded to zero at small entanglement — it switches on immediately; what is graded is the *deficit* $\mu(s) - 1$, the minimal mass that must ride on a singular carrier component, interpolating from 0 at the product state (where the product of one-qubit models is the extremal solution, mass exactly one) to 1 at maximal entanglement (fully deterministic carrier relation). Entanglement equals singular carrier mass, quantitatively. The conjectured attainability of the deficit by a graph-supported singular component plus an a.c. background is the row's next target.

Artifacts: `models/QWardAcCarrier.lean`, `q_ward_ac_carrier.py`, with Mathematica certifying the closed form and the deficit's positivity exactly.

Progress on the bipartite row: the steering dilation closes the constructive side (compat notebook, Phase 4). The complement to Phase 3's no-go: an exact, *absolutely continuous*, setting-conditioned model exists for every pure two-qubit state. Alice carries the branch mixture of dipole basins (weights $(1 \pm c)/2$, parent Phase 8); Bob's carrier is drawn from the dipole basin of the steered vector $m_\alpha = (cz + \alpha Ta)/(1 + \alpha c\, a_z)$, which is exactly unit ($|cz + \alpha Ta|^2 = (1 + \alpha c\, a_z)^2$, kernel-checked), so the joint law telescopes to the quantum form identically and the dilation is a working sampler (Monte Carlo verified, including maximal entanglement, where matched settings give perfect correlation from an a.c. measure — the setting-conditioned family evades the support law measure-by-measure). Its no-signalling is a *conditional* telescoping cancellation, not parity protection, and it is provably non-equivariant: the odd-odd pairing $E[\mathrm{sign}(z \cdot n_A)\, \mathrm{sign}(x \cdot n_B)]$ takes different values at Alice settings of equal polar angle (0.159 vs. 0.000), exactly the setting-injection Phase 3 requires of any a.c. exact model. Post-first-outcome states are pure products, so the parent sequential argument attaches verbatim. The resulting classification: product — equivariant a.c.; maximal — equivariant singular or non-equivariant a.c.; interior — non-equivariant a.c. only, with the equivariant-singular case open but rigid (forced odd-odd total variation one, even-even = |odd-odd|, orbit saturation annihilating mixed-parity points — kernel-checked). Trade-off, informally: *exactness forces singular carriers or setting-injected odd sectors.* For Question 6.1 this bifurcates the protection: parity covers the extreme points of the entanglement interval; the interior admits only conditional cancellation in the a.c. class, and whether that cancellation is itself the fixed locus of a symmetry of the boundary conditioning is now the sharpest open form of the question. Artifacts: `models/SteeringDilation.lean`, `steering_dilation.py`, with Mathematica residuals vanishing identically.

Progress on the bipartite row: the equivariance collapse completes the classification (compat notebook, Phase 5). The equivariant-singular case closes negatively, and the mechanism is rigidity rather than search. The correlation pairing reads only the odd-odd sector, whose total variation is at least one by the Schmidt-axis dual ($E(z, z) = 1$, $|g_z| \le 1$); orbit nonnegativity dominates it by the even-even sector, whose mass is exactly one for every setting pair; and a dominated family with equal mass *equals* its dominator (kernel-checked). So equivariance of the odd sectors alone forces the entire measure to be setting-independent — a local hidden-variable model, bounded by $|S| \le 2$ (pointwise dichotomy and averaged bound kernel-checked), against the quantum value $S = 62/25 > 2$ at rational settings for $(c, s) = (3/5, 4/5)$ and $2\sqrt{1 + s^2} > 2$ in general (Gisin). **Equivariant models exist iff the state is a product.** A corrective by-product: mapped to the two-carrier space, even the parent singlet model's odd-odd sector is setting-dependent ($G(z) = 0$ vs. $G(\mathrm{skew}) = 0.200$), so no entangled state has an equivariant model of any kind. The row's structural story is now complete: *entanglement is exactly the property that forces the boundary law to write the settings into the parity-odd sectors of the carrier measure* — the Phase-3 deficit prices the absolutely continuous refusal, this phase outlaws refusal altogether, and the steering dilation shows compliance is cheap. For Question 6.1, parity protection is confined to unentangled states (and the trivially-zero marginal sectors of maximal ones); the single remaining form of the question is whether the conditional telescoping arrangement is itself symmetry-protected. Artifacts: `models/EquivarianceCollapse.lean`, `equivariance_collapse.py`, with Mathematica confirming the reclassification integrals.

Progress on entanglement swapping: the stress test closes positively (compat notebook, Phase 6). The architecture's sharpest feared failure — that the common-past selector cannot be replaced when the entangled pair shares no source — is refuted constructively. The protocol's statistics factor through one-qubit steering alone: Alice's outcome steers B_1 to $m_1 = \alpha\, T(\Phi^+)a$; the Bell-state measurement outcome is a *uniform common-future selector* ($P(\mu) = \frac{1}{4}$ independent of all settings and outcomes — the event-ready property, kernel-checked); and the teleportation composition steers C to $\alpha\, T(\mu)a$, with correction table $R_\mu = T(\mu)T(\Phi^+) = \{I, Z, X, Y\}$ exactly the Pauli Bloch maps. The resulting dilation — Alice's dipole branches, two fresh tape bits at the station, Charlie's dipole basin of the composed steering vector — reproduces the swapped joint law $(1/16)(1 + \alpha\gamma\, a{\cdot}T(\mu)c)$, verified by direct four-qubit computation, by Monte Carlo sampling of the dilation, and symbolically in Mathematica; conditional on $\mu = \Psi^-$ the never-interacting pair violates CHSH at rational settings ($S = 14/5 > 2$, Z3 LHV-infeasible), with n_A and n_C correlated only through the boundary data. The necessity statement is the phase's sharpest point: $\sum_\mu T(\mu) = 0$, so marginalizing the station record leaves the swapped pair exactly uncorrelated — the common-future conditioning does not *replace* a common source, it *constitutes* the correlation. Scope: the station is modeled as the Bell-basis device only; a native carrier-level theory of general entangled measurements is not claimed. Artifacts: `models/SwappingDilation.lean`, `swapping_dilation.py`.

Progress on the POVM row: the smoothing theorem (compat notebook, Phase 7). The conjecture of Section 9 item 4 is proven constructively for the qubit-pair sector. Unsharp effects are exactly the classical mixture $E_\pm = \eta\Pi_\pm + (1-\eta)I/2$ (kernel-checked), and the steered carrier density for a mixed Bloch vector of length η has its odd part forced to $2\eta t$ with two-sided positivity forcing the even floor $e \geq 2\eta|t|$, of mass exactly η. Below sharpness one the slack $1-\eta$ admits smooth, strictly positive, full-support densities with *exact* statistics (explicit construction, verified as a working sampler), and the conditioning affinity is finite: $\kappa(\eta) = \log \frac{1+3\eta}{1-\eta}$ — the parent finite-affinity no-go (Theorem 2.2) dissolves off the sharp limit. At $\eta = 1$ the floor's mass saturates the budget and the same saturation lemma that drives the Phase-3 deficit and the Phase-5 collapse forces the density uniquely to the dipole basin — support zero on half the carrier space, cusp included: *the cusp is the signature of ideal measurements.* Since real detectors are strictly unsharp, the two-boundary law required to describe actual experiments is finite-affinity and smooth; the hard boundary belongs to the idealization. Scope: the density class only (records still read out by sign responses, governed by the finite-gain floors of the deviation dictionary); non-covariant effects and $d \geq 3$ remain open. Artifacts: `models/PovmSmoothing.lean`, `povm_smoothing.py`, with Mathematica confirming the saturation mass, Z_δ, and both identities.

Progress on the multipartite row: the GHZ chain (compat notebook, Phase 8). The all-or-nothing case closes constructively. The GHZ joint law is $(1/8)[1 + \alpha\beta a_z b_z + \alpha\gamma a_z c_z + \beta\gamma b_z c_z + \alpha\beta\gamma\, \mathrm{Re}(\hat{a}\hat{b}\hat{c})]$ with $\hat{a} = a_x + i a_y$, and the sequential steering chain reproduces it exactly: Alice's fair record conditions BC into the phase-generalized two-qubit family with Schmidt bias αa_z and coherence $W = \alpha\hat{a}$; the steered-norm identity of the Phase-4 dilation is *phase-independent*, so it certifies the new family unchanged; the chain telescopes, the transverse composition yields $\mathrm{Re}(W\hat{b}\hat{c})$, and Alice's record squares it into the three-point term (all kernel-checked). The dilation is a working sampler and carries the Mermin paradox exactly — $E(XXX) = 1$, $E(XYY) = E(YXY) = E(YYX) = -1$, $M = 4$ against the kernel-checked local bound 2 (pointwise dichotomy over all 64 deterministic strategies; Z3 LHV-infeasible with control). Monogamy appears as architecture: every two-party marginal is the classical $\mathrm{diag}(0, 0, 1)$, carried by the shared branch, while the tripartite coherence rides exclusively on the record-conditioned chain. With Phases 4 and 6, steering chains now cover arbitrary pure two-qubit states, entanglement swapping, and GHZ — the chain is emerging as the constructive normal form of the SD-B sector on finite qubit networks, with general networks following by induction (each outcome conditions the remainder into a pure state of fewer qubits) and only the uniqueness/forcing side left open. Artifacts: `models/GhzChainDilation.lean`, `ghz_chain_dilation.py`, with Mathematica confirming the symbolic three-point identity, the Mermin value, and the pairwise correlations.

