

The Price of Locality

Exact Minimum Measurement Dependence for Multipartite Quantum Correlations, and a Race Among Lawful Mechanisms

Aaron Alai, with Claude (Anthropic) as mathematical collaborator — July 18, 2026

Abstract

Bell derivations rest on locality, determinism, and measurement independence. Hall [2] priced the third assumption exactly for the singlet state: $(\sqrt{2} - 1)/3 \approx 13.81\%$ of measurement independence must be surrendered for a local deterministic model to reproduce all singlet correlations. We extend that program to multipartite stabilizer scenarios and compute the exact minimum measurement dependence required to reproduce the full faithful statistics of GHZ/Mermin experiments — the optimal-model problem Hall posed and left open in 2011, having priced the tripartite perfect correlators alone at $1/3$ [11]: $F(3) = F(4) = 1/3$, $F(5) = F(6) = 2/5$, $F(7) = F(8) = 4/9$, $F(9) = F(10) = 8/17$, $F(11) = F(12) = 16/33$, $F(13) = 32/65$ **[LP-EXACT, two-machine verified]**, with hand proofs at $n = 3, 4$; the reduction theorem shows the faithfulness constraints are free, so Hall’s correlator-only tripartite threshold is promoted to the faithful value, and everything at $n \geq 4$ is new. All eleven points obey the closed law $F = R/(2(R+1))$, with $R = 2^{\lfloor (n-1)/2 \rfloor}$ the Mermin violation ratio [15]; a proven combinatorial lower bound (Theorem 5) is tight on every computed core and homogeneous family; and a universal ceiling $\mathbf{F} \leq 1/2$ **[PROVEN]** shows the statistics never *require* total superdeterminism — a faithful model surrendering no more than half exists at every size. The optimal hidden-variable densities have a closed physical form: uniform measures on the contextual ground states of the prepared state’s frustrated stabilizer Hamiltonian — an inheritance principle confirmed out-of-sample on cluster states. The landscape is glassy (energy is spectral, supported on stabilizer characters) and provably beyond any pairwise interaction. We then pose the supply question as an open race: every published lawful-measurement-dependence program — invariant-set theory, cellular-automaton interpretations, future-input dynamics, all-at-once retrocausality, and the displacement-condensate framework of the present authors — is graded against this one answer key, steelmanned in its own formalism, with several opening calculations executed here on rivals’ behalf. The condensate mechanism’s consolidated derivational debt is paid in-model: under exact $U(1)$ conservation with no external phase reference, a collective emission leaves a superselected two-branch record whose coherence support equals the prepared state’s stabilizer group, forcing the effective energy’s term content to coincide with the required Hamiltonian’s — winding- n coherence amplitude $1/2$ exactly, partial-winding contamination identically zero, residual exposures quantified against known physics, and one named test (the continuum finite- β vise) that could kill it. Operationally, the floors are counterfeiting thresholds for multipartite device-independent certificates. Version 1’s central numerical claim is withdrawn and replaced by a principle; the correction is documented in Part 0. Every tagged result is reproducible from the linked repository in under one minute.

Status tags used throughout: **[PROVEN]** (proof included or sketched), **[LP-EXACT]** (exact rational optimum of a finite linear program, machine-verified, key values reproduced on independent hardware/solvers), **[CONJECTURE]** (closed form fitting all computed points), **[TOY]** (mechanism-level result in an explicit toy model).

Part 0 — Version note and correction of record

Version 1 of this paper — and the companion measurement-independence paper [25], which this work supersedes — argued that the DST condensate coupling $g_0^2 = 9/64 \approx 14.06\%$ approximately saturates Hall’s exact bipartite minimum of $(\sqrt{2} - 1)/3 \approx 13.81\%$ (Hall [2]; the variational-distance bound, distinct from the mutual-information bound of Barrett-Gisin [3] — the two must not be conflated), and offered the near-coincidence as structural evidence. The multipartite computations below refute that reading: the required dependence is not a constant near 14% but climbs lawfully in pairs toward $1/2$, so a coupling that saturates the bipartite floor funds nothing at $n \geq 3$, and the coincidence carried no evidential weight.

Version 1’s central numerical claim is therefore withdrawn, and the companion paper’s saturation argument is withdrawn with it; the bipartite framework is no longer a viable candidate for the supply story. What survives of the companion work is

what was independent of the coincidence: Hall’s bound itself as the correct bipartite price, and the accounting-relation reading of effective-correlation corrections (extremal quantities such as Tsirelson’s bound are fixed by the geometry; corrections land in bookkeeping between theoretical quantities, never in laboratory statistics) — a reading that §11.5 now elevates from interpretive caution to the mechanism’s explicit obligation. The withdrawn claim is replaced by the *flooding postulate* (§10): a physical medium supplying dependence lawfully and uniformly naturally supplies the saturation value of the ceiling theorem, $\sim 1/2$, sufficient for every stabilizer experiment of every size — a postulate this paper subsequently converts into a computed limit (§11.5: the ceiling is the hot limit of the supply, the LP optimum its ground state). We consider the exchange an upgrade — a coincidence has been replaced by a principle with derivational obligations, most of which are now paid — and we document it because visible self-correction is the only credential that cannot be counterfeited.

1. Introduction

Bell’s inequalities follow from locality, determinism, and measurement independence — the statistical independence of setting choices from the hidden state of the measured system. Experiments violate the inequalities; one assumption must yield. The first two exits are heavily populated (nonlocal hidden variables; abandonment of hidden values). The third exit — *lawful, partial* measurement dependence, as distinguished from initial-condition conspiracy — has a small but serious literature: Brans’s existence proof [4]; Hall’s exact bipartite pricing [2] and its per-observer refinement by Friedman, Guth, Hall, Kaiser and Gallicchio [5], with the Hall-Branciard synthesis [12] and Kochen-Specker extension [11]; Barrett-Gisin’s mutual-information bound [3]; Hossenfelder-Palmer’s programmatic defense [6] and the Donadi-Hossenfelder toy dynamics [18]; Palmer’s invariant-set and rational-quantum-mechanics program [7]; ‘t Hooft’s cellular-automaton interpretation [8]; the all-at-once formulations of Wharton, Adlam, and collaborators [9]; taxonomy work by Waegell and McQueen distinguishing conspiratorial from nomic dependence [10]; the fine-tuning analysis of Wood-Spekkens [19]; and the explicit superdeterministic models of Sen-Valentini [20].

