Frame-Aligned Records: Supplementary Information

Calculations, study definitions and reproducibility

Supplementary Information v1.2 | 26 September 2026

James McGaughran

Scope. This supplement supports the methods paper. Scientific calculations and model definitions come first; dated lineage is in the final provenance section. A reported historical run, a fresh same-code replay and a separate numerical implementation have different evidential meanings.

Version crosswalk. Portfolio v1.8 identifies the distribution; *Frame-Aligned Records* v0.5, this supplement v1.2, the results log v1.2, protocol reader edition v1.3 and *Absolute Infinite Union* v1.4 identify its components. CAL-1.6 and similar identifiers name fixed calculation suites, not alternative current manuscripts. A historical label means a preserved prior state, even when its date is only one day earlier. The authoritative executed specification remains protocol v1.2 and its original pre-execution receipt; later reader edits do not modify that record. Finalization retains these version labels and records changes through source hashes and a dated change record.

S1. Numerical provenance for the methods paper

The source run labels S01 to S06 refer to the preserved six-script calibration suite under historical/supporting_lineage/v1_6_code/legacy_v1_5/legacy_claude_v0_2_0/. The original display values and sampling histories are retained, not all rerun during this editorial release. The newer calibration and pilot paths are listed explicitly below.

Statement	Value(s)	Script and output key
Lemma 1 counterexample	$L = 401$	S02 L1_counterexample_published_expression_without_identity
Lemma 3 invariance; generic change	$\leq 6 \times 10^{-17}$; 0.144	S02 L3_*
Proposition 1(c)	1.0	S02 P1_commuting_unitary_changes_pure_ground_state_by
Proposition 1, Remark 2	$1.0, 0.399, 0, 0.399, 1.0$ bits	S02 P1_stationary_average_conceals_conditional_dynamics
Proposition 2	0.5270653410; 0.3400520132; -0.7615941560; 0; $1 \to 0$ bits	S01 P2_*
Lemma 4 example	0.6428571429 (formula and enumeration)	S01 hypergeometric_*
Proposition 3 checks, Table 1	as tabulated	S03 A_haar_fragments, B_dynamics
Table 2	as tabulated	S04 duplicate_inflation
Hoeffding half-widths	$0.200, 0.100, 0.050$	S01 hoeffding_half_width
Supplement Tables S4 to S6	as tabulated	S05 e1, e2, e3, e4 in the exploratory output
Supplement Section S2.5	as stated	S05 replication output
Legacy S06 span checks	0; 7×10^{-16}; 0.4166 versus 0.4170	S06
Mixed twirl, historical v1.5	2,000 draws; mean 0.0557464152; SE 0.0007264476	Archived W03 v1.1, Section 6; not the CAL-1.6 sample
Mixed twirl, CAL-1.6	1,200 draws; mean 0.05550839355; SE 0.00093479233; exact 0.056	historical/supporting_lineage/v1_6_data/results.json, mixed_summary; full draws in mixed_draws.csv
Proposition 5a endpoint constant	$\sqrt{1.0025} = 1.0012492197$ at $n = 2, 4, 6$	historical/supporting_lineage/v1_6_data/localized_constants.csv
Proposition 6 spectator	error difference approximately 8.44×10^{-15} against a loose bound 78.53981634	historical/supporting_lineage/v1_6_data/results.json, spectator
Unequal certificate example	0.0995 and 0.104, same true error 0.099	Section 7.3.1, direct arithmetic; not two physical observations
Proposition 7, Figure 6	capped bound 1 through $n = 18$; log10 bound approximately -56.85 at $n = 19$	data/prop7_thresholds.csv; direct envelope, not state simulation
Source-off feasibility pilot	192 independent units in two size strata; separate outcome definitions	data/pilot_run/results.json, per_unit.json; results log, not an LRA test
Post hoc sign analysis	primary raw $p = 5.4592 \times 10^{-5}$; four-row Holm $p = 1.6377 \times 10^{-4}$	data/review_additions/sign_tests.csv; code/analyze_review_additions.py
Range-only precision budget	32,791 units for half-width 0.015, not a power claim	data/review_additions/review_analysis.json, precision

Statement	Value(s)	Script and output key
Initial eligibility and transient access	5 and 10 eligible site/unit pairs in native arms briefly exceed 0.8	data/review_additions/initial_floor_diagnostics.csv
Source encoding in clean transport control	phase code $d = \|f\|$; vacuum/excitation code $d = \|f\|^2$	data/review_additions/transport_encodings.csv; separate full-space checks