Progress on Question 6.1: record-affinity, the third protecting structure (compat notebook, Phase 9). The refined question left by Phase 5 — what protects the telescoping cancellation carrying no-signalling in all three dilations — is answered: the subnormalized conditioning is *affine in the record*, $P(\alpha|a)\, m_\alpha = (cz + \alpha T a)/2$ exactly (the steered denominator equals the record probability), so the record average telescopes with uniform weights and the setting-dependent, record-odd steering part cancels identically regardless of outcome bias. Quantum grounding: this is Bayes consistency plus completeness ($\sum_\alpha \mathrm{Tr}_A[(\Pi_\alpha \otimes I)\rho] = \rho_B$), transported to the carrier level. Its symmetry form is the record-flip involution $(\alpha, c) \to (-\alpha, -c)$: exact at unbiased marginals (recovering the Phase-1 parity Ward identity at the dilation level), state-covariant in general — the breaking is carried by the state parameter, never the setting. The perturbation dichotomy is kernel-checked and demonstrated in a working sampler: record-odd deformations of the steering — however setting-dependent — do not signal (the no-signalling fixed locus is a *manifold* of measurement-dependent models, the quantum point one point of it), while record-even setting-dependence signals linearly, with the shift formula $(ca_z)\delta(b{\cdot}D) + \varepsilon(b{\cdot}E)$ exact. The Q-ward ledger for the verified sector is now complete: source parity (marginals), branch normalization (censoring), record-affinity (steering) — every exact no-signalling identity in the program is a fixed locus, none a tuning. What remains is the S1 hinge, sharpened: derive record-affinity from a microdynamics, and extend to $d \geq 3$ where implementation conjugacy is its general form. The safe-manifold finding also re-poses the information premium: what distinguishes the quantum point *within* the fixed locus is the next quantitative target. Artifacts: `models/QWardRecordFlip.lean`, `q_ward_record_flip.py`, with Mathematica (closed out after a server outage) independently confirming the symbolic record average cz, the linear-response residual, and the partial-trace completeness residual, all exactly zero.

Progress on the information premium: one-wing rigidity (compat notebook, Phase 10). The question of Theorem 4.6(iii) closes with a stronger result than conjectured. In the one-wing class — Bob's carrier density conditioned on Alice's data only, Alice's record any stochastic function of the carriers — exactness forces the entire model uniquely: the correlated density's odd part is forced, the antipodal floor and mass saturation pin the carrier density to $2|a \cdot n|$ (the saturation mechanism's fifth appearance), and a two-sided squeeze pins the record response deterministic. The parent one-sided Liouville model is the *only* exact one-wing model, so its setting-carrier information $\ln 2 - \frac{1}{2}$ nats = 0.2787 bits is forced, not chosen; the premium over Hall's two-setting minimum (0.0280 bits, reproduced to ten digits) is $\Delta = 0.2506$ bits. The kernel-checked KL chain rule locates it exactly: at every relative angle the total splits into sector information plus within-sector information, the Haar average of the sector term *is* Hall's cost, and the premium is the average within-sector structure that a one-wing density cannot discard because it must serve every remote setting at once — *the price of remote-setting universality.* Combined with Phase 9: within the no-signalling safe manifold, the quantum point of the one-wing class is not merely distinguished, it is alone. The result also quantifies a boundary trade: terminal conditioning that references both wings' settings jointly can be ten times cheaper in setting-information than conditioning threaded through one wing's records — a fixed, larger information signature for locally-steered models, added to the deviation dictionary. Artifacts: `models/InformationPremium.lean`, `information_premium.py`, with Mathematica confirming the forced value $\log 2 - \frac{1}{2}$, Hall's integral to twenty digits, and the chain-rule factorization symbolically.

Progress on Question 8.1: the deviation–spread law (compat notebook, Phase 11). The convergence question is answered for the counting class, in an unexpected direction: *there is no $\Theta(f(K))$ decay — the trade-off variable is the spread fraction $\sigma = t/K$, not the resolution K.* The parent's six-cycle orbit bound, which proves the one-third law under exact repeatability, fails without it (the alternating occupation cheats it to $D = 3$); exhaustive enumeration over all 512 occupation/window vertices — kernel-checked and independently confirmed in Mathematica — yields the tight patched bound $D \leq 1 + \frac{2}{3}R$, pricing the cheat exactly. Summed over the 60° orbits this gives a *universal deviation floor*: every counting model at spread fraction $\sigma > \frac{1}{3}$, at every resolution, with stochastic responses allowed, deviates from the Born curve by at least $(\sigma/2 - 1/6)/(\frac{8}{3}\sigma)$ — zero exactly at the one-third point, 3/80 at $\sigma = \frac{5}{12}$, 1/16 at $\sigma = \frac{1}{2}$. Exact LP upper bounds bracket the true curve and are K-independent to $\sim 10^{-3}$ between $K = 24$ and $K = 48$ (0.054 and 0.104 at the two tested spreads, against the linear-profile anchor 0.10526). The conclusion sharpens this document's Planck-scale discussion: for ontically finite programs of this class, *refining the ontic grid is not a route to the Born statistics*; the deviation is controlled by the naturalness parameter alone, and vanishes only as epistemicity is surrendered toward the one-third point. Open for part 2: raising the floor with the trine-orbit (120°) machinery and lowering the ceiling beyond arc supports. Artifacts: `models/DeviationSpreadLaw.lean`, `deviation_spread_law.py`, with Mathematica confirming the enumeration constants, the anchor, and the floor values.

Progress on Question 8.1, part 2: the six-cycle is exhausted (compat notebook, Phase 12). Both refinement directions of the Phase-11 bracket return informative negatives plus a new certificate. The *exact* six-cycle relaxation — the 512-atom LP over the full configuration polytope, constraining the occupation, repeatability, 60°, and 120° moment groups jointly — reproduces the Phase-11 floor identically at every tested spread (0, 1/48, 3/80, 1/16 at $\sigma = \frac{1}{3}, \frac{3}{8}, \frac{5}{12}, \frac{1}{2}$): the patched inequality was already six-cycle-optimal, and the remaining floor–ceiling gap (0.0625 vs. 0.1036 at $\sigma = \frac{1}{2}$) is provably beyond-six-cycle territory. Dual extraction produced a second, independent certificate at $\sigma = \frac{1}{2}$ in which the 120° moments participate — $3S_1 - 3S_2 + 4S_0 \leq 18$ per orbit, kernel-checked over all 512 vertices, closing arithmetic $21 - 48\varepsilon \leq 18$ giving $\varepsilon \geq 1/16$ exactly (Mathematica concurs on both) — so the floor is doubly certified by distinct extremal mechanisms. Hill-climbing over supports at $K = 24$ found arcs locally optimal: the ceiling stands. The structural conclusions of the deviation–spread law — a positive, resolution-independent floor for all $\sigma > \frac{1}{3}$ — are now as robust as orbit-local tools can make them; closing the residual bracket would need the 2^{18}-atom twelve-cycle relaxation (coupling in the 90° moments) and is recorded as a secondary refinement. Artifacts: `models/DeviationSpreadLaw2.lean`, `deviation_spread_law2.py`.