Adjacent prior work delimits our novelty claim precisely, and the closest point of contact is Hall himself. In the Kochen-Specker section of Ref. [11] (2011), Hall priced the tripartite Mermin scenario — but for the four *perfect correlators alone*, with no faithfulness constraints: his Table I model achieves $M = 2/3$, a surrendered fraction of $1/3$, and he explicitly left open the problem of “a similar optimal model for Mermin’s state” itself. Our $F(3) = 1/3$ answers that fifteen-year-old question with the same number, and Theorem 1 explains why the numbers coincide rather than merely happening to: the fiber-uniform lift makes the marginal

constraints free at every n , so the correlator-only and full-faithful optima are provably equal. Hall’s 2011 threshold is thereby promoted to the faithful value; everything at $n \geq 4$ — the staircase, the ceiling, the landscape and its inheritance structure — has no precedent we can find. (A convention caution: Hall’s symbol F in [11] denotes the *retained* fraction $1 - M/2$; ours denotes the *surrendered* fraction $M/2$.) Hossain et al. [13] derived relaxed *inequality bounds* for the tripartite Mermin, Svetlichny, and NS_2 inequalities — the dependence sufficient to reach the inequalities’ quantum *values*. Roy et al. [14] priced tripartite equatorial measurement classes under per-party bookkeeping. Most recently, Pal, Sharma, and Podila [26] introduced *benchmark dependence*, a task-level generalization of measurement dependence, proving a universal additive bound on game scores ($S_\eta \leq S_{cl} + \eta$) and applying it to CHSH, the GHZ/Mermin game, and the magic-square game; their bound governs reproduction of a game’s *score* in a total-variation-from-marginal measure (related to Hall’s pairwise M by $M/4 \leq \eta \leq M/2$), and for a perfect $n = 3$ Mermin score yields the generic lower bound $1 - S_{cl} = 1/4$ — weaker than, and incomparable in bookkeeping to, the exact faithful $1/3$ computed here. None of [13], [14], [26] computes the minimum overall variational dependence for full faithful reproduction of the statistics (every correlator, every marginal), which is this paper’s object. The approaches meet consistently at the boundary: our exact tripartite optimum, allocated symmetrically (one-sided parameters $M_1 = M_2 = M_3 = 2/3$), saturates Hossain et al.’s [13] necessary condition $2M_i + M_j = 2$ with equality — independent methods corroborating one another.

Provenance of the question. This paper did not begin as a Bell paper. It began as an attempt to narrate the complete operational stack of a superconducting quantum computer — microwave pulse to bitstring — under a measurement-dependent ontology, to locate where such an ontology would actually differ. The exercise kept returning the same verdict: nowhere operational, since an interpretation that reproduces quantum statistics exactly offers hardware no error to shave (indeed, the world’s quantum processors run a continuous null experiment on any *excess* settings-dependence, and report none). What survived the filtering was a single quantitative question: not *whether* the measurement-independence exit is taken, but *how much* it costs — and the shared expectation was that the answer was already known, either Hall’s bipartite 13.81% (which the condensate coupling $g_0^2 = 9/64 \approx 14.06\%$ appeared to saturate, the coincidence advertised in version 1) or, failing that, everything ($F \rightarrow 1$, the conspiracy horn). The linear program of §3 was written to arbitrate between those two priors. It refused both: the floors climb a staircase toward one-half — a number large enough to kill the coincidence, lawful enough to kill the conspiracy reading, and unpublished enough that no mechanism in the literature had ever been asked to fund it. Everything in this paper is downstream of that refusal.

This paper has three parts. **Part I** (§§2–9) is interpretation-neutral mathematics: the exact demand curve of the measurement-dependence exit — floors, law, ceiling, and the physical form of the optimal models. **Part II** (§§10–11) poses the supply question as a race: one answer key, every published mechanism steelmanned in its own language, several of their opening calculations executed here on their behalf. **Part III** (§12) develops the displacement-condensate candidate in detail, installments and wounds alike. The Discussion (§13) gives the operational reading and the epistemic scoring rule for empirically tied rivals.

2. Scenario and definitions

For $n \geq 3$ parties sharing the n -partite GHZ state, each party measures X or Y . The Mermin settings are strings $s \in \{X, Y\}^n$ with an even number of Y ’s (2^{n-1} of them). Quantum

mechanics predicts with certainty $E(s) = +1$ for $\#Y \equiv 0 \pmod{4}$ and -1 for $\#Y \equiv 2 \pmod{4}$, with every proper-subset correlator vanishing. A local deterministic model assigns each hidden state λ pre-set answers $(a_i, b_i) \in \{\pm 1\}^2$ per party (4^n strategies) and, per settings string s , a density $\rho_s(\lambda)$. Hall's measure: $M := \max_{\{s, s'\}} \sum_{\lambda} |\rho_s - \rho_{s'}|$; the fraction of measurement independence given up is $F := M/2$. **Faithfulness** requires reproducing the deterministic full correlators *and* the vanishing marginals. Floors below are minima of F over all faithful models.

3. The linear program [PROVEN formulation]

Variables $\rho_s(\lambda) \geq 0$; equality constraints: normalization, full correlators $= E(s)$, all proper-subset correlators $= 0$; objective: minimize the maximum pairwise L1 distance (standard auxiliary-variable encoding). Rational data $\Rightarrow$ exact rational optima.

4. The reduction theorem [PROVEN]

Theorem 1. With $a_i = (-1)^{\{\alpha_i\}}$, $b_i = (-1)^{\{\beta_i\}}$, $P := \bigoplus \alpha_i$, $\gamma_i := \alpha_i \oplus \beta_i$, the LP is exactly equivalent to the same LP on the $2^{(n+1)}$ classes (P, γ) , with the marginal constraints deleted. *Proof sketch:* every Mermin correlator is the parity $P \oplus \bigoplus_{i \in S} \gamma_i$, so constraints depend on λ only through its class; each class is a $2^{(n-1)}$ -strategy fiber, and the fiber-uniform lift annihilates every proper-subset correlator (a non-constant linear functional on the fiber's affine subspace averages to zero) while preserving correlators; joint fiber-uniformization is group averaging and cannot increase L1 distances. ■ Practically: $n = 7$ collapses from 16,384 strategies to 256 classes, making the exact frontier commodity-computable. **Theorem 1b (symmetrization, lossless).** The scenario is invariant under party permutations (settings permute within fixed Hamming weight; targets depend only on $|s| \pmod{4}$; the class action is $(P, \gamma) \mapsto (P, \pi\gamma)$); the feasible set is convex and the objective $T = \max$ pairwise L1 is convex and permutation-invariant; hence group-averaging any optimum yields a covariant optimum, and restricting to S_n -covariant densities — functions of $(|s|, P, |\gamma \cap s|, |\gamma \cap \bar{s}|)$ only — is exact. Variable count drops from $O(4^n)$ to $O(n^3)$: $n = 9$ falls from 33 million variables to 3,673, and the $n = 13$ floor solves in twenty seconds. ■

5. Exact values [LP-EXACT]

Scenario	F_min	s (max satisfiability)	F·s	Status
CHSH @ Tsirelson ($n = 2$)	$(\sqrt{2}-1)/3 = 0.138071$	—	—	LP-EXACT; reproduces Hall's continuum minimum from four settings
Mermin $n = 3$	1/3	3/4	1/4	PROVEN + LP
Mermin $n = 4$	1/3	3/4	1/4	PROVEN + LP
Mermin $n = 5$	2/5	5/8	1/4	LP-EXACT
Mermin $n = 6$	2/5	5/8	1/4	LP-EXACT (two independent solver methods)
Mermin $n = 7$	4/9	9/16	1/4	LP-EXACT
Mermin $n = 8$	4/9	9/16	1/4	LP-EXACT (sym-reduced, two-machine verified)
Mermin $n = 9$	8/17	17/32	1/4	