Additional run crosswalk

Target	Source and data	Boundary
Factorial four/six/eight-qubit means	data/reviewer_replay/factorial_results.json; factorial_n*.npz	Same-code reproduction of an exploratory review extension
Larger source-off sample	data/reviewer_replay/reimpl_id1_results.json	Fresh-stream reimplementation, not the original sample
Third-frame means	data/reviewer_replay/factorial_neutral.json	Post hoc comparator, no guaranteed neutrality
Paired range and finite-family intervals	data/factorial_analysis/analysis.json; code/analyze_factorial.py	Derived secondary sampling bounds
Alternative numerical method	verification/factorial_independent_checks.json	Selected trajectories, not a second full population study

S2. Historical continuously driven simulation

The material below retains the complete exploratory model and outcomes for reproducibility. It is not another experiment or a retrospective confirmation of the newer source-off target. References to original Sections 2 to 5 concern the methods framework; references to Section 6.x are mapped to S2.x below. Its reported causal interpretations are descriptive and family-specific, not general theorems.

S2.0 Status

Status: EXPLORATORY. This section reports computed results for one small model class. The configuration was changed after the first results were inspected (Section S2.2), so they are not a confirmatory test of the prospectively stated population conjectures. They can nevertheless supply counterexamples to a universal claim and guide later design; exploratory does not mean scientifically worthless. They are included because they show concretely how the design choices of Sections 2 and 5 decide outcomes. The frozen configuration was then re-run with fresh seeds (Section S2.5).

The notation below uses K_C for the source-coupling operator, reserving C for an alternate frame label. Preserved scripts and original outputs still use their original variable C; this is a notation change only.

S2.1 Model

- **System and source.** $n = 3$ system qubits and a binary source with equal priors; $H_x = H_S + (-1)^x K_\mathrm{C}$.
- **Internal Hamiltonian.** H_S is a random nearest-neighbour 2-local chain: Gaussian coefficients on all two-body Pauli products on adjacent pairs, plus one-body terms at half weight. It is normalised to Frobenius norm λ.
- **Coupling and time window.** The source coupling is the star operator $K_\mathrm{C} = \sum_{j=1}^{3} X_j$, of Frobenius norm $\sqrt{24} \approx 4.90$. Times are $t \in \{0.6, 0.7, 0.8, 0.9, 1.0\}$.
- **Candidates.** There are $N = 20$ candidates: the identity (the frame in which H_S and K_C are local), a preparation frame W, and 18 Haar-random unitaries.
- **Locality.** L is the internal score L^{int} on the span of Pauli strings of weight at most 2, identity included.
- **Records (simplified).** The record score is the time-averaged number of single qubits whose Helstrom error is at most 0.1. This is a simplification of the disjoint, persistent redundancy of Section 2.5.
- **Near-optimal sets.** $\mathcal{N}_L$ requires $L \leq \min L + 0.02$ and $L \leq 0.6$; $\mathcal{N}_R$ requires $R = \max R$ and $R \geq 0.5$.
- **Arms.** In the *shared-frame* arm the initial state is $|000\rangle$. In the *decoupled-preparation* arm it is $W\,|000\rangle$ with W Haar-random, so by Section 4.3 it is Haar distributed. The coupling stays in the identity frame in both arms.

S2.2 Development history

The first configuration used a single-qubit coupling $X \otimes I \otimes I$, $\lambda = 4$ and nine times in $[0.3, 1.5]$. Records qualified in 1% of shared-frame models and in none of the decoupled models ($A = 0.010$ and 0.000). The coupling was then changed to the star operator, the time window was narrowed, and a grid over λ was added. All results below use the changed configuration, which was frozen on 25 September 2026.

S2.3 Shared versus decoupled preparation

Table S4. Association results by arm (200 models per cell; Hoeffding 95% intervals clipped at $1 - 1/N = 0.95$).