Question 8.1, part 3: the orbit program converges (compat notebook, Phase 13). The twelve-cycle relaxation — 262,144 configurations deduplicating to 289 distinct moment vectors, coupling the 90° moments — does *not* raise the floor, and the reason is structural: the optimal integer duals have zero weight on the 90° moments (the parent's "90° comes free from antipodality" rediscovered as a vanishing dual variable), and by antipodality their coefficient forms collapse onto the two interleaved six-cycles — the twelve-cycle dual *is* a pair of six-cycle duals. Both extracted duals, $(4, 3, 0, -3; 36)$ at $\sigma = \frac{1}{2}$ and $(4, 1, 0, -5; 12)$ at $\sigma = \frac{5}{12}$, are kernel-checked in the decoupled form (per-position separation $qx \leq \max(0, x)$ plus eight-vertex max-sum bounds) with closing arithmetic re-deriving 1/16 and 3/80 exactly (Mathematica concurs on all four values). The deviation–spread floor is now *triply certified and final for the entire class of shift-orbit relaxations over rational Born targets*: 30° and 150° have irrational targets, 90° is free, and higher cycles decompose. The residual floor–ceiling gap is attributable to the all-angle constraints only global or irrational-target methods can see; Question 8.1 is settled to the reach of this program's tools. Artifacts: `models/DeviationSpreadLaw3.lean`, `deviation_spread_law3.py`.

Progress on the $d \geq 3$ row: the qudit chain dilation (compat notebook, Phase 14). The last heavyweight row opens constructively. Every binary projective split of a pure state is an *effective* *qubit* — with $p = \|\Pi\psi\|^2$ the state lives in the 2D span of $\Pi\psi$ and $(I - \Pi)\psi$ with Bloch vector $(\sin 2\theta, 0, \cos 2\theta)$, $\cos^2\theta = p$ — so the parent dipole-basin machinery applies verbatim and a d-outcome measurement becomes a *chain* of basin measurements with Lüders updates. The chain probabilities telescope to $\|\Pi_i\psi\|^2$ independently of the ordering and the completing basis (kernel-checked; all six $d = 3$ ordering values return exactly $\{p_1, p_2, p_3\}$ in Mathematica), and additivity over the context is a certified identity: requirements (a) and (b) of the Busch/CFMR frame hold *by algebra*. The Kochen–Specker split is realized exactly as Section 5 predicted: the carrier basins for a shared effect differ across contexts and chain positions (contextual outcome sets) while their measures agree (shared-effect check: 0.6512–0.6520 against 0.6519 across two completing bases and three chain positions). The dilation runs as a working sampler in $\mathbb{C}^3$ and $\mathbb{C}^4$ with exact two-step Lüders sequences. The verified sector now spans all pure two-qubit states, swapping, GHZ, unsharp qubit pairs, and all finite-dimensional pure projective measurements with sequences — the chain is the constructive normal form of the SD-B sector, four times over. What remains of the S1 program, and of the compatibility question's heavyweight list, is exactly one item: *hard* implementation conjugacy — a canonical physical automorphism intertwining implementations, derived from a microdynamics rather than granted by the abstract isomorphism theorem. Artifacts: `models/QuditChainDilation.lean`, `qudit_chain_dilation.py`.

Progress on the hard-conjugacy item: the canonical microdynamics (compat notebook, Phase 15). The program's single remaining open item splits, and its larger half closes. The microdynamics was already in the archive: the parent Phase-8 hybrid, taken seriously as *one* law — state space $X \times \{0, 1\}$ with the apparatus axis an *ontic* coordinate of X, measure Liouville$\times$Haar$\times$uniform-bit, and one measurement map, the record toggle on the basin indicator $M(n, a, r) = (n, a, r \oplus [a \cdot n \geq 0])$. The map is an involution and an automorphism of the law (kernel-checked as an `Equiv` with pushforward invariance; it fixes every fiber-uniform measure *pointwise* — the proof is `rfl`, since the map does not read the record it toggles). Two theorems follow. *Record-affinity is derived*: for one measure conditioned on the record partition, the subnormalized conditionals sum to the marginal — partition additivity, the law of total probability (kernel-checked, Z3 UNSAT) — so Phase 9's protecting structure is automatic for every single-law model, and the verified SD-B sector *is* single-law (the parent's max-caliber phase exhibits the boundary conditioning as conditioning one reference measure on the compatibility event). *Implementation conjugacy is derived on rotation orbits*: the lifted rotation $\Phi_R(n, a, r) = (Rn, Ra, r)$ preserves the law (cotangent lifts are symplectic — parent Phase 8; Haar is rotation-invariant) and intertwines the one measurement map with itself, $M \circ \Phi_R = \Phi_R \circ M$, because the basin indicator is equivariant: $(Ra) \cdot (Rn) = a \cdot n$ (kernel-checked for the generating rotation; symbolically for the general axis-angle rotation by exact rational parameterization; exact on 3000 random rotations). Outcome masses are conjugacy invariants (kernel-checked), so implementations of rotated effects have equal statistics under transport — measure noncontextuality by the S1 mechanism, now with *canonical physical* automorphisms rather than abstract isomorphisms; transitivity of $SO(3)$ puts all qubit rank-one effect implementations in a single conjugacy orbit (per effective qubit, every link of the $d \geq 3$ chains), with stabilizer rotations covering same-effect different-history implementations (both verified as samplers). The honest residue, sharper than before: canonical intertwiners *across chain skeletons* — two orderings or completing bases implementing the same $d \geq 3$ effect have equal masses (Phase 14) and abstract equal-mass isomorphisms, but a microdynamically canonical intertwiner between skeletons is not constructed; the identified route is the parent's Rosenblatt rank charts. Artifacts: `models/CanonicalMicrodynamics.lean`, `canonical_microdynamics.py`, and `canonical_microdynamics.wl` (the general rotation identity closed symbolically).

Progress on the hard-conjugacy item, concluded: the rank-chart intertwiner (compat notebook, Phase 16). The item Phase 15 left open closes by the route it named. The structural fact doing the work: in the record-conditioned steering-chain presentation — the sector's own normal form — conditioning on the record path makes the traversed steps *independent* dipole lobes, so the cell-conditional carrier law is an exact product and the per-step ranks $(t_j^2, \varphi_j/2\pi)$ are i.i.d. uniform (the probability integral transform; the chart density identity $d(t^2) = 2t\,dt$ and its strict monotonicity are kernel-checked). The *rank chart* of a cell is the nested composition of these ranks in chain order — every ingredient read off the skeleton's own dynamics: the cell decomposition is the record, the coordinate order is the sampler's draw order, the digit priority is the chain's time order (the filtration). The cross-skeleton intertwiner is $\Phi = \mathrm{chart}_B^{-1} \circ \mathrm{chart}_A$ cell by cell. It is *well-defined* because cell masses agree across skeletons — Phase 14's telescoping reused as boundary alignment (kernel-checked; symbolic for all six $d = 3$ orderings; exact rational for all orderings at $d = 3, 4$); *measure-preserving* by partition additivity and the chain-order lane merge (a bijection of dyadic grids, kernel-checked with exhaustive and 20000-round-trip confirmation); and *canonical* by rigidity rather than preference: given the filtration order, a monotone self-correspondence of the rank axis is the identity (kernel-checked; Z3-enumerated), so exactly one chart exists and no abstract-isomorphism choice survives — with the chart data built from pairing invariants, the assignment skeleton $\mapsto$ chart is equivariant under the Phase-15 canonical automorphism group. Verified running: same-context reorderings at $d = 3$ transport cell-to-cell exactly (the Φ-image of one ordering's native ensemble reproduces the other's law: Born frequencies, per-step moments $E[t] = 2/3$, $E[t^2] = 1/2$, uniform azimuth ranks, vanishing cross-step covariance); across contexts the shared effect's cell transports exactly (the KS-relevant statement — measure noncontextuality upgraded to *canonical-isomorphism* noncontextuality) while the complement moves as one lump with the sub-record as a discrete chart coordinate — records of distinct effects are transported by measure, never pretended to correspond, which is exactly the strength Kochen–Specker permits. With Phases 15 and 16 together, the record-affinity/implementation-conjugacy structure of the SD-B sector is derived end to end: one invariant law, its canonical automorphisms, and a unique canonical transport between any two implementations of the same effect. Artifacts: `models/RankChartIntertwiner.lean`, `rank_chart_intertwiner.py`, `rank_chart_intertwiner.wl`.