Scenario	$F_{\min}$	s (max satisfiability)	$F \cdot s$	Status
				LP-EXACT (sym-reduced, two-machine verified) — first $R = 16$ point
Mermin $n = 10-13$	8/17, 16/33, 16/33, 32/65	17/32, 33/64, 33/64, 65/128	1/4	LP-EXACT (sym-reduced, two-machine verified) — staircase exact through thirteen parties
C4 linear cluster, AVN core	1/3	3/4	1/4	LP-EXACT — out-of-sample confirmation (§8.3)
C5 ring cluster, AVN core	1/3	3/4	1/4	LP-EXACT — third entanglement class
C4 core + 2 compatible stabilizers (padding)	1/3	5/6	5/18	LP-EXACT — scope counterexample : padding raises s , leaves F fixed (§9)
C6 ring, minimal 5-element core	1/4	4/5	1/5	LP-EXACT — discriminator : refutes 1/(4s) (5/16), confirms Theorem 5 (1/4); ansatz optimal

All floors are for the stated finite setting sets; richer scenarios containing them can only raise the values.

6. Hand proofs at $n = 3$ and $n = 4$ [PROVEN]

Theorem 2 ($n = 3$ floor = $1/3$). The four Mermin parities multiply to $+1$ identically while the targets multiply to -1 ; hence every strategy satisfies exactly three or one of the conditions. Upper bound: classes C_i (satisfy all but condition i) with weight $1/3$ each on the three admissible classes per context give every pairwise fraction exactly $1/3$, with all marginals vanishing by the fiber argument (this construction coincides, after fiber-uniformization, with Hall’s Table I model for the perfect correlators [11]; the lift is what supplies faithfulness at no cost). Lower bound: each λ supports at most three of the four densities; summing pairwise overlaps over the six pairs yields Σ overlaps ≤ 4 , so some pair overlaps $\leq 2/3$, i.e., some fraction $\geq 1/3$. ■

Theorem 3 ($n = 4$ floor = $1/3$). $GF(2)$ relations with odd target-sums (e.g., $p_{(12)} + p_{(13)} + p_{(14)} + p_{(1234)} \equiv 0$ with target-sum -1) forbid joint satisfaction; optimal strategies satisfy exactly six of eight conditions, with missed pairs partitioning the settings into four complementary pairs; settings sharing a missed pair may share distributions, and settings across pairs reproduce the three-of-four combinatorics of Theorem 2, giving $1/3$ by the identical counting bound and construction. ■

7. The universal ceiling [PROVEN]

Theorem 4. $F_{\min}(n) \leq 1/2$ for every $n \geq 3$. *Proof:* each setting’s satisfying set is an affine coset of a hyperplane in $GF(2)^{(n+1)}$ (density $1/2$); distinct settings have distinct functionals, so any two satisfying sets intersect in density exactly $1/4$; letting each p_s be uniform on its satisfying set makes every pairwise fraction exactly $1/2$. ■ Lawful measurement dependence therefore has a *bounded price schedule*: the statistics never force it toward total superdeterminism — a model paying exactly one half exists at every n , so the conspiracy limit is never required (though, as with any bound on necessity, never excluded). Satisfying one

parity constraint only ever confines the hidden state to half the space, and halves of halves always share a quarter.

8. The optimal densities: landscape, inheritance, and glassiness

8.1 Contextual ground states [LP-EXACT $n = 3-7$]. Define the frustration Hamiltonian $H(\lambda) :=$ number of Mermin constraints violated. The GHZ contradiction says H has no zero-energy state. The **minimal-frustration ansatz** — ρ_s uniform over ground states of H within the sector satisfying s — achieves the exact optimum at every computed size and evaluates to $R/(2(R+1))$ at least through $n = 11$ (ground-shell miss counts 1, 2, 6, 12, 28, 56, 120, 240, 496 = the fraction $(R-1)/(2R)$ of settings).

8.2 Stabilizer reading and the required spectrum [PROVEN]. The Mermin constraints are classical shadows of the GHZ state's stabilizers; H is the classicalized stabilizer Hamiltonian, and its frustration is the GHZ contradiction. Fourier analysis over $\text{GF}(2)^{(n+1)}$: H has nonzero coefficients **only on the stabilizer characters, each of magnitude exactly $1/2$** (degrees 1, 3, ..., $n+1$). Corollary (**pairwise no-go**): the residual above degree 2 is nonzero at every n — *no pairwise interaction can generate the landscape*. Equivalently, $H_{\text{required}} = \text{const} - \frac{1}{2} \sum_S t_S \chi_S$: the classical shadow of the very Hamiltonian whose unique quantum ground state is the prepared state. **Inheritance principle [CONJECTURE]:** the hidden variables inherit the energy landscape of the state that wrote them.

8.3 Out-of-sample confirmation: the C4 cluster [LP-EXACT]. The 4-qubit linear cluster state (not local-Clifford equivalent to GHZ_4) has minimal all-versus-nothing core $\{\text{IZXZ}, \text{IZYY}, \text{ZYYZ}, \text{ZXYX}\}$ with signs $(+, +, +, -)$. The measurement-dependence LP with full marginal constraints gives $F = 1/3$ exactly; the frustration ansatz is feasible and optimal; $s = 3/4$ and $F \cdot s = 1/4$. The inheritance principle and the $1/4$ invariant hold beyond the GHZ family. The C5 ring cluster (a third entanglement class) confirms the same values on its minimal core. Upgraded and *scoped* conjecture: **for irreducible stabilizer AVN scenarios — sets in which every constraint participates in the contradiction — $F_{\min} = 1/(4s)$** . The scoping is forced by an explicit counterexample we constructed deliberately: padding the C4 core with two compatible stabilizers raises s to $5/6$ while the floor stays at the core's $1/3$ (the padding constraints are satisfiable by the optimal densities at no cost), so $F \cdot s = 5/18$ there. General scenarios conjecturally inherit their floor from the tightest embedded irreducible sub-scenario. Every minimal core found to date has exactly four elements with three-of-four structure ($s = 3/4$, $F = 1/3$); the full Mermin sets are irreducible scenarios realizing the other rungs of the staircase.

Theorem 5 (general lower bound) [PROVEN]. For any subset S' of deterministic-target constraints with maximum satisfiable count $m(S')$: $F_{\min} \geq (|S'| - m(S'))/(|S'| - 1)$. Hence $F_{\min} \geq B := \max_{S'} (|S'| - m(S'))/(|S'| - 1)$. *Proof sketch:* every λ supports at most m' of S' 's densities; for values $v_1 \leq \dots \leq v_k$, $\sum_{\text{pairs}} \min(v_a, v_b) \leq ((k-1)/2) \sum v_a$ by Chebyshev's sum inequality (centered coefficients decrease while values increase); summing over λ bounds the total pairwise overlap by $|S'|(m'-1)/2$, so some pair overlaps at most $(m'-1)/(|S'|-1)$. ■

Status of the exact law. Theorem 5's bound is *exhaustively tight* on Mermin-3, 4, 5 ($B = F$ over all subsets, verified by complete enumeration) and reproduces the entire staircase on the full Mermin sets: $(16-10)/15 = 2/5$, $(64-36)/63 = 4/9$, and $(N-m)/(N-1) = R/(2(R+1))$ identically on that family. It is not tight universally: deleting a single constraint from Mermin-5 leaves the floor at $2/5$ (computed) while B falls to $4/11$ — so the exact floor lies, in general, between Theorem 5's proven lower bound and the frustration-ansatz upper bound, which

coincide on every homogeneous family and core computed to date. Closing that gap — or characterizing exactly where the ansatz is optimal — is the sharpest open problem this work poses, with LP dual certificates available at every computed size as raw material.