λ	Arm	Records qualify	$\bar{A}$	$\bar{q}$	$\hat{d}$	95% interval
0.3	shared	1.000	1.000	0.050	0.950	[0.758, 0.950]
0.3	decoupled	0.095	0.050	0.005	0.045	[-0.147, 0.237]
1.0	shared	1.000	1.000	0.050	0.950	[0.758, 0.950]
1.0	decoupled	0.060	0.025	0.003	0.022	[-0.170, 0.214]
4.0	shared	0.320	0.255	0.016	0.239	[0.047, 0.431]
4.0	decoupled	0.030	0.005	0.002	0.004	[-0.189, 0.196]

In the shared arm, the most local candidate hosts the records in every model at $\lambda \leq 1$, and the descriptive estimate is 0.95 in these sampled cases. This agreement is facilitated by the shared inputs, not proved for every model by a coupling bound. At $\lambda = 4$ the internal dynamics is comparable in norm to the coupling, and records qualify in only 32% of models. In the decoupled arm, records qualify rarely. In the reported cases they preferentially sit in the identity frame, where the coupling remains local. This association is consistent with the supplied interaction structure, not an isolated causal estimate.

S2.4 Moving the preparation frame and the coupling frame

Table S5. Preparation frame rotated by $W = e^{-i\theta G}$ away from the dynamics frame (G a random Hermitian of unit spectral norm; $\lambda = 0.3$; 100 models per row). The columns give the fraction of models in which records qualify anywhere, and in which the record-optimal set contains the dynamics frame, the preparation frame, or neither.

θ	Records anywhere	In dynamics frame	In preparation frame	Elsewhere
0.0	1.00	1.00	1.00	0.00
0.1	1.00	1.00	1.00	0.00
0.2	1.00	1.00	0.97	0.00
0.4	1.00	1.00	0.44	0.00
0.8	0.86	0.82	0.06	0.00
1.6	0.13	0.10	0.00	0.03

At $\theta = 0$ the "preparation frame" candidate equals the identity, a duplicated class of exactly the kind Lemma 3 requires to be merged. As the ready resource is rotated away, records do not follow it. They stay in the frame of the coupling and then disappear, in this specific exploratory family. Finite-angle exponentials of normalised random Hermitians are not Haar draws, and no Haar limit for this parameterisation has been established.

Table S6. Coupling and ready resource moved together ($K_{\mathrm{C}} \mapsto U_c K_{\mathrm{C}} U_c^\dagger$, initial state $U_c\,|000\rangle$, $U_c = e^{-i\varphi G}$; 60 models per row). The internal-locality minimum was at the identity in every model.

λ	φ	Records anywhere	In internal-locality frame	In coupling frame	Elsewhere
0.3	0.4	1.000	0.583	1.000	0.000
0.3	0.8	1.000	0.017	1.000	0.000
1.0	0.4	1.000	0.333	0.917	0.000
1.0	0.8	1.000	0.033	1.000	0.000
4.0	0.4	0.317	0.133	0.183	0.067
4.0	0.8	0.200	0.033	0.100	0.067

The total record-generating Hamiltonian H_x was more local, by L averaged over x, in the coupling frame than in the internal-locality frame in 100%, 100% and 95% of models at $\lambda = 0.3$, 1.0 and 4.0 ($\varphi = 0.8$, 60 models each).

S2.5 Fresh-seed repeats of the frozen exploratory configuration

With the configuration frozen, the main contrasts were re-run with seeds not used during development.

- **Seeds 20260926 and 20260927, $\lambda = 0.3$, 200 models per arm.** The shared arm again gave $\bar{A} = 1.000$, $\bar{q} = 0.050$ and $\hat{d} = 0.950$. The decoupled arm gave $\hat{d} = 0.072$ $[-0.120, 0.264]$ and 0.040 $[-0.153, 0.232]$.
- **Three-frame design, $\lambda = 0.3$, $\varphi = 0.8$, 100 models per seed.** Records qualified in every model. They lay in the coupling frame in every model and in the internal-locality frame in 0% and 1% of models.
- **Total Hamiltonian.** H_x was more local in the coupling frame in 100% of models for both seeds.
- **Larger model-sample decoupled run, still three qubits.** One further run was declared before execution (seed 20260928, 2,952 models, chosen for a Hoeffding half-width of 0.05). It gave records in 10.7% of models, $\bar{A} = 0.0515$, $\bar{q} = 0.0055$ and $\hat{d} = 0.046$ $[-0.004, 0.096]$.