Progress on the swapping row: the Bell-basis chain removes the black box (compat notebook, Phase 17). The Phase-6 scope note — the station modeled as the Bell-basis device only — dissolves. A Bell-state measurement is a projective measurement on the co-located pair, and the Phase-14 chain never cared whether the basis vectors are entangled: with station input $|m_1\rangle \otimes |\Phi^+\rangle$ (pure in $\mathbb{C}^8$) the BSM is the degenerate projective family $\{\Pi_\mu \otimes I\}$, realized as three binary splits with dipole-basin carriers. The new structural fact: every Bell state's one-qubit marginal is $I/2$, so each cell mass is $\frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$ for *every* input (kernel-checked; symbolic in Mathematica), hence the chain's conditional profile is *universally* (1/4, 1/3, 1/2) with effective Bloch data (−1/2, −1/3, 0) — the event-ready property is derived at the carrier level and strengthened: the station's basins require no setting, outcome, or state information whatsoever, so station-side no-signalling is manifest rather than arranged, and the common-future selector is one fixed carrier object for all protocols of this shape — the sharpest form yet of the anti-conspiracy story for swapping. The steered amplitudes realize $\{I, Z, X, iY\}$ exactly (Mathematica residuals zero), i.e. the correction table $R_\mu = T(\mu)T(\Phi^+) = \{I, Z, X, Y\}$, with $T(\Phi^+)^2 = I$ and $\sum_\mu T(\mu) = 0$ re-derived through the chain (all kernel-checked). End to end — Alice's fair branch, the Bell chain on literal dipole carriers, Charlie's steered basin — the dilation reproduces the swapped joint law $(1/16)(1 + \alpha\gamma\, a\cdot T(\mu)c)$ in all sixteen cells with *no black box*, and conditional on Ψ^- the never-interacting pair gives CHSH $S = 2.798$ (target 14/5) against the exhaustive LHV bound 2. A scope observation closes the general question: any entangled basis on co-located systems *is* a projective measurement on the joint space, hence a chain; entangled measurements across spacelike-separated wings are not local operations in quantum mechanics either, so co-location restricts the physics, not the model. The chain normal form now covers entangled bases — its fifth deployment. Artifacts: `models/BellBasisChain.lean`, `bell_basis_chain.py`, `bell_basis_chain.wl`.

Progress on the bipartite row: the two-wing steering chain in every dimension (compat notebook, Phase 18). The row's constructive content extends from Schmidt rank 2 to arbitrary pure states of $\mathbb{C}^{d_A} \otimes \mathbb{C}^{d_B}$. Alice's measurement on the joint pure state is the degenerate projective family $\{P_i \otimes I\}$: her record-conditioned chain runs on her wing with split masses $p_i = \sum_k \lambda_k |\langle e_i | u_k \rangle|^2$ — functions of her basis and the Schmidt data only, so her basins are local and Alice-side no-signalling is structural — and her record steers Bob to the pure state $|\varphi_i\rangle = (\langle e_i| \otimes I)|\psi\rangle / \sqrt{p_i}$, where Bob's chain attaches verbatim. The joint law telescopes to the Born cell exactly (kernel-checked; exact for $\mathbb{C}^3 \otimes \mathbb{C}^3$ and the asymmetric $\mathbb{C}^2 \otimes \mathbb{C}^4$; the full sampler matches all nine cells of a random Schmidt-rank-3 state). The protecting structure simplifies as it generalizes: Bob's marginal after steering is $\sum_i p_i |\varphi_i\rangle\langle\varphi_i| = \rho_B$ — record-affinity's d-dimensional form — and the identity holds because Alice's overlap matrix has orthonormal columns: *unitarity itself protects* (kernel-checked at ranks 2 and 3 by `linear_combination` in the Gram constraints; polynomial residuals zero in Mathematica; Z3 UNSAT). Completeness of the local context — the Busch/CFMR requirement (b) — is the no-signalling protection, so conspiracy would require *breaking* unitarity, not adding structure. The Schmidt-rank-3 witness: the chain dilation of the maximally entangled qutrit pair with the standard settings reproduces CGLMP $I_3 = 2.8717$ (sampled) against the exact quantum value $4(3+2\sqrt{3})/9 = 2.87293\ldots$ (Mathematica closed form) and the local bound 2 — exhaustive over all 81 deterministic strategies and *kernel-checked by* `decide`. The chain is now one architecture from qubit pairs to arbitrary dimensions — its sixth deployment; whether the rank-2 structure theory (deficit, collapse) persists at rank ≥ 3 is recorded as open structure theory, not a compatibility gap. Artifacts: `models/BipartiteQuditChain.lean`, `bipartite_qudit_chain.py`, `bipartite_qudit_chain.wl`.

Progress on the multipartite row: the induction certified on instances (compat notebook, Phase 19). The general-network claim rode on a sketch — each outcome conditions the remainder into a pure state of fewer qubits — executed only at GHZ; it is now a verified mechanism. The induction step: a first local measurement is a binary split of the joint pure state (an effective qubit, carrier local to its wing), each record steers the remainder to a *pure* state, and at two qubits the recursion lands exactly on the Phase-4 base case; the telescoping is kernel-checked at depths two and three, wing-by-wing no-signalling is Phase 18's unitarity protection, and the recursive law equals the joint Born law exactly for random three- and four-qubit states ($< 10^{-12}$). Two instances of opposite monogamy architecture run as samplers: the 4-qubit linear cluster (all sixteen cells, depth-two recursion) and the W state (all eight cells), whose steering decomposition is exact — record 0 (mass 2/3) steers the pair to Ψ^+, record 1 (mass 1/3) to a product. The W monogamy story is the anti-GHZ: the pair's reduced correlation matrix is the record mixture $\mathrm{diag}(2/3, 2/3, -1/3)$ (exact partial trace), so Horodecki gives unconditional CHSH at most $4\sqrt{2}/3 = 1.8856 < 2$ — entangled (concurrence 2/3) yet CHSH-*local* — while conditioning on the third record upgrades the pair to the *maximal* $2\sqrt{2}$. One sampler at one set of settings realizes both numbers (1.888 and 2.829; the arithmetic $32/9 < 4 < 8$ and $4\sqrt{2}/3 < 2 < 2\sqrt{2}$ kernel-checked). GHZ hid the tripartite coherence in classical pairwise marginals; W hides it in entangled-but-local pairs; in both, the record conditioning is where the nonclassical correlation lives. With this, the entanglement completion program (Phases 17–19) is done: entangled measurements, arbitrary bipartite dimensions, multipartite induction. Artifacts: `models/MultipartiteChainInduction.lean`, `multipartite_chain_induction.py`, `multipartite_chain_induction.wl`.

Progress on mixed states: the label is gauge (compat notebook, Phase 20). The remaining state-space caveat — mixed states handled by convex branch labels, noted at Phase 14 but not separately certified — closes with three reductions. *Purification*: a mixed state is a wing of a pure entangled state, i.e. already the Phase-18 sector; the convex label *is* the ancilla's record, the choice of decomposition *is* an ancilla measurement basis (Hughston–Jozsa–Wootters: all decompositions arise as ancilla bases on one purification), and decomposition-independence of the physics *is* the unitarity-protected no-signalling already certified — verified as marginal invariance over random ancilla bases ($< 10^{-12}$), with the kernel carrying the bilinear Gram invariance (orthonormal-column reshuffles preserve every second moment; ensembles of different *cardinality* included via isometries). The hidden-common-cause worry dissolves: the label is gauge, and canonically so, since ancilla-basis changes are wing-local unitaries conjugate under the Phase-15/16 canonical machinery. *Sequential exactness*: label-Bayes equals Lüders — the posterior label mixture of conditional states telescopes to $\Pi\rho\Pi$ (kernel-checked; verified exactly on qutrit mixtures through a degenerate rank-2 projector, where post-states genuinely differ per label; exact rationals in Mathematica). *The Werner witness*: $W(\eta)$ is the Bell-label mixture with weights $(1+3\eta)/4$ on Ψ^- and $(1-\eta)/4$ on each other Bell state, each label carrying its Phase-4/17 dilation, and the weighted transfer matrices collapse — $\sum_k w_k\, T(\mu_k) = -\eta I$ (kernel-checked) — so $E(a,b) = -\eta\, a\cdot b$ and CHSH $S = 2\sqrt{2}\,\eta$ (symbolic in Mathematica, with the partial-transpose spectrum $\{(1-3\eta)/4,\ (1+\eta)/4^{\times 3}\}$ exact). The sampler traces the law across the threshold ($S = 2.399$ at $\eta = 0.85$, violating; $S = 1.692$ at $\eta = 0.6$, local yet entangled), and the thresholds are kernel arithmetic: violation iff $\eta > 1/\sqrt{2}$, entanglement iff $\eta > 1/3$, and $1/3 < 1/\sqrt{2}$ — the band of entangled-but-CHSH-local mixed states is nonempty and the same label model covers it constructively, Werner's own 1989 observation realized inside the SD-B normal form. Mixed states are not an extension of the verified sector; they are its marginals. Artifacts: `models/MixedStateLabels.lean`, `mixed_state_labels.py`, `mixed_state_labels.wl`.