The discriminating scenario [LP-EXACT]. The C6 ring cluster’s stabilizer group contains 1,920 *minimal five-element* contradictions (and none of size three — extending the observed impossibility of triple contradictions). A minimal 5-core has $s = 4/5$ by minimality, so the two candidate formulas diverge for the first time: $1/(4s) = 5/16$ versus Theorem 5’s $(5-4)/4 = 1/4$. The faithful LP returns **exactly $1/4$** , with the frustration ansatz feasible and optimal. Verdict: the $F \cdot s = 1/4$ invariant is special to the Mermin family — where $N(N-1) = 4m(N-m)$ makes the two formulas coincide — and the surviving general candidates are: Theorem 5’s bound as the exact floor on cores and homogeneous families (all instances tight), and the frustration ansatz as the exact optimum everywhere tested without exception (nine scenarios, three entanglement classes, cores of two sizes).

8.4 The landscape is glassy [PROVEN by exhaustion at $n = 4$; sampled at $n = 6$]. No bit-hierarchy ultrametric (0 of 120 orderings at $n = 4$, exhaustive) grades the violation shells; even Hamming distance fails, because maximal-violation configurations lie at distance one from the ground set. Energy here is spectral, not proximity: any supply story of the form “deviation from lawfulness, graded by closeness” is structurally excluded, and relaxation dynamics on this landscape require collective (≥ 2 -site) moves (§11.3).

9. The staircase law [CONJECTURE — exact fit to all computed points]

With $R(n) := 2^{\lfloor (n-1)/2 \rfloor}$ (the Mermin violation ratio) and $s(n) := (R+1)/(2R)$ (the classical bound, re-expressed; matching all hand-computed class counts):

$F_{\min}(n) = R/(2(R+1)) = 1/(4 \cdot s(n))$, equivalently $F \cdot s = 1/4$ exactly.

Verified: $n = 3$ –13 (eleven exact points, every pairing and every step, including the first $R = 8$ and $R = 16$ values) and the C4 core; geometric convergence to the ceiling $1/2$, never attained. Proof program: upper bound by generalizing the ground-state overlap counting of Theorems 2–3; lower bound from the LP dual certificates (available from the solver output at every computed n).

Part II — The race

10. One answer key, every mechanism steelmanned

Rubric (Part II grading): A. mechanism named? B. category per Waegell-McQueen (fine-tuned initial conditions / fluke / nomic-lawful)? C. quantitative against the price list? D. parity scope actually achieved? E. divergent predictions? F. published wounds engaged?

Answer key (Part II-B, universal requirements): (K1) context-conditional densities on satisfying sectors, with a principled reason context enters; (K2) the frustration landscape *generated*, not postulated; (K3) uniform context weights; (K4) budget curve and ceiling respected; (K5) the reach problem solved — every mechanism has one (mediator range, initial-condition correlation, or boundary data); (K6) zero dependence wherever quantum mechanics exhibits none. We record a general result on K6: the inheritance principle delivers it automatically for unentangled stabilizer preparations (jointly satisfiable stabilizers = zero

frustration floor = zero required dependence), so no landscape-adopting mechanism generically contaminates unentangled physics. And a general filter on K5, from experiment: loophole-free Bell tests switch settings faster than light can travel from the source, so **any retarded dynamical mediation of settings-dependence is already excluded by published data** — every surviving mechanism must supply dependence as standing structure, initial condition, or boundary data (Appendix B).

Invariant Set Theory (Palmer), steelmanned. Payment concept: frustration as distance from the invariant set. Priced here: no bit-hierarchy ultrametric — nor Hamming proximity — grades the shells (§8.4), so the payment cannot come from generic off-set distance and must be built from fractal structure tuned to the stabilizer characters. K5 free (state-space geometry needs no mediator). Gift, offered sincerely: the exact floors are rational — $1/3$, $2/5$, $4/9$ — the most Rational-Quantum-Mechanics-consonant quantitative structure we know of.

Cellular Automaton Interpretation ('t Hooft), steelmanned — half-paid here. The reduced class space with the frustration energy is a finite deterministic ontic state space beneath GHZ statistics — the most CAI-native object this analysis produces. Executed on the program's behalf: a local stochastic automaton (Metropolis, 1- and 2-site flips) relaxes to uniform contextual ground shells (spreads ≤ 0.003) and *generates* $F \approx 0.337$ versus the exact $1/3$ **[TOY]**. Finding: single-site updates are non-ergodic on constraint sectors (the glassy landscape freezes them) — the automaton requires ≥ 2 -site rules. Owed: the deterministic version, and K5 (local update rules push settings-correlation into initial conditions, the category the program must escape).

Future-input-dependent dynamics (Hossenfelder-Palmer [6]; Donadi-Hossenfelder [18]), steelmanned. K1 axiomatic (settings are inputs — the bullet bitten openly). Mapping executed here: their hidden variable $\leftrightarrow$ our λ ; their detector-dependent attractor flow $\leftrightarrow$ contextual $T \rightarrow 0$ relaxation; their Born-rule requirement $\leftrightarrow$ our faithfulness constraints. The relaxation half of K2 is genuinely in hand; owed: the landscape half, K3, and multipartite scale.

All-at-once / retrocausal (Wharton; Adlam; Ridley-Adlam), steelmanned — the strongest structural rival. Frictionless translation: histories weighted by an action = the frustration count, past boundary = the preparation's stabilizer conditions (the emission phase is boundary data, which this program alone treats as primitive), future boundary = the realized context; $T \rightarrow 0$ gives the optimal densities exactly. K1, K5 structural; K6 from the landscape. Owed: why *this* action. Full construction: Appendix A.

Scorecard.

Mechanism	K1	K2 landscape	K3	K4	K5 reach	K6
DST condensate	paid (U(1))	derived in-model: inheritance theorem — writable channels = state's stabilizer group; vertex toy-exact at all computed n	paid (phase-blind rings, sharpened to Lemma B's identity)	consistent + continuum PASS + thermal axis computed: ceiling = hot limit, LP floor = ground state	emission-imprinting (Appendix B); massive channel dead; drive records at the WAY/ Ozawa floor ($\eta^2 = \pi^2/64N_d$)	delivered for stabilizer preparations (theorem); open beyond
All-at-once	structural		owed	untested	free (block)	

Mechanism	K1	K2 landscape	K3	K4	K5 reach	K6
		most natural reading				partial (via inheritance)
CAI	plausible	half-paid here (automaton)	half-paid (uniformity emerges)	untested	initial-condition wound	partial
Future-input	axiomatic	relaxation half in hand	owed	untested	inherited	partial
IST	plausible	cheap route closed here ; fractal route open	owed	untested	free (no mediator)	partial

The table is the thesis without adjectives: one mechanism has now paid its principal in-model — installments proven, exposures bounded against known physics, its one free scale quarantined — and carries every scar on the public record; the others have been handed their strongest openings by an adversarial collaboration. The lead is one any contender can erase by computing: the answer key is an entry bar, not a finish line, and §11.5's remaining test names the place our own mechanism is most likely to die.