S2.6 What the calculation shows, and what it does not

1. **Shared frames produce enrichment by construction.** The estimate $\hat{d} = 0.95$ reflects the design, not a regularity of the model class.
2. **Records track the frame in which the ready resource and the source coupling are aligned.** A ready resource rotated away from the coupling produces few records, and they do not follow the resource (Table S5).
3. **The verdict depends on which Hamiltonian defines locality.** When the coupling and the resource move together, records move with them, away from the internal-locality frame, and into the frame in which the total Hamiltonian is most local (Table S6 and Section S2.5). Its descriptive pattern differs for internal versus branchwise total locality. Neither reading is a calibrated verdict on LRA-int or LRA-tot, because the exploratory time-averaged record outcome and label reference are not the prospective persistent-record test.
4. **The "decoupled" arm is only partly decoupled.** Its coupling remains local in the dynamics frame. Its small descriptive enrichment ($\hat{d} = 0.046$) is compatible with influence from the shared coupling frame, but the toy does not isolate that cause; and its interval, $U = 0.096$, lies just below a margin of $\Delta = 0.1$. This is not evidence about any conjecture.

The calculation has three qubits, 20 candidates, Haar candidates not demonstrated to form an exhaustive set of equivalence classes, a simplified record score and thresholds chosen during development. The coupling has a larger declared norm than the internal Hamiltonian at $\lambda \leq 1$; that norm comparison alone is not a record-persistence guarantee. Its role is to show that the questions in Section 7 need to be posed with frames declared separately.

S3. Post hoc sign tests and window diagnostics

The four-row directional analysis uses the immutable per-unit contrast data and verifies them against the full trajectories. Its null is Pr(z>0 given z is nonzero)=1/2 for independent signs; it is not a zero-mean null. No ties were observed. Exact binomial p-values and a four-row Holm adjustment are reported in the results log. The Holm family was chosen during review, not in the pre-execution plan. Dependence between the paired arms is not treated as extra sample size.

The script also reports initial-eligibility and transient-access counts by frame, and compares both source encodings in an ideal exchange-chain transport control. The phase-superposition pilot code has singleton distinguishability |f|, while the vacuum-versus-excitation control has |f| squared. These are not interchangeable time-window diagnostics. The original study's thresholds, windows and data remain unchanged.

```
OPENBLAS_NUM_THREADS=1 OMP_NUM_THREADS=1 python code/analyze_review_additions.py --root . --out
↪  /path/to/fresh_review_analysis
```

The code refuses a nonempty output folder. It writes review_analysis.json, sign_tests.csv, transport_window_checks.csv, transport_encodings.csv and initial_floor_diagnostics.csv. The last file counts site/unit pairs, not independent observations. Mathematical derivations, same-code replay, separate-method checks and external replication remain distinct.

S3.1 Initial eligibility and transient strong access

Qubits	Arm	Frame	Eligible / total site-unit pairs	Ever reaches d >= 0.8
4	native	L	288 / 384	5
4	native	C	192 / 384	0
4	rotated	L	96 / 384	0
4	rotated	C	288 / 384	0
6	native	L	480 / 576	10
6	native	C	384 / 576	0
6	rotated	L	288 / 576	0
6	rotated	C	480 / 576	0

A site-unit pair is not an independent model. These are post hoc diagnostic counts from existing trajectories. A transient threshold crossing does not establish full-window persistence or redundant copying. The strict persistent-record result remains zero under the original definition.

S4. Factorial extension and paired uncertainty

S4.1 Design and replay

The review extension crosses injection L/C with Hamiltonian orientation L/C, with $H_C = U_C H_L U_C^\dagger$. Every unit supplies all four outcomes. The time grid is 0.05 through 2.5 plus the switch-off readout; the common injection code and two-layer circuit follow the original source-off conventions. The main paper and results log define z and e_a. We retain supplied inputs in feedback/current/ without editing their scripts or their predictions.

code/replay_reviewer.py runs the supplied functions with their original seeds and sample sizes and retains per-unit observables, rather than only aggregate summaries. It replays the 41,000 factorial units, the 60,000-unit reimplementation and its smaller checks, and the 6,000-unit third-frame comparison. Each arm in a unit remains paired. The statement that two routines were written independently refers to source implementation, not statistical independence of their authors' reasoning or external peer review.