Progress on the POVM row, concluded: general POVMs by the Naimark chain (compat notebook, Phase 21). The row's remaining caveat — non-covariant effects — closes by one reduction: every POVM is a projective measurement on system $\otimes$ ancilla (Naimark), and a projective measurement on a pure state in any finite dimension is a Phase-14 chain. For a rank-one POVM the Naimark isometry has rows $\langle e_i|$, and it *is* an isometry iff $\sum_i E_i = I$: POVM completeness is the isometry property — the same completeness/unitarity mechanism Phase 18 identified as the no-signalling protection, now with no covariance assumption anywhere (verified: on an entangled pair, Bob's marginal is identical whether Alice performs the trine, the SIC, or a projective measurement). The ancilla is allocated locally and unitarily, realizing the intuition recorded at the fork that unitary quantum mechanics always allocates an ancilla. Two canonical non-covariant witnesses. *The trine* — the parent program's original obstruction object: parent Phase 1 proved fixed-cell models cannot reproduce sequential trine statistics, forcing hidden-measurement structure; the Naimark chain now realizes the trine constructively with apparatus carriers, its completeness kernel-checked ($\sqrt{3}$ algebra), its sampler exact, and its *sequential* statistics exact via the chain's Kraus post-states (symbolic residuals zero) — the arc that opened the parent program closes, with the obstruction real and the ancilla chain its resolution. *The SIC tetrahedron*: completeness exact, probabilities $(1 + b\cdot t_k)/4$, and informational completeness as a kernel theorem — the radical-free integer frame identity $\sum_k (b\cdot u_k) u_k = 4b$ — demonstrated end to end by reconstructing the Bloch vector of a *mixed* state (Phase-20 labels) from sampled chain frequencies to within 0.0033. The POVM row is now closed in full generality, and the protection ledger gains nothing new because nothing new is needed: completeness is the isometry is the no-signalling identity. Artifacts: `models/NaimarkChain.lean`, `naimark_chain.py`, `naimark_chain.wl`.

Progress on the DI row: the price of defeating certification (compat notebook, Phase 22). The row's open demand — quantify the minimum setting–hidden correlation needed to defeat device-independent certification — closes with a sharp law. Let a local deterministic model fake CHSH through setting-dependent hidden-variable distributions $p(\lambda|xy)$, with measurement dependence D the maximal pairwise total variation among the four conditionals. Then $S \leq 2 + 6D$, and the bound is attained: $D^*(S) = (S-2)/6$ exactly. The upper bound is kernel-certified in three pieces: every deterministic strategy's CHSH contribution is exactly ± 2 (the coefficient pattern $(s_{11}, s_{12}, s_{21}, -s_{22})$ has product -1, hence an odd number of minus signs — `decide` over all sixteen); against the barycenter each strategy contributes at most $2q$ plus its total deviation (the $|c| \leq 1$ pointwise bound); and total variation to the barycenter is at most $\frac{3}{4}D$ by convexity, assembling to the constant six. Attainment is a five-strategy gadget — weight $1-3d$ on the all-plus strategy plus d on each of the three bad-elsewhere strategies per setting pair — giving $E = (1, 1, 1, 1-6d)$, $S = 2+6d$, with *every* pairwise total variation equal to d (kernel-checked; exact rationals). Sharpness is certified on both sides: Z3 proves the 64-variable LP achieves $S = 2+6D$ and cannot exceed it by 10^{-3} at four exact rational points, and Mathematica's exact simplex independently returns $D^*(14/5) = 2/15$, $D^*(16/5) = 1/5$, $D^*(4) = 1/3$. Corollaries: defeating a Tsirelson-point certification costs $(\sqrt{2}-1)/3 = 0.1381$ of total variation (the explicit forger at that budget produces $S = 2.8286$, indistinguishable from quantum mechanics to any DI certifier); the house rational violation 14/5 costs exactly 2/15; the algebraic maximum costs 1/3. In the information currency the TV-optimal forger carries 0.172 bits of setting–hidden mutual information at the Tsirelson price, against Hall's information-optimal 0.0280 bits (Phase 10) — different corners of the same trade-off, both now quantified. Read against the constructive side: the SD-B models that actually reproduce quantum mechanics pay nothing in this currency (their no-signalling is identity, protected by the ledger); the $(S-2)/6$ price is what a flat-space LHV *forger* pays, and it is the amount of measurement independence every DI protocol implicitly assumes. Artifacts: `models/DiCertificationPrice.lean`, `di_certification_price.py`, `di_certification_price.wl`.

Progress on the normal form: chain forcing (compat notebook, Phase 23). The last structural question — is the steering chain merely *a* normal form or *the* normal form? — closes with a uniqueness theorem assembled from certified pieces plus new glue. In the filtration-exact sign-readout class (records revealed one at a time, each conditionally exact, outcomes read as carrier signs): (i) fresh per-step carriers are *forced* by sequential exactness — one carrier draw cannot serve all follow-up bases (the parent archive's hidden-measurement theorems); (ii) the conditional masses are *forced*, uniquely — exactness at the Born leaves inverts the telescoping, a triangular system with exactly one solution branch (kernel-checked; Mathematica `Solve` returns one branch with zero residuals), whose Jacobian determinant is the product of survivals — symbolically $(1-q_1)^2(1-q_2)$ at depth three — nonvanishing at every interior profile (kernel-checked; Z3 UNSAT for two distinct solutions): exactness pins every parameter, with no flat directions; (iii) the per-split carrier law is *forced* to the dipole basin by the saturation mechanism (Phases 7 and 10); and (iv) the residual freedom is *exactly* the canonical gauge — skeleton choice (the Phase-16 rank-chart orbit) and the per-split azimuth stabilizer (Phase 15), a group (kernel-checked closure), demonstrably statistics-free (random azimuthal rotations of every split's Bloch vector leave the joint law unchanged; the $y = 0$ in Phase 14's frame was a gauge fixing, now classified). **The theorem**: every exact model in the class is the steering chain up to canonical gauge. The program's three structural theses close into one: the chain is the normal form (existence), the gauge is canonical (transport), and the form is forced (uniqueness). Honest scope: sign-readout is where per-split forcing is proven; arbitrary stochastic readouts remain open structure theory. Artifacts: `models/ChainForcing.lean`, `chain_forcing.py`, `chain_forcing.wl`.

Progress on computation at scale: circuits on the chain and the caliber bookkeeping (compat notebook, Phase 24). The row's two clauses close together. *Coverage*: a circuit is deterministic unitary transport punctuated by measurements, i.e. chains (Phase 14) with recursion (Phase 19); the witness is Grover search, whose amplitude dynamics is an exact *rational* reflection recursion — one step from the uniform start lands on $(5/2,\, 1/2)$ in $1/\sqrt{8}$ units, a second on $(11/4,\, -1/4)$, with success probabilities exactly 25/32 and 121/128 (kernel-checked; Z3; Mathematica confirms both the recursion and the closed form $\sin^2((2k+1)\arcsin(1/\sqrt{M}))$ exactly) — and the carrier chain reproduces them, finding the marked item in 20000/20000 runs of the exact two-qubit instance and at 0.7803 and 0.9452 for the three-qubit targets, with an intermediate-measurement variant exercising the recursion at circuit level. *The* 4^{-N} *bookkeeping*: the compatible fraction tensorizes to 4^{-m} for m measurements (the per-measurement quarter is the fair record–success pair) — exponentially small, as Theorem 4.5 states — but the *cost* the boundary pays is the log-affinity, and it is additive: $\log 4^m = m \log 4$, exactly two bits per measurement, linear in circuit volume, with per-split density ratios bounded by the Phase-7 affinity (all kernel-checked; path-mass products verified at circuit scale). Nothing diverges at any finite scale; the only divergence in the program remains the projective idealization's cusp, per-measurement rather than per-scale. *Accounting*: the chain's parameter count per N-qubit round is the geometric sum $2^N - 1$ (kernel-checked), $(2^N)^T$ record paths over T rounds — the exponential sits in the boundary data, the quantum state itself, exactly where quantum mechanics puts it, while the dynamics stay local and per-split constant. No efficient classical sampler falls out, so compatibility does not collapse sampling-hardness beliefs: SD-B pays for computation in boundary data, not in dynamics. L1 holds sector-by-sector because every sector is the same chain. Artifacts: `models/CircuitCaliber.lean`, `circuit_caliber.py`, `circuit_caliber.wl`.