Part III — The condensate mechanism

11. Supply from the displacement condensate

11.1 The standing-reference formulation. The condensate's broken $U(1)$ furnishes at every point a phase reference — the structure against which rotational displacements have definite phase at all. A Bell experiment samples this one reference twice: the pair's λ acquires its value against it at emission; each apparatus's quadrature choice (X versus Y is a homodyne phase) is defined against it at detection. The correlation is not an influence propagated between events but a shared dependence on a common reference — the repaired common cause: what Bell's theorem excludes is a common cause acting on outcomes alone; what the flooding postulate asserts is a common reference legible to outcomes *and* settings, at exactly the priced budget, bounded by Theorem 4.

11.2 Installments paid. *K1*: the condensate $U(1)$ forbids system-only couplings to collective phase sums; only system-plus-apparatus couplings are invariant, so a context-free energy is symmetry-forbidden — measurement dependence as a consequence of no-external-observer **[PROVEN, given the phase encoding]**. *K3*: integrating out the radial field (quadratic coupling = $\text{Tr} \ln \Rightarrow$ ring tower) yields coefficients depending only on amplitudes and geometry; Mermin contexts differ purely in quadrature phase; hence uniform weights are forced **[PROVEN, structural]**. *K2 (toy)*: the minimal invariant coupling built from party-local vertices — $F(\theta_i, \varphi_i) = e^{(-i\varphi_i) \cdot \sqrt{2} \cos(\theta_i - \varphi_i)}$, the published hemispheric projection times the apparatus reference's own phase — gives $E = -\sum_s \text{Re} \prod_i F_i = 2H(\lambda) - 2^{n-1}$ **exactly, at every computed $n = 3-7$ including $n = 6$, machine-verified to 10^{-14} [TOY]**. The Mermin target signs emerge as $\cos(\Sigma \varphi_i)$ — products of local apparatus phases, never inserted — and no emission-phase parameter exists in the construction. This supersedes an earlier first-harmonic collective-branch analysis whose apparent cancellation at $n \equiv 2 \pmod{4}$ is thereby revealed as a truncation artifact: the full complex vertex product contains all interference branches, and their sum is the frustration Hamiltonian identically. The complex factor is

physically transparent — the apparatus field’s own phase, carried automatically by any complex ring product and discarded by any real pairwise truncation, which is exactly why pairwise mechanisms fail the Fourier no-go (§8.2) while this vertex passes it. (Corollary retained: the Mermin frustration count is an affine function of the collective-phase cosine.) *Continuum consistency*: the hemispheric supply curve $F(\varphi) = \varphi/\pi$ funds the pairwise demands of Hall’s optimal model at every tested angle (worst differential utilization 0.858, at the degenerate geometry; worst finite 0.534) **[TOY, machine-verified]**.

11.3 Wounds, filed. (i) *Range*: the massive channel dies across laboratories — $M = \sqrt{g \cdot v}$ at any keV scale gives $e^{(-ML)}$ with $ML \sim 10^{11}-10^{12}$; a reaching mediator would need $M \lesssim 10^{-7}$ eV, i.e., the massless sector — and the timing filter (§10, Appendix B) then excludes *retarded* massless mediation too. Resolution: the mechanism survives only in emission-imprinting form against the standing reference (11.1), with the Goldstone vacuum as candidate carrier; the flooding postulate is thereby promoted from interpretation to load-bearing dynamics. (ii) *Bookkeeping*: a g -normalization ambiguity between companion papers (dark-matter mass 511 versus ~ 44 keV depending on which g is meant; the proposed bare/dressed reconciliation is unproven) — flagged wherever the papers are co-cited, with the following standing sentence: *where this paper is co-cited with the DST companion papers, note that the proposed resolution — neutral φ_r excitations carry no $U(1)$ charge, undergo no vacuum-polarization dressing, and sit at bare $g = 1$ ($m_{DM} = 511$ keV), while the dressed $g = \alpha$ applies only to charged-sector excitations — is a structural proposal, not yet derived; no result in the present paper depends on that normalization, since the supply mechanism enters only through $g_0 = 3/8$, the vertex structure, and the ratio $\beta = J/T_{eff}$.* (iii) *Underscored*: the magnitude of the ring coefficient — its term content (which characters the effective energy may contain, with what selection) is now derived in-model (11.5); the magnitude sets T_{eff} and the near-optimum signature. (iv) *Dynamics*: the glassy landscape (§8.4) requires the condensate’s relaxation to support collective moves — the same constraint discovered for the automaton.

11.4 The ring derivation: status. The required coupling factorizes as $E_s = -J \cdot \cos(\sum_i \varphi_i - n\theta_c) \cdot \prod_i \cos(\theta_i - \varphi_i)$, and its derivation from the ring expansion decomposes into three parts, two now proven:

Lemma A (equal branch weights) [PROVEN]. The excitations’ phases are global properties, not functions of position, so they factor out of every spatial ring integral: the n -ring linking one system-apparatus interference bilinear per detector yields exactly $J_0 \cdot \prod_i \cos(\theta_i - \varphi_i)$, with J_0 purely geometric. The product of cosines is the equal-weight sum over all 2^n interference branches — the structure the Fourier spectrum demands is automatic, not engineered.

Lemma B (context-uniform coefficient) [PROVEN]. J_0 contains no phases; it cannot know the context. (This sharpens the earlier K3 installment to an identity.)

Step C (collective modulation) — reduced to one expectation value. Dressing the ring with one condensate-apparatus bilinear per site generates orientation branches $\cos(\sum_i \varepsilon_i \varphi_i - (\sum \varepsilon') \theta_c)$; a GHZ preparation deposits an n -quantum phase coherence into the standing structure, so $\langle e^{(ik\theta_c)} \rangle$ over the imprinted condensate vanishes except at $k = \pm n$ — annihilating every mixed branch and selecting $\cos(\sum \varphi_i - n\theta_c)$ uniquely. The flooding postulate thereby acquires its concrete mechanism: *settings-legibility is the coupling of the apparatus array’s collective phase to the winding- n coherence left by emission*. The construction predicts the emission phase locks to the condensate reference ($\Theta_{em} = 0$), in agreement with the toy of 11.2, which required no phase parameter. *Both the selection rule and the expectation value are now derived: §11.5.*