S4.2 Range and finite-family bound

At fixed injection, the initial gap cancels in the difference between dynamics arms:

$$e_a = \tfrac{1}{2}(g_L^a - g_C^a) \in [-1, 1].$$

Let s_a^2 be the unbiased sample variance of e in N iid units. Theorem 4 of Maurer and Pontil (2009) applies to variables in [0,1]. Rescale by $Y = (e+1)/2$, use both tails, and allocate total error α over K specified targets. A simultaneous interval for each mean is its sample mean plus/minus

$$r_N = \sqrt{\frac{2s_a^2 \log(4K/\alpha)}{N}} + \frac{14 \log(4K/\alpha)}{3(N-1)}.$$

Derivation. The theorem's one-tail bound is $\sqrt{2V_N \log(2/\delta)/N} + 7\log(2/\delta)/(3(N-1))$. Set $\delta = \alpha/(2K)$, note $V_N(Y) = s_a^2/4$, multiply by two to return to e, and union over the 2K events. No independence between targets is required. The iid assumption is within each model-unit sample; paired arms are used to form one e. The variance identity $V_N = N^{-1}(N-1)^{-1}\sum_{i<j}(Y_i - Y_j)^2$ equals the usual unbiased variance. Numerical rounding is not covered by the sampling theorem.

For K=6 and alpha=0.05, all four intervals at N=20,000 exclude zero. The two eight-qubit intervals at N=1,000 do not. Full normal, Hoeffding and Bernstein intervals are in data/factorial_analysis/factorial_intervals.csv. The method was chosen during review; this is not an external preregistration or a claim that arbitrary post hoc selections retain coverage.

S4.3 Important interpretation checks

The original negative change relative to frozen dynamics and the positive Hamiltonian-intervention contrast concern different counterfactuals. Neither invalidates the other. Isospectrality preserves the energy eigenvalues, not the initial energy distribution or all resources relative to the fixed input. The third-frame comparator is an additional sampled intervention, not a certified link-only control. A Lieb-Robinson bound is not a sign theorem for e. Regression on an initial gap must account for the gap's explicit minus-one-half contribution to z.

Reproduction commands, from the extracted release root

```
OPENBLAS_NUM_THREADS=1 OMP_NUM_THREADS=1 python code/replay_reviewer.py --root . --out
 ↪ /path/to/fresh_reviewer_replay --task all
python code/analyze_factorial.py --root . --out /path/to/fresh_factorial_analysis
OPENBLAS_NUM_THREADS=1 python code/validate_factorial.py --root . --out
 ↪ /path/to/fresh_factorial_checks.json
```

The analysis reads the delivered data/reviewer_replay by default. To analyse a replay, use a separate copy with the fresh replay at that relative location. Do not overwrite the delivered evidence. The README records exact environment and tool versions; portability to an arbitrary numerical environment is not assumed. All original scripts can also be run directly in a copied input folder, with the single-thread setting stated by their author.

S5. Historical execution and provenance

The following subsections preserve the earlier account, with their historical dates and sample sizes. These are not fresh counts of independent physical experiments.

S5.1 Scope

This is supporting material for *Frame-Aligned Records*, review manuscript v0.5, not an independent research paper. It records new executions, supplied-script reproduction, prior reported results and what was not done. The original LRA population study remains unexecuted. A separately defined source-off feasibility pilot, P01-ID1, has now run; see the dedicated results log and protocol. It is not silently relabelled as the original study. No participant data, cosmological likelihood fit or unrestricted subsystem search was analysed. The mathematical propositions are assessed through their proofs; numerical checks are finite safeguards, not universal demonstrations.

Table S1. Execution levels must remain separate

Layer	Work done in v1.6	Limit
Supplied post-v1.5 Claude audit	Executed both numerical scripts in a fresh directory and compared their two JSON outputs	Same-code replay, not a new external implementation
New CAL-1.6 suite	174 assertions, including random-case checks and exact small examples	Related assertions, not 174 independent experiments
Separate CAL-1.6 checker	46 checks using independent small formulas and enumeration; does not import the main script	Same-project second implementation, not external scientific replication
Prior v1.5 results	Preserved unchanged with source files and historical reports	Its large historical assertion counts are not fresh v1.6 executions
Proposed LRA-count experiment	Specification and uncertainty handling revised	Not run, registered or scientifically confirmed

S5.2 Replay of the supplied audit

The exact delivered audit scripts are retained under historical/supporting_lineage/v1_6_code/supplied_claude_audit/. Both ran successfully from an initially empty output directory. v15_independent_checks.py took about 43 seconds and coupling_dominance.py about 2 seconds in this environment. Runtime is descriptive, not a portable performance promise. Captured output and environment differences are recorded in historical/supporting_lineage/v1_6_verification/claude_replay/.