Progress on the QFT row: opened by the two-mode squeezed vacuum (compat notebook, Phase 25). The last untouched row acquires its first verified territory — opened, not closed. The two-mode squeezed vacuum $|\psi(r)\rangle = \sqrt{1-t}\, \sum_n t^{n/2} |n, n\rangle$ ($t = \tanh^2 r$) has *geometric* Schmidt data $\lambda_n = (1-t)t^n$, and three consequences follow. *Truncation control*: the truncated mass is exactly $1 - t^N$ with tail t^N (kernel-checked; symbolic including the infinite sums), so every finite truncation is exactly a Phase-18 bipartite pure state with exponentially controlled error (2×10^{-10} by $N = 16$ at $t = 1/4$). *Reeh–Schlieder's operational shadow*: every Schmidt weight is strictly positive (kernel-checked) — the reduced state is thermal and full-rank, no local operator annihilates the vacuum side, with mean photon number $t/(1-t) = \sinh^2 r$; the state is the thermofield double, so the Phase-20 label machinery covers thermal field states verbatim. *Field-sector Bell violation on the chain*: parity pseudospin gives $E(a, b) = \cos\theta_a \cos\theta_b + K \sin\theta_a \sin\theta_b$ with the coherence $K = 2\sqrt{t}/(1+t) = \tanh 2r$ derived from the state symbolically; parities match exactly on the Schmidt diagonal (kernel-checked, $E(z, z) = 1$); at the rational squeezing point $t = 1/4$, $K = 4/5$ and CHSH $S = 2\sqrt{1+K^2} = 2\sqrt{41}/5 = 2.5612$ (kernel-checked, violating since $\sqrt{41} > 5$), approaching Tsirelson as $K \to 1$ (kernel-checked); and the *two-wing carrier sampler* — the same chain as everything else, with degenerate binary pseudospin splits — gives $S = 2.559$ against the LHV bound 2. Nothing about field statistics at the two-mode operational level resists the SD-B architecture. Honest scope, recorded: algebraic QFT proper — type III von Neumann algebras, exact microcausality, genuine Reeh–Schlieder cyclicity, Tsirelson from first principles, multimode and curved-background structure — remains beyond the toolset; the row stays open with its frontier named. Artifacts: `models/TmsvChain.lean`, `tmsv_chain.py`, `tmsv_chain.wl`.

A second protecting symmetry, already in the archive. The parent's censoring line (Phases 21–26) is the dynamical counterpart of the parity result and belongs to Question 6.1's ledger. Any finite-time smooth implementation of the terminal condition leaves a no-record sector $\perp$, and the postselected law is Born tilted by the cumulative survival fraction — exactly Born only when that fraction is history-independent. The Riccati-calibrated multibaker achieves the exact conditional sequential law at every depth, and the calibration is protected not by tuning but by a symmetry: *branch* *normalization*, one outcome-blind completion rule on the normalized outcome factor whose pullbacks automatically yield the proportional censoring widths $\alpha = q\varepsilon$, $\beta = (1-q)\varepsilon$, with uniqueness of the proportional censor Z3-certified. Parity protects the marginals (no-signalling); branch normalization protects the conditional Born law under filtering. The sharpened conjecture: every exact operational identity of the settings-dependent theory is the fixed locus of an identifiable symmetry. A corollary for the audit's honesty: in every finite-time smooth implementation, "exact compatibility" means exact *conditional on resolved records*; the unconditional theory is distinguishable from quantum mechanics through the no-record rate, which is this class's universal experimental signature alongside the small-angle slope of Corollary 2.4 and the derived filter scale of Remark 2.5.

The audit's verdict

With the twenty-five phases in place, the founding question has its answer, graded by model class as Section 1 demanded. *SD-F*: exactly incompatible (Section 2), with the smoothing theorem locating the entire hard-boundary structure in the projective idealization — real, strictly-unsharp detectors need only smooth densities and finite affinity. *SD-C*: incompatible at every resolution, now quantitatively — the deviation–spread law is a resolution-independent floor, so ontically finite programs face a fixed naturalness–accuracy exchange rate rather than a convergence path. *SD-B*: compatible, constructively and at quantified cost, across a verified sector spanning all pure two-qubit states, entanglement swapping, GHZ, unsharp qubit pairs, and all finite-dimensional pure projective measurements with Lüders sequences. Three structural theses carry the sector. First, the *record-conditioned steering chain* is its constructive normal form: one mechanism, four times deployed (bipartite, swapping, GHZ, qudit), whose telescoping is simultaneously the exactness proof, the no-signalling proof, and the measure-noncontextuality proof. Second, the *saturation mechanism* — domination plus equal mass forces equality — carries every rigidity result in the program (the a.c. deficit, the equivariance collapse, the sharp-limit basin, one-wing rigidity: five appearances). Third, the *protecting-structure ledger is complete, and now grounded*: source parity, branch normalization, and record-affinity exhibit every exact no-signalling identity of the verified sector as a fixed locus, and Phases 15–16 upgrade the ledger from inventory to derivation — record-affinity is the law of total probability for the canonical single-law microdynamics; implementation conjugacy holds on the full rotation orbit by its canonical symplectic automorphisms; and between any two implementations of the same effect — different chain orderings, different completing bases — there is a *unique* canonical transport, the rank-chart intertwiner, whose canonicity is a rigidity theorem (given the skeleton's intrinsic filtration order, the monotone chart is the identity's only companion) rather than a choice. The conspiracy objection is thereby answered in the standard way physics answers fine-tuning, with the protecting structures themselves theorems of one invariant law plus its symmetry group. The hard implementation-conjugacy item — the audit's last heavyweight — is closed. What remains open is only the residual deviation floor–ceiling bracket (beyond orbit tools) — entangled measurements at the carrier level closed with the Bell-basis chain (Phase 17), the bipartite row extended to every dimension by the two-wing chain (Phase 18), the multipartite induction certified on instances (Phase 19), mixed states derived as purification marginals with the convex label proven gauge (Phase 20), the POVM row concluded in full generality by the Naimark chain (Phase 21), the DI row quantified by the sharp forgery price $D^*(S) = (S-2)/6$ (Phase 22), the chain proven to be THE normal form — unique up to canonical gauge in its class (Phase 23), computation at scale closed with a linear caliber budget and the exponential located in boundary data (Phase 24), and the QFT row opened at its two-mode operational core — the squeezed vacuum on the chain, Bell-violating to the Tsirelson limit (Phase 25). What remains open: the deviation floor–ceiling bracket, and algebraic QFT proper. Everything this audit set out to decide is decided.

Provenance: Parent Phases 21–26; `parent/S4RiccatiCalibratedMultibaker.lean`, `parent/S4BranchNormalizedSelector.lean` and paired scripts.

8 Planck scale and the convergence program

By Corollary 3.6, whether a Planck-scale framing is finite-precision depends on whether the Planck scale is read ontically or operationally. Under the ontic reading every framing is L2, quantum mechanics' own continuum statistics become an idealized limit, and the fine-tuning debate loses its asymmetry: all parties hold effective theories, and the comparison shifts to deviation structure. The resolution-independent results of Section 3 then become the main payload, since they characterize the deviation signature of the entire finite class: the five-angle rationality pattern (Theorem 3.2), the forced hidden-measurement structure (Theorem 3.3), the one-third spread ceiling (Theorem 3.4), the broken-covariance requirement (Theorem 3.5), and the small-angle cusp slope (Corollary 2.4).

The deviation dictionary

The archive's obstruction phases, reread as approximation theory, assemble into a dictionary of universal deviation laws for the finite-resource class. Each is a sharp floor with an attaining model, so the class predicts not only *that* finite implementations deviate but *where*, *which way*, and *how fast*:

Signature	Law	Provenance
No-record floor	$\varepsilon_T \geq \frac{2G}{L_R L_\gamma} e^{-B_T}$ with $B_T = \int_0^T \|DX_{H_t}\|\, dt$; a target rate η costs $B_T \geq \log(C/\eta)$; sharp — the hyperbolic selector attains it	Phases 17–19
Bandwidth form	$\|\mathrm{Hess}\, H\| \leq N^2 A_1$; a bandlimited implementation needs latency $T \geq \log(C/\eta)/(N^2 A_1)$	Phase 20
Postselection bias	one-stage residual $\delta(2q-1)/(1-2\delta)$; bias per unresolved mass $\to (2q-1)/2$; universal total-variation ceiling η; zero only at balanced splits or under branch-normalized censoring	Phases 21–23, 25–26
Censoring-evasion cost	hiding the residual by proportional censoring needs selector curvature of order the inverse unresolved width squared	Phase 24
Pointer width	false-negative probability $\delta^2/6$ for an independent ready rectangle at width-to-gain ratio δ	Phase 11
Repeatability neck	unresolved comparator mass linear in inverse gain at every interior overlap ($M_T \varepsilon \to 4G\sqrt{1-d^2}/\pi$), quadratic exactly at repeatability; overlap-uniform floor $\varepsilon_d(\delta) \geq \delta^2$	Phases 31–32
Singlet cusp slope	$E_\gamma(\theta) = -1 + \frac{2r_\gamma}{\pi}\theta + o(\theta)$, $r_\gamma \sim 2/\sqrt{\pi\gamma}$	Phase 42
Filter gap scale	the finite-depth survival filter deviates at the spectral-gap rate of the survival kernel (Remark 2.5)	Phase 7
A.c. entanglement deficit	any equivariant absolutely continuous two-carrier profile falls short of the exact bipartite correlations by mass $\geq (s^2/c)\,\mathrm{artanh}\, c$	compat Phase 3
Affinity cost of sharpness	exact statistics at sharpness η need conditioning affinity $\kappa(\eta) = \log \frac{1+3\eta}{1-\eta}$, finite below and divergent at the projective limit	compat Phase 7
One-wing information floor	any exact one-wing-conditioned singlet model is unique and carries exactly $\ln 2 - \frac{1}{2}$ nats of setting-carrier information; premium over the two-setting minimum 0.2506 bits	compat Phase 10
Deviation–spread floor	counting models at spread fraction $\sigma > \frac{1}{3}$ deviate by $\geq (\sigma/2 - 1/6)/(\frac{8}{3}\sigma)$ at every resolution (3/80 at $\sigma = \frac{5}{12}$, 1/16 at $\frac{1}{2}$); resolution does not shrink it; triply certified and final for all shift-orbit relaxations	compat Phases 11–13

The linear-to-quadratic crossover at repeatability and the sign structure of the postselection bias are parameter-free *shape* predictions; the exponential budget laws convert every proposed implementation resource (gain, latency, bandwidth, curvature) into a floor on observable anomaly rates. These are the continuum-side entries of the $\Theta(f(K))$ dictionary sought by Question 8.1; the arithmetic floors of Section 3 are its counting-side entries. Jointly they make the finite-resource superdeterministic class phenomenologically over-determined — the opposite of unfalsifiable.

Question 8.1 (Two-sided convergence bound)**.** Determine $f(K)$ such that the best achievable uniform Born deviation of a counting model at ontic resolution K is $\Theta(f(K))$. The constructive cell-splitting lift (Phase 3) and the $(t, M) = (8, 8)$ witness with deviation $\leq \frac{1}{16}$ give upper-bound machinery; the $K/3{+}1$ UNSAT certificates, Theorem 3.3, and Theorem 3.2 give lower-bound machinery. Whether the deviation tends to zero at all under the full constraint set, and at what rate, is open; its answer calibrates every ontically Planck-finite proposal, including [28, 24].

9 Ordered program

This program has been executed in full: every item below is closed or answered, and the audit then continued past the list through Phases 17–25 (Section 7). What remains open in the entire program is the deviation floor–ceiling bracket (beyond orbit tools) and algebraic QFT proper. The list is preserved as originally posed — each item named its success criterion and was gated on an audit row or a standing question — with its disposition appended in brackets.

1. **Entanglement swapping.** Extend the SD-B architecture to the event-ready configuration: two sources, a Bell-state measurement in the common future, terminal conditioning including the BSM record. Success: the swapped-pair Lüders law from local interactions; failure mode: a no-go showing common-future conditioning cannot replace a common source, which would kill the architecture cheaply. *[Closed positively, Phase 6; the station itself realized on carriers, Phase 17.]*
2. **S1 conjugacy microdynamics ($d \geq 3$).** Construct a microdynamics whose implementations of the same effect form single conjugacy orbits of (Ω, T, μ), discharging measure-noncontextuality and, via [8, 9], the Born rule in every dimension. Success: the qubit constructions of Section 4 recovered as the $d = 2$ case. *[Closed, Phases 14–16: the chain realizes measure noncontextuality; the canonical microdynamics derives the conjugacy; the rank charts supply the cross-skeleton intertwiner.]*
3. **Partial entanglement.** Test whether weighted two-basin mixtures reproduce $\cos\xi\,|00\rangle + \sin\xi\,|11\rangle$; the Haar-moment method of Theorem 4.4 and the trace formula of Theorem 2.1 give fast no-go machinery, and the biased marginals identify the first obstruction. *[Closed, Phases 2–5: two-carrier forcing, exact a.c. deficit, steering dilation, equivariance collapse; extended to every dimension, Phase 18.]*
4. **POVM smoothing theorem.** Prove or refute: every POVM with sharpness bounded away from projective admits an exact SD-F realization (smooth positive densities, no cusp), with the cusp of Theorem 2.1 arising only in the projective limit. This would recast the cusp as the signature of ideal measurements and connect directly to the finite-resolution experiments of the parent archive. *[Proven, Phase 7; the POVM row concluded in full generality, Phase 21.]*
5. **Multipartite gate analysis.** Generalize the moment method of Theorem 4.4 to N wings; determine whether exclusive selection survives GHZ constraints and what monogamy costs in selector information. *[Closed, Phase 8 (GHZ, monogamy as architecture); the induction certified on W and cluster instances, Phase 19.]*
6. **Symmetry protection (Question 6.1).** Formalize the Wood–Spekkens objection as a measure-zero statement over model parameters and search for the protecting Ward identity; either outcome resolves the conspiracy question. *[Answered, Phases 1 and 9: the protecting Ward identity exists (source parity; record-affinity), and Phase 15 derives it from the canonical microdynamics.]*
7. **Convergence bound (Question 8.1).** Settle $\Theta(f(K))$ for the counting class. *[Settled in an unexpected direction, Phases 11–13: there is no decay in K — the deviation–spread law is a resolution-independent floor, final for shift-orbit relaxations.]*
8. **Information premium.** Prove or refute: among models whose Born sector masses are derived from fiber-volume projections (rather than imposed), the setting–hidden information is bounded below by a constant exceeding Hall's 0.028 bits; Theorem 4.6(iii) frames the gap. *[Proven, Phase 10: the unique one-wing model carries a forced 0.2506-bit premium.]*

As planned at the fork, items 1–3 were attempted first and mined existing machinery; item 6 — the decisive question for the program's philosophical standing — resolved in favor of symmetry protection rather than tuning; and items 4, 7, and 8 resolved as theorems (the smoothing theorem, the deviation–spread law, one-wing rigidity).

A Provenance map

All theorems above are restatements of machine-checked results. The table maps each to its origin; Lean files verify in the `borncount` Mathlib project, Python scripts run under `python3.12` with Z3. Paths are relative to this document's directory, and every referenced file ships with it: the parent-archive artifacts cited below are copied verbatim into `parent/` (phase numbers in this table are the parent archive's; entries given as bare theorem names, `no_exact_trine_cells` and `witness_certified`, are remote kernel jobs recorded in the parent notebook, reachable through `models/NOTES.md`), while the compatibility program's own proofs accumulate in `models/` with their phase notebook `models/NOTES.md`.

Result	Phase	Lean artifact
Thm 2.1	39	`S4OneSidedLiouvilleSinglet.lean`
Thm 2.2	47	`S4SoftTerminalAffinity.lean`
Thm 2.3	48	`S4MeasureClassNoGo.lean`
Cor 2.4	42	`S4BoundaryApparatus.lean`
Thm 3.3	1	`no_exact_trine_cells`
Thm 3.4	2–4	`SixCycle.lean`, `witness_certified`
Thm 3.5	5	`S4Topology.lean`
Thm 4.1	8–9	`S4HamiltonianRecord.lean`, `S4Sequential.lean`
Thm 4.2	39–40	`S4EndogenousSettingDilation.lean`
Thm 4.3	44	`S4CommonSelectorBoundary.lean`
Thm 4.4	45–46	`S4ExclusiveSelectorUniqueness.lean`, `S4CommonFutureCompiler.lean`
Thm 4.5	43	`S4CanonicalHistoryMeasure.lean`
Thm 4.6	36–39	`S4FiniteSpeedBell.lean`, `S4MinimumInformationBell.lean`, `S4ContinuousHallBri`
Thm 4.7	41	`S4InterventionalFork.lean`
Conjugacy transport (§5)	74–75	`S4DynamicConsistencyBoundary.lean`, `S4NaturalSupportSelection.lean`

The compatibility program's own artifacts (Section 7) live in `models/`, verified in the remote `compat` Mathlib project (kernel job ids are recorded per phase in `models/NOTES.md`):