Residual audit — executed [machine-verified where stated]. Residual couplings sort into three classes. **Class 1** (per-context constants: apparatus-only and amplitude terms): harmless by construction — a constant within a sector never moves its argmin. **Class 2** (condensate-system terms of winding $k \notin \{0, \pm n\}$): annihilated by the Step-C selection rule — the imprint populates no such channel. **Class 3** (winding-zero system-system couplings, $\cos(\theta_i - \theta_j)$): *lethal if present* — perturbing the draw energy by $\varepsilon \cdot \sum \cos(\theta_i - \theta_j)$ at $|\varepsilon| = 0.01$ already breaks faithfulness (proper-subset marginals cease to vanish) and inflates the floor (to $3/7$, $1/2$, or worse at $n = 3, 4$); a symmetry-orbit defense fails (the scenario symmetry group is verified non-transitive on the ground shells). Class 3 is excluded exactly when the selection rule generalizes to the system sector: **the effective energy contains only terms coupling to coherence channels the imprint populates** — and a GHZ emission populates the collective winding- $\pm n$ channel and no pairwise channel. This is the inheritance principle in its sharpest form: the hidden variables inherit the prepared state's landscape *because the writable channels are the state's coherence channels*. Under it, the payload survives, Classes 2 and 3 vanish for one reason, and the entire remaining bill consolidates into a single derivation: *show the emission model's effective energy inherits the imprint's coherence support* (which contains the former item (1) as its winding- n component). Fallback if the rule is only approximate: a hierarchy requirement $J_2 \ll T_{\text{eff}} \ll J$, with deviations from exact quantum statistics scaling as J_2/T_{eff} — the mechanism's own quantitative, falsifiable exposure, recorded as such. A proof that the imprint's coherence support does not control the energy's term content kills the supply story cleanly, as a good mechanism should die. *That derivation is executed next.*

11.5 The inheritance theorem: the bill, paid in-model [PROVEN in-model; matrix elements machine-verified]. The consolidated derivation of 11.4 — *show the emission model's effective energy inherits the imprint's coherence support* — is a theorem in an explicit emission model, with two honest exposures, both quantified.

The model and its hypotheses. n system modes (the emitted excitations, ontic phases θ_i), one condensate collective mode (the standing reference), n macroscopic apparatus modes (homodyne phases ϕ_i). **(H1)** Exact $U(1)$ conservation with no external phase reference: the condensate phase is relational. This is not a new assumption — it is K1's no-external-observer principle working a second shift. **(H2)** The preparation is a collective emission: the only phase-sensitive system-source coupling is the n -quantum vertex $(\prod_i a_i^\dagger)c^n + \text{h.c.}$

Step 1 — the record is superselected. Under H1 the source cannot hold a phase-definite coherent state (a coherent condensate is an eigenstate of c , takes no imprint, and would populate every winding channel — H1 is where the selection lives or dies). A GHZ emission leaves the joint record $|\Psi\rangle = (|0\dots 0\rangle|N\rangle + |1\dots 1\rangle|N-n\rangle)/\sqrt{2}$, and every question about which λ -characters the standing structure can energetically distinguish becomes matrix-element algebra in this state.

Step 2 — the dangling-winding lemma. Every term the ring expansion can write couples a λ -character $e^{i(\sum \varepsilon_i \theta_i)}$ to environment operators transferring $\varepsilon_i \in \{0, \pm 1\}$ units of system winding at site i and q units of condensate winding. Sweeping **all** $3^n \cdot (2n+5)$ channel assignments at $n = 3, 4$ **[machine-verified]**: the record's expectation is nonzero for exactly three — ($\varepsilon = 0, q = 0$), the diagonal channel (per-context constants, Class 1), and ($\varepsilon = \pm 1^n, q = \pm n$), the full collective flip. Any term flipping a proper subset of sites maps both branches to states orthogonal to both — Classes 2 and 3 vanish *identically, to all orders in the imprint channel*, for a single reason: no partial-winding coherence exists to pay them. The audit's lethal $\cos(\theta_i - \theta_j)$ is a $(+1_i, -1_j)$ channel; its record amplitude is exactly zero.

Step 3 — the amplitude (the former principal). The winding- n coherence of the record is **1/2 exactly, independent of N [machine-verified, $N = 6-40$]**. The one expectation value the bill compressed to is computed: maximal for an equal-branch record, $O(1)$, unsuppressed. Two corollaries fall out. $\Theta_{em} = 0$ is no longer a prediction but an identity — relational phases admit no emission-phase parameter, which is what the toy of 11.2 discovered numerically when it needed none. And the first-order source term, $-J_{src} \cdot \cos(\Sigma \theta_i - n \theta_c)$, is context-independent — common to every ρ_s , it is the λ -level record of the GHZ coherence itself, absorbed into the prior; its consistency with faithfulness is what the toy’s exactness already verified.

Step 4 — multi-site connectivity. Classical linkage cannot fake connectivity: rings “connected” through the order parameter (a c -number) factorize, and disconnected pieces contribute additively to the energy — site-local terms only, a free-energy identity. Genuine multi-site terms require quantum lines: the massive radial channel between detectors (dead, $e^{(-ML)}$, $ML \sim 10^{11}$ — the range wound of 11.3(i), now load-bearing as a defense) or the imprint, which Step 2 restricts to the collective channel. The payload $\cos(\Sigma \phi_i - n \theta_c) \cdot \prod_i \cos(\theta_i - \phi_i)$ is the *unique* closure: converting the source’s $e^{(i \Sigma \theta_i)}$ into apparatus-referenced phases requires every site’s local $(\theta_i - \phi_i)$ rotation, and a missing site leaves a dangling single-site winding, which Step 2 prices at zero. Full weight is not selected — it is forced.

Step 5 — generalization to the stabilizer family. For any stabilizer state, the set of Pauli operators with nonzero expectation is exactly its stabilizer group **[machine-verified: GHZ $n = 3-5$, whose X/Y-containing support is purely full-weight with every pairwise XX/XY/YY channel zero; C4-chain and C5-ring clusters, whose supports equal their own stabilizer groups, mixed weights included]**. The record of a stabilizer preparation therefore populates precisely the characters from which Part I’s Fourier theorem builds the required Hamiltonian — writable = required, scenario by scenario, including the cluster scenarios that confirmed the inheritance principle out-of-sample in §8. The principle is now a theorem of the emission model rather than a name for a pattern, and Part I and Part III meet in one sentence approached from opposite directions: *the hidden variables inherit the prepared state’s landscape because the writable channels are the state’s coherence channels.*