Recursive comparisons of dictionaries, lists and numerical leaves found agreement with both supplied JSON outputs using relative tolerance 10^{-8} and absolute tolerance 10^{-10}. Key sets and list lengths must agree; timing and unrelated files are not used to infer numerical correctness. The comparison report is historical/supporting_lineage/v1_6_verification/claude_replay_comparison.json.

The replay includes the supplied exact four-label counterexample, mixed twirl, gauge invariances, local-persistence example, symbolic examples and the reported permutation Monte Carlo. For 30 units, the displayed derangement rejection estimate is approximately 0.060 at a nominal 0.05 threshold; that is a finite Monte Carlo estimate with sampling uncertainty, not an exact universal size. The exact four-label example separately establishes a failure of the proposed general validity claim.

The 9,000 perturbation checks all satisfy their bound within the script tolerance. Those are fragment-time checks across related models, not 9,000 independent experiments. They confirm the stated implementation of the inequality, not a theorem that coupling-frame records must maximize a locality-record score.

S5.3 New suite and effective configuration

The new executable source is historical/supporting_lineage/v1_6_code/checks_v0_3_0.py. It accepts one complete JSON configuration, rejects missing or unused keys, checks numerical ranges, refuses a nonempty output directory, and writes the effective configuration and its canonical-JSON hash with every run. The version is CAL-1.6; the code version is 0.3.0. This is calibration code, not the future LRA experiment implementation.

Table S2. Complete new calibration configuration

Field	Value	Use
version	CAL-1.6	Validate scope
seed	20260926	NumPy default RNG stream
random_cases	32	Each random matrix/invariance/perturbation check family
mixed_draws	1200	Fixed-spectrum twirl Monte Carlo
word_length	5	Symbolic support enumeration
chain_length	30000	Finite golden-mean example
bound_epsilon	0.1	Adaptive-family probability envelope
net_radius	0.05	Operator-norm cover radius
angle_limit	pi	Declared circuit parameter interval

Field	Value	Use
max_bound_qubits	32	Envelope evaluation through 32 qubits; no 32-qubit state simulation

The canonical effective-configuration SHA-256 is f5f8d1633e36e402922eefb5fe8a24640b890fa028e0 985856a753f779cabc2b. Numerical environment: Python 3.13.5, NumPy 2.3.5 and SciPy 1.17.0. The supplied Claude audit originally reports Python 3.11.15, NumPy 2.4.4 and SciPy 1.17.1, so agreement is assessed with tolerances, not assumed from version identity. The code's environment report supplies the actual current values rather than promising installation compatibility with every platform.

Table S3. Data files and their unit of observation

File in historical/supporting_lineage/v1_6_data/	Contents	Interpretation
per_case_checks.csv	32 cases, gauge and stabilizer residuals, local bound and sampled slope	Numerical checks of deterministic inequalities
perturbation_checks.csv	32 differences and perturbation bounds	New small finite checks, distinct from the replayed 9,000 checks
localized_constants.csv	Global and endpoint touching constants for 2, 4, 6 qubits	One explicitly specified bounded-strength model
mixed_draws.csv	1,200 squared Hilbert-Schmidt deviations	Individual draws, not only an aggregate
symbolic_support.csv	All 32 binary length-five words and their model probabilities	Exact support and finite-sample occurrence are distinct columns
adaptive_cover_bound.csv	Log probability envelopes for n=4 through 32	Evaluation of a theorem bound, not physical observations
results.json	Assertions, selected summaries, environment and code hash	Failures cannot be hidden by counting only passed examples

The old constant-window example retains its analytic certificate 0.0097751350, before adding a numerical error budget. The new endpoint chain has $K_F = \sqrt{1.0025}$ for every plotted size. The supplied random-chain example is a different construction, not a failed attempt to reproduce this value.

S5.4 Independent small verification and reproduction instructions

The second checker does not import the main checker. It reconstructs exact permutation ranks, contracts the twirl in a different form, evaluates endpoint spectra analytically, and checks the circuit-cover arithmetic. Its 46 checks are another implementation path within the same session, not a second research team.