Phase	Artifacts (`.lean` / `.py`)
1	`QWardParityDichotomy` / `q_ward_parity_dichotomy`
2	`QWardEquivariantBipartite` / `q_ward_equivariant_bipartite`
3	`QWardAcCarrier` / `q_ward_ac_carrier`
4	`SteeringDilation` / `steering_dilation`
5	`EquivarianceCollapse` / `equivariance_collapse`
6	`SwappingDilation` / `swapping_dilation`
7	`PovmSmoothing` / `povm_smoothing`
8	`GhzChainDilation` / `ghz_chain_dilation`
9	`QWardRecordFlip` / `q_ward_record_flip`
10	`InformationPremium` / `information_premium`
11	`DeviationSpreadLaw` / `deviation_spread_law`
12	`DeviationSpreadLaw2` / `deviation_spread_law2`
13	`DeviationSpreadLaw3` / `deviation_spread_law3`
14	`QuditChainDilation` / `qudit_chain_dilation`
15	`CanonicalMicrodynamics` / `canonical_microdynamics`
16	`RankChartIntertwiner` / `rank_chart_intertwiner`
17	`BellBasisChain` / `bell_basis_chain`
18	`BipartiteQuditChain` / `bipartite_qudit_chain`
19	`MultipartiteChainInduction` / `multipartite_chain_induction`
20	`MixedStateLabels` / `mixed_state_labels`
21	`NaimarkChain` / `naimark_chain`
22	`DiCertificationPrice` / `di_certification_price`
23	`ChainForcing` / `chain_forcing`
24	`CircuitCaliber` / `circuit_caliber`
25	`TmsvChain` / `tmsv_chain`

The parent archive (its own document and 134-phase notebook, reachable through the linkage recorded in `models/NOTES.md`) remains the complete record, including the selector, record-carrier, clock, and gravitational-consistency analyses not mined here; those address the internal consistency of the terminal-support axiom and are orthogonal to the compatibility audit above. The parent artifacts cited in this document are copied verbatim into `parent/`, so the audit is self-contained as distributed.

References

[1] J.-M. A. Allen, J. Barrett, D. C. Horsman, C. M. Lee, and R. W. Spekkens, "Quantum common causes and quantum causal models," *Phys. Rev. X* **7**, 031021 (2017), arXiv:1609.09487.

[2] J. Barrett and N. Gisin, "How Much Measurement Independence Is Needed to Demonstrate Nonlocality?" *Phys. Rev. Lett.* **106**, 100406 (2011), arXiv:1008.3612.

[3] J. S. Bell, "On the Einstein Podolsky Rosen Paradox," *Physics Physique Fizika* **1**, 195–200 (1964), doi:10.1103/PhysicsPhysiqueFizika.1.195.

[4] J. S. Bell, "On the Problem of Hidden Variables in Quantum Mechanics," *Rev. Mod. Phys.* **38**, 447–452 (1966), doi:10.1103/RevModPhys.38.447.

[5] J. S. Bell, "Free variables and local causality," *Epistemological Letters* **15** (1977); reprinted in *Speakable and Unspeakable in Quantum Mechanics*, Cambridge University Press (1987).

[6] The BIG Bell Test Collaboration, "Challenging local realism with human choices," *Nature* **557**, 212–216 (2018), doi:10.1038/s41586-018-0085-3; arXiv:1805.04431.

[7] C. H. Brans, "Bell's Theorem Does Not Eliminate Fully Causal Hidden Variables," *Int. J. Theor. Phys.* **27**, 219–226 (1988), doi:10.1007/BF00670750.

[8] P. Busch, "Quantum states and generalized observables: a simple proof of Gleason's theorem," *Phys. Rev. Lett.* **91**, 120403 (2003), arXiv:quant-ph/9909073.

[9] C. M. Caves, C. A. Fuchs, K. K. Manne, and J. M. Renes, "Gleason-Type Derivations of the Quantum Probability Rule for Generalized Measurements," *Found. Phys.* **34**, 193 (2004), arXiv:quant-ph/0306179.

[10] G. S. Ciepielewski, E. Okon, and D. Sudarsky, "On Superdeterministic Rejections of Settings Independence," arXiv:2008.00631 (2020).

[11] R. Colbeck and R. Renner, "No extension of quantum theory can have improved predictive power," *Nature Communications* **2**, 411 (2011), doi:10.1038/ncomms1416; arXiv:1005.5173.

[12] J. Degorre, S. Laplante, and J. Roland, "Simulating quantum correlations as a distributed sampling problem," *Phys. Rev. A* **72**, 062314 (2005), doi:10.1103/PhysRevA.72.062314; arXiv:quant-ph/0507120.

[13] S. Donadi and S. Hossenfelder, "A toy model for local and deterministic wave-function collapse," *Phys. Rev. A* **106** (2022), arXiv:2010.01327.

[14] A. S. Friedman, A. H. Guth, M. J. W. Hall, D. I. Kaiser, and J. Gallicchio, "Relaxed Bell inequalities with arbitrary measurement dependence for each observer," *Phys. Rev. A* **99** (2019), arXiv:1809.01307.

[15] A. M. Gleason, "Measures on the Closed Subspaces of a Hilbert Space," *J. Math. Mech.* **6**, 885–893 (1957), doi:10.1512/iumj.1957.6.56050.

[16] M. J. W. Hall, "Local Deterministic Model of Singlet State Correlations Based on Relaxing Measurement Independence," *Phys. Rev. Lett.* **105**, 250404 (2010), doi:10.1103/PhysRevLett.105.250404; arXiv:1007.5518.

[17] M. J. W. Hall, "The significance of measurement independence for Bell inequalities and locality," arXiv:1511.00729 (2015).

[18] J. R. Hance, S. Hossenfelder, and T. N. Palmer, "Supermeasured: Violating Bell-Statistical Independence without violating physical statistical independence," *Found. Phys.* **52**, 81 (2022), arXiv:2108.07292.

[19] L. Hardy, "Quantum ontological excess baggage," *Stud. Hist. Phil. Mod. Phys.* **35**, 267–276 (2004), doi:10.1016/j.shpsb.2003.12.001.

[20] S. Hossenfelder and T. N. Palmer, "Rethinking Superdeterminism," *Frontiers in Physics* **8**, 139 (2020), doi:10.3389/fphy.2020.00139; arXiv:1912.06462.

[21] S. Kochen and E. P. Specker, "The Problem of Hidden Variables in Quantum Mechanics," *J. Math. Mech.* **17**, 59–87 (1967), doi:10.1512/iumj.1968.17.17004.

[22] A. Montina, "Exponential complexity and ontological theories of quantum mechanics," *Phys. Rev. A* **77**, 022104 (2008), doi:10.1103/PhysRevA.77.022104.

[23] I. Niven, *Irrational Numbers*, Carus Mathematical Monographs 11, Mathematical Association of America (1956).

[24] T. N. Palmer, "Discretised Hilbert Space and Superdeterminism," arXiv:2204.05763 (2022).

[25] C. Pfister and S. Wehner, "An information-theoretic principle implies that any discrete physical theory is classical," *Nature Communications* **4**, 1851 (2013), doi:10.1038/ncomms2821; arXiv:1210.0194.

[26] I. Sen and A. Valentini, "Superdeterministic hidden-variables models I: non-equilibrium and signalling," *Proc. R. Soc. A* **476** (2020), arXiv:2003.11989.

[27] I. Sen and A. Valentini, "Superdeterministic hidden-variables models II: conspiracy," *Proc. R. Soc. A* **476**, 20200214 (2020), arXiv:2003.12195.

[28] G. 't Hooft, *The Cellular Automaton Interpretation of Quantum Mechanics*, Springer (2016), doi:10.1007/978-3-319-41285-6; arXiv:1405.1548.

[29] G. 't Hooft, "Explicit construction of Local Hidden Variables for any quantum theory up to any desired accuracy," arXiv:2103.04335 (2021).

[30] K. B. Wharton, "The Universe is not a Computer," in *Questioning the Foundations of Physics*, A. Aguirre, B. Foster, and Z. Merali (eds.), Springer (2015), arXiv:1211.7081.

[31] K. B. Wharton, "Lagrangian-Only Quantum Theory," arXiv:1301.7012 (2013).

[32] K. B. Wharton, *A New Class of Retrocausal Models*, arXiv:1805.09731 (2018).

[33] K. B. Wharton and N. Argaman, "Colloquium: Bell's theorem and locally mediated reformulations of quantum mechanics," *Rev. Mod. Phys.* **92**, 021002 (2020), doi:10.1103/RevModPhys.92.021002; arXiv:1906.04313.

[34] C. J. Wood and R. W. Spekkens, "The lesson of causal discovery algorithms for quantum correlations: causal explanations of Bell-inequality violations require fine-tuning," *New J. Phys.* **17**, 033002 (2015), doi:10.1088/1367-2630/17/3/033002.