Exposures, quantified. (E1) H2 beyond single-vertex emission — narrowed: (i) any gate acting as a unitary on the system alone transforms the record covariantly, $|\Psi\rangle \rightarrow (U \otimes 1)|\Psi\rangle$, so record support = state support is preserved *exactly*, at every step of the circuit. Graph-state preparations — per-site emissions (whose records populate the single-site channels the product state’s own stabilizer group contains) followed by number-diagonal CZ gates (which commute with every winding operator and exchange nothing with the reference) — are therefore protected by gate structure, and every stabilizer scenario is local-Clifford equivalent to a graph state with the local rotations folding into the measurement settings: **H2 is discharged for the entire stabilizer family at the level of ideal gate structure.** (ii) The residual is winding bookkeeping — *now computed in-model*: $U(1)$ conservation forces single-site rotations to exchange winding with drive fields. Modeling the drive under H1 itself (a superselected laser: a Poissonian number mixture, no absolute phase) with the number-conserving exchange $U: |0, n\rangle \rightarrow \cos(gt/\sqrt{n})|0, n\rangle - i \sin(gt/\sqrt{n})|1, n-1\rangle$, the $\sqrt{n}$ Rabi dependence is the physical origin of the record: number branches rotate by slightly different angles, leaving residual drive-system entanglement. The relational winding-coherence deficit per $\pi/2$ rotation is **$\eta^2 = \pi^2/(64 N_d) + O(N_d^{-2})$ [machine-verified: $\eta^2 \cdot N_d \rightarrow 0.154213 = \pi^2/64$; log-log slope -1.0000 over $N_d = 25-10^5$]**, deficits adding over gates. This is the Wigner-Araki-Yanase mechanism in the present language — conservation laws bound the accuracy of any

operation that fails to commute with the conserved charge, with the same inverse-variance scaling (Ozawa [17]) — so the residual record is not an artifact of this mechanism but a floor that standard quantum mechanics already imposes on every gate: it is small for exactly the reason quantum gates work at all. The kill condition is now quantitative: the mechanism dies here only if $T_{\text{eff}}/J < \eta \sim 1/\sqrt{N_d}$ — impossible for macroscopic drives unless T_{eff}/J is itself extreme, tying this exposure's fate to the one remaining number, J_0 . The fallback is unchanged: $J_2 \sim \eta J$, deviations $\sim J_2/T_{\text{eff}}$. **(E2)** Channels outside the record entirely: static Goldstone exchange between localized, globally neutral wavepackets is derivative-coupled and multipole-suppressed, $J_2/J \lesssim (d/L)^p$ with $p \geq 3$ — at packet scale $d \sim \mu\text{m}$ and detector separation $L \sim \text{m}$, $\lesssim 10^{-18}$ against the audit's machine-verified lethality threshold of 10^{-2} — sixteen orders of margin. For emission-type preparations the hierarchy $J_2 \ll T_{\text{eff}} \ll J$ is thereby *derived*, not assumed. Second-cumulant contributions are products of populated channels — doubled characters and apparatus constants, Class 1 on the hemispheric sector variables.

The last number, quarantined. J_0 — the magnitude of the geometric ring coefficient — is not symmetry-forced: it is an overlap integral over wavepacket and condensate geometry, and we declare it the mechanism's single free scale, entering all λ -statistics through one dimensionless ratio, $\beta = J/T_{\text{eff}}$. Its empirical reach is then sharply bounded by a computation **[machine-verified, n = 3-5]**: in the context-hard Gibbs family $\rho_s \propto e^{(-\beta H) \cdot [\text{sat}_s]}$ on the reduced classes, the full correlators equal their targets *exactly* and every proper-subset marginal vanishes *at every* β — the free scale is invisible in every stabilizer experiment, precisely the experiments Part I prices. Only the hidden dependence moves: $F(\beta)$ runs monotonically from **exactly 1/2 at $\beta = 0$ to the LP floor at $\beta \rightarrow \infty$** , with cold-side approach $e^{(-R\beta)}$: the satisfying-sector spectral gap of the frustration count equals the Mermin ratio R *exactly, at every n* **[PROVEN]** — a corollary of the same product identity that yields the classical bound ($W = \sum_s t_s \cdot \text{sign}_s = \text{Re} \prod_j (x_j + iy_j)$, modulus $2^{(n/2)}$, phase quantized in $\pi/4$; Mermin [15], repurposed), verified exhaustively over all 4^n assignments for $n \leq 9$, with the closed-form grounds recovering every known satisfiability $s = 1/2 + W_{\text{max}}/2^n$. (Scoping kept honest: the *gap* is proven; that the fitted $F(\beta)$ decay rate equals it requires the leading-order coefficient to be nonzero — machine-verified at $n = 3-5$, generic but not separately proven.) The flooding postulate thereby acquires a temperature axis: the ceiling value is the hot limit of the supply, the LP optimum its ground state, and the v1 note's principle is now a computed interpolation rather than a posit. The gap law is the second place R appears for a provable reason — a thermal echo of the staircase. (One curiosity remains filed, now with a sharper target: the $F(\beta)$ curves for $n = 3$ and $n = 4$ coincide digit-for-digit at every temperature despite different level *counts* — two levels versus three at the same gap — so the conspiracy, if it is one, lives in the multiplicities; checkable.)

What remains — the continuum finite- β test, scoped against ourselves. Stabilizer targets are extremal (± 1), so the cold limit is natural and overshoot impossible; non-extremal targets force finite β , and the quarantine ends exactly where extremality does. The simplest realized-context family on the $n = 2$ reduced classes is closed-form **[machine-verified]**: $E(s; \beta) = \tau_s \tanh \beta$ with vanishing marginals, hence $S = 4 \tanh \beta$ and $F(\beta) = \frac{1}{2} \tanh \beta$ — local at $\beta = 0$, **Tsirelson at $\beta^* = \text{arctanh}(1/\sqrt{2}) \approx 0.8814$** , PR-box as $\beta \rightarrow \infty$, paying $F(\beta^*) = 1/(2\sqrt{2}) \approx 0.354$ against Hall's floor of 0.138 (valid, $2.6\times$ wasteful). The stakes, stated as sharply as we can against our own candidate: a generic free β passes *through* the quantum bound, and $S > 2\sqrt{2}$ is excluded at $\sim 30\sigma$ (Poh et al. [16]: $S = 2.82759 \pm 0.00051$). The mechanism therefore owes exactly one of two derivations: that the ring-generated continuum functional is *cos-saturating* — correlators bounded by the quantum envelope with Tsirelson as its extremal value,

consistent with the companion framework in which the condensate geometry fixes Tsirelson exactly and corrections land in accounting relations rather than laboratory statistics — or a dynamical pinning $\beta \leq \beta^*$. Failing both, precision data closes a two-sided vise: visibility budgets bound β from below, Tsirelson from above, and the window can shut. This is the mechanism’s most dangerous surviving test; we schedule it as such, and note that the wasteful F of the naive family suggests the true functional — if it exists — is simultaneously safer and cheaper. Beyond it, the ledger holds only bookkeeping (the g -normalization sentence co-cited across companion papers).

12. Discussion

Operational corollary. $F_{\min}(n)$ is the counterfeiting threshold for multipartite device-independent certificates: an adversary — or a physics — commanding that fraction of settings-dependence simulates the n -party violation classically and locally. The closed form says multipartite certificates never cost the counterfeiter more than half of measurement independence, a security-relevant statement independent of interpretation.

The scoring rule for empirically tied rivals. Inside a system one cannot exit, every account of the hidden layer rests on at least one unmeasurable inference — irreducible chance, the invariant set, the block, the flooded condensate. Nobody plays inference-free; the legitimate tiebreaker is *smallest inference, largest audited coverage* (the criterion by which special relativity displaced the empirically identical Lorentz ether theory). Part I is the audited coverage; Part II prices the inferences; Part III’s postulate is one inference with structural obligations, currently the most audited and the most wounded — which is what being tested looks like.