From the extracted release root, use initially absent output directories:

```
python historical/supporting_lineage/v1_6_code/checks_v0_3_0.py --config
 ↪  historical/supporting_lineage/v1_6_code/config.json --output /path/to/fresh_calibration
python historical/supporting_lineage/v1_6_code/second_checks_v0_3_0.py --results
 ↪  /path/to/fresh_calibration
python historical/supporting_lineage/v1_6_code/plot_results.py --data /path/to/fresh_calibration
 ↪  --output /path/to/fresh_figures
```

The legacy scripts are preserved for provenance and old calibration replay. They are not silently relabelled as corrected production software. In particular, the earlier hard-coded defaults and aggregate-only outputs remain historical limitations. New matrix cases and mixed draws have individual rows; a future model-search implementation must additionally log every input frame, rejection, tie, search failure and certificate interval.

For raw floating-point checks, a small residual is not a verified interval enclosure. The many-body continuum certificate requires a defensible numerical allowance or interval method. Without one, report numerical evidence and a certificate conditional on the allowance, not rigorous machine certification.

S5.5 Conditional pairing retained only as an alternative

The v1.6 primary prospective target compares normalized counts within an input-eligible model. It uses no cross-model pairing. The previous pairing method remains meaningful only under its own contract: specify the independent unit, frame registration, conditioning variables, admissible group action, why the null is invariant, tie handling and missing-output treatment. A complete permutation group cannot create exchangeability in a physical population that lacks it.

Fixed common frames can make matched and mismatched distances identical. Randomly varying shared orientations can produce an apparent pairing signal inherited from supplied structure. These are valid diagnostic examples, not reasons to ban every permutation test. Where the null cannot be justified, the appropriate status is NO_IDENTIFIED_PAIRING_NULL. The original conditional wording is preserved in editorial/retained_pairing_contract.md.

S5.6 Delivery and scientific readiness

The present supplement is complete for inspection at its declared finite scope. The full release contains the exact new code and data plus the supplied audit and baseline archive. A short manuscript-only review necessarily cannot certify those absent files, and no claim in the text turns that limitation into disconfirmation. Source parity, render inspection, archive integrity and scientific correctness are separate checks, recorded in the release verification report.

The new uniform-cover theorem is a conditional mathematical corollary of known concentration and covering arguments. Its novelty and physical relevance remain open. The broad AIU inquiry neither inherits empirical confirmation from it nor is refuted by a poor finite-size bound.

S5.7 Maintenance patch and separate source-off execution

Appendix C of the methods paper retains the archived v1.5 2,000-draw twirl result (mean 0.0557464152; standard error 0.0007264476). CAL-1.6 used 1,200 draws instead (mean 0.05550839355; standard error 0.00093479233), with exact target 0.056. These are two calibration samples, not a conflict or one enlarged study. The historical data paths identify them separately.

P01-ID1 is a new finite computational pilot with 96 independent model units at each of four and six sites. Each unit supplies two paired injection arms and paired evolution-on/frozen controls. Its primary outcome is a baseline-subtracted checkpoint-average local-readability contrast, not persistent redundancy, the older LRA-count target, or a recovered TPS. Full trajectories, model coefficients, random-stream identifiers, configuration and code hashes are in data/pilot_run/; the specification and local pre-execution receipt are in pilot/. The specialist review was not sent and no external validation is claimed.

The main solver diagonalizes each internal Hamiltonian. code/validate_injection_pilot.py independently checks selected trajectories with exponential action and a full-density-matrix partial trace. It also checks exact transport, source-free global distinguishability, zero evolution, passive conjugation and closed stationary records. Its finite-precision agreement is not a rigorous roundoff enclosure. Persistence classifications are conditional on the explicitly empirical numerical guard; failed certificates are not absent records.

Reproduction of the new pilot from an empty directory:

```
OPENBLAS_NUM_THREADS=1 OMP_NUM_THREADS=1 python code/run_injection_pilot.py --config pilot/config.json
 ↪  --out /path/to/fresh_run
python code/validate_injection_pilot.py --root /path/to/extracted_patch
```

The validator reads the delivered data/pilot_run by default. To validate a replay, place that replay at the same relative path in a separate extracted copy. Do not overwrite the delivered evidence. The code and effective configuration specify all model parameters; no unused fields are silently accepted.

References

Maurer, A., and Pontil, M. (2009). Empirical Bernstein Bounds and Sample Variance Penalization. COLT 2009. https://arxiv.org/abs/0907.3740