Why the mechanism grants no simulation superpower — and why that is a consistency success. A natural hope: if each particle’s state is defined at emission, simulating n qubits should cost n vectors, not 2^n amplitudes. The inheritance theorem (11.5) is the no-go: the hidden variables are cheap, but the *density* over them is the $T \rightarrow 0$ Gibbs state of an energy whose term content equals the prepared state’s coherence support — polynomial for stabilizer states (2^n signs, n generators), generically exponential ($\approx 4^n$ independent Pauli expectations) otherwise. The exponential does not vanish; it relocates into the standing structure’s record. The alignment is exact and uninstructed: where the record compresses (the stabilizer fragment) classical simulation is already easy — Gottesman-Knill [21], of which our reduction theorem (§4) is the hidden-variable shadow — and where the record explodes (magic-state resources) classical cost is known to grow exponentially (stabilizer-rank bounds, Bravyi-Gosset [22]). A mechanism that compressed general quantum simulation would collapse BQP classically and be refuted by published sampling-supremacy experiments; a mechanism whose cost profile independently reproduces the known easy/hard boundary of quantum computation is instead passing a consistency test it was never shown the answers to.

Outlook. The staircase now stands exact through $n = 13$; the symmetry-reduced program makes further steps commodity computations; further graph states stress the $1/4$ law toward theoremhood; the LP dual certificates open the lower-bound half of its proof; and the answer key stands as an entry bar for any future mechanism — including ours.

Appendix A — The all-at-once formulation (fair-play construction)

Histories = complete counterfactual outcome assignments (the strategy space of §3). Past boundary: the preparation’s stabilizer conditions $\{(S, t_S)\}$ (emission phase = boundary data). Future boundary: the realized context. Action: $S[h|s] = J \times (\text{violated stabilizer conditions})$, context’s condition hard. Rule: $P(h|s) \propto e^{-(S/T)}$, $T \rightarrow 0$ — exactly the optimal densities, floors, staircase, and ceiling of Part I. K1, K5 structural; K6 via the inheritance principle. Owed: why this action. We judge this the most frictionless translation into any rival formalism and therefore the condensate mechanism’s strongest structural rival.

Appendix B — The timing filter and the fork

Loophole-free Bell tests switch settings spacelike-separated from emission; retarded mediation (massive or massless) cannot establish the settings- λ correlation in time. Consequence, general: dynamical-mediation supply stories are excluded by published data for every mechanism; dependence must be supplied as standing structure, initial condition, or boundary data. Consequence, DST-specific: the Goldstone-mediation route survives only in its static limit, which is the standing reference of §11.1 — the fork resolves to emission imprinting, and the flooding postulate becomes dynamical, not decorative.

Code and data availability

Every **[LP-EXACT]** and **[machine-verified]** tag in this paper corresponds to an assertion in a runnable script. The complete verification suite — the fiber-reduced linear program, the symmetry-reduced program of Theorem 1b (which recomputes the entire staircase $n = 3\text{--}13$ in under one minute on commodity hardware), and the mechanism checks of §11.5 (inheritance_check.py, eta_check.py, finite_beta.py, vise_toy.py, delta_R_check.py) — is available at:

https://github.com/RandomInternetPreson/moire-phase-space-sampler/tree/main/DST_Bell_MI/Price_of_Locality_Source_Material

Readers are invited to rerun everything.

References

1. J. S. Bell, “On the Einstein Podolsky Rosen paradox,” *Physics Physique Fizika* 1, 195 (1964).
2. M. J. W. Hall, “Local deterministic model of singlet state correlations based on relaxing measurement independence,” *Phys. Rev. Lett.* 105, 250404 (2010).
3. J. Barrett and N. Gisin, “How much measurement independence is needed to demonstrate nonlocality?” *Phys. Rev. Lett.* 106, 100406 (2011).
4. C. H. Brans, “Bell’s theorem does not eliminate fully causal hidden variables,” *Int. J. Theor. Phys.* 27, 219 (1988).
5. 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, 012121 (2019).
6. S. Hossenfelder and T. N. Palmer, “Rethinking superdeterminism,” *Front. Phys.* 8, 139 (2020).

7. T. N. Palmer, “The invariant set postulate,” *Proc. R. Soc. A* 465, 3165 (2009); arXiv:1507.02117; *Universe* 10, 47 (2024); J. R. Hance, T. N. Palmer, and J. Rarity, *Phys. Scr.* 100, 065123 (2025); T. N. Palmer, *PNAS* (2026).
8. G. ‘t Hooft, *The Cellular Automaton Interpretation of Quantum Mechanics* (Springer, 2016).
9. K. Wharton and N. Argaman, “Colloquium: Bell’s theorem and locally mediated reformulations of quantum mechanics,” *Rev. Mod. Phys.* 92, 021002 (2020); E. Adlam, *Entropy* 20, 41 (2018).
10. M. Waegell and K. McQueen, arXiv:2509.23047; M. Waegell, arXiv:2604.00311.
11. M. J. W. Hall, “Relaxed Bell inequalities and Kochen–Specker theorems,” *Phys. Rev. A* 84, 022102 (2011).
12. M. J. W. Hall and C. Branciard, *Phys. Rev. A* 102, 052228 (2020).
13. S. Hossain, S. Molla, P. Kundu, A. Paul, S. Chattopadhyay, and R. Sarkar, arXiv:2107.12495.
14. A. Roy, S. S. Bhattacharya, A. Mukherjee, and M. Banik, arXiv:1310.8570.
15. N. D. Mermin, “Extreme quantum entanglement in a superposition of macroscopically distinct states,” *Phys. Rev. Lett.* 65, 1838 (1990).
16. H. S. Poh, S. K. Joshi, A. Cerè, A. Cabello, and C. Kurtsiefer, “Approaching Tsirelson’s bound in a photon pair experiment,” *Phys. Rev. Lett.* 115, 180408 (2015).
17. M. Ozawa, “Conservative quantum computing,” *Phys. Rev. Lett.* 89, 057902 (2002).
18. S. Donadi and S. Hossenfelder, “Toy model for local and deterministic wave-function collapse,” *Phys. Rev. A* 105, 022210 (2022).
19. 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).
20. I. Sen and A. Valentini, “Superdeterministic hidden-variables models,” papers I and II, *Proc. R. Soc. A* 476 (2020).
21. D. Gottesman, “The Heisenberg representation of quantum computers,” arXiv:quant-ph/9807006 (1998).
22. S. Bravyi and D. Gosset, “Improved classical simulation of quantum circuits dominated by Clifford gates,” *Phys. Rev. Lett.* 116, 250501 (2016).
23. A. Alai, “Displacement Spacetime,” ai.viXra:2604.0009 (2026).
24. A. Alai, “The DST Mass Gap,” ai.viXra:2604.0056 (2026).
25. A. Alai, ai.viXra:2604.0075 (2026) — superseded by the present work (see Part 0).
26. P. Pal, G. Sharma, and R. Podila, “Loophole-Robust Certification of Quantum Advantage,” arXiv:2607.13090 (2026).