

Computational Bandwidth Limits Unify Gravitational and Quantum Dynamics in a Discrete Spacetime

Jason Merwin

July 2026

Abstract

We demonstrate that gravitational and quantum dynamics can be unified as emergent consequences of strictly bounded computational bandwidth within a discrete generative spacetime (the Distinction Engine Universe). In this framework, the fundamental phenomena of physics arise from scheduler load: persistent topological defects impose chronic bandwidth demand (analogous to gravitational mass), while acute logical conflicts trigger saturation-limited resolution events (analogous to quantum state collapse). Through pre-registered experiments on a growing triangulated foam, we derive a unified two-queue ledger law—governing forced backlog and standing vacuum demand—that dictates this load response with zero fitted parameters. This single internal law accurately brackets the chronic starvation collapse threshold (a gravitational-regime constant of $x_{50} = 1.30$) and completely reproduces acute Zeno-pattern suppression in temporally structured loads. While the central claim is a model-internal unification of these two operating regimes, we provide external bridge diagnostics—including a CMB amplitude basin, preliminary GW150914 ringdown-compression signatures, and formal equivalence to standard quantum and relativistic clock laws—as falsifiable targets for physical equivalence.

1 Introduction and scope

It is known that Einstein’s General Relativity—as usually understood at the present time—treats spacetime as a fundamentally continuous and smooth manifold. In contrast, this work models spacetime as the output of a finite generative process: a distinction combinatorial graph [4]. Within this framework, we demonstrate that the fundamental phenomena of physics—gravitation and quantum state resolution—are emergent consequences of scheduler load: persistent defects (mass) impose chronic bandwidth demand, while acute paradoxes trigger saturation-limited resolution events. This paper presents a unified, pre-registered experimental study of these load economics, establishing the quantitative laws governing both the gravitational collapse threshold and the suppression of quantum Zeno patterns.

This paper separates three evidential tiers. The first tier is the measured engine: a frozen substrate, scheduler, and relief mechanism are subjected to chronic, acute, and temporally structured load, with pre-registered predictions and held-out confirmations wherever a quantitative claim is made (Sections 4–5). The second tier is the derived engine: Section 6 reduces the same runs to a five-rule two-queue law whose one internal constant is derived rather than fitted. The third tier is not part of the proof but is important for interpretation: Section 8 and Appendices A–C collect external bridge diagnostics and formal textbook substitutions from companion audits. These bridges are reported as evidence of reach and as falsifiable targets, not as claims that the DEU has derived nature.

2 The engine

2.1 Substrate

The substrate of the distinction engine is a growing triangulated 2-complex [2]. Faces carry a refinement depth k ; a face of depth k has weighted area 3^{-k} and weighted edge length $3^{-k/2}$, so total weighted area is conserved under the 1→3 refinement move. Face types $\{S, I, G\}$ drive growth through a frustration rule (S faces with a G neighbor and no I neighbor split); child types are assigned by random permutation. Degree regulation is provided by 2–2 edge flips accepted only when they reduce the local cost $\sum_v (\deg v - 6)^2$ and join faces of equal depth.

Both substrate rules were forced by measurement, not chosen a priori. Sorted type assignment concentrates half of all faces on a single immortal vertex (degree 6,318 of 12,636 faces in a reference run), destroying all metric structure. Depth-unrestricted flips collapse the weighted diameter to ≈ 0.53 at every flip cadence tested: a flip joining unequal depths rewires adjacency across scales and acts as a metric wormhole. With both constraints imposed, the vacuum passes its geometric entrance exam: ball-growth dimension $d = 2.00 \pm 0.11$ across seeds at $\sim 3 \times 10^4$ faces, weighted diameter 1.0–1.5, maximum vertex degree bounded at 25–28.

We state the design law these constraints exemplify, because it has now recurred independently *five* times in this program (flips, excision scale, collapse orientation, observable definition, and—in Section 6.1—instrument timing):

Every legal move, and every observable, must respect the structure of the metric it acts on—its scale, its radial fibration, its material identity, *and its timing relative to the events it measures*. Metric-blind moves destroy emergent geometry; coordinate-fixed observables misread material surgery; event-blind instruments misread their own mechanism.

2.2 Native scheduler

A core window W of weighted radius r_c surrounds the defect anchor. Its per-epoch capacity is Γ , the live count of its boundary edges—counted by the scheduler at runtime, never supplied as a parameter. Forced defect operations and vacuum maintenance (frustrated splits) compete for Γ in randomized order; unserved forced operations persist as a literal backlog B . The exterior runs on a sub-extensive global budget $N^{3/4}$ (a declared simplification). Sub-extensive capacity is itself empirically forced: with extensive capacity the defect accelerates local activity and the starvation channel inverts sign; positive time-dilation-like response exists only for capacity exponents $p \lesssim 0.75$.

2.3 Relief valve

When $B \geq \Gamma$, one relief event fires. Two mechanisms are compared: an oracle (delete $\lfloor T/6 \rfloor$ of the T deepest defect faces by fiat) and the collapse valve: link-conditioned edge collapses [6] restricted to circumferential defect–defect edges at the defect’s own metric scale, voiding one deferred obligation per removed defect face. In $> 10^4$ executed collapses across all experiments the edge-incidence audit recorded zero manifold violations.

2.4 Instrumented variants and certification

Sections 5–6 required event-level observables not present in the base engine’s log. Each instrumented variant (v21i1: tagged population and per-event voiding; v21i2: population read *at trigger*; v21i3: standing vacuum demand) adds only pure reads—no dynamics, no random draws. We adopted the rule that *instruments are engines*: every variant must reproduce, bit-exactly, a fixed panel of ten archived runs spanning both load regimes, both outcome classes, and the extremes of the capacity lottery, across all four scalar observables, before any of its data are used. All variants passed 10/10; certification scope (scalar pipeline observables) is declared, and geometry-sector claims from these variants are out of scope.

2.5 Reproducibility

All engines are deterministic given a seed. Code, per-run data, pre-registration documents with frozen timestamps, and the continuity ledger sufficient to reproduce every table and figure are provided in the supplementary archive.

3 Methods: measures, metrics, and pre-registration

Three methodological rules produced every surviving result and caught every retracted one. (1) **Declared measures**: on a multi-scale foam, node-count, face-count, area-weighted, and length-weighted statistics of the “same” quantity differ by orders of magnitude; one early “geometry-collapse” finding was an artifact of switching measures between sessions and is retracted in Section 7. This campaign forced a new clause: an observable must also declare its *timing* relative to the event it measures (served vs. demanded; pre-trigger vs. post-relief). (2) **Pre-registration**: quantitative predictions and fit procedures were frozen before the runs that tested them; amendments are permitted only with the prior failure and its diagnosis on the record first. (3) **Bookkeeping identities as instrument checks**: the measured late-time backlog slope equals $m - \bar{s}_F$ to within a few percent in every cell of every sweep.

The primary chronic observable is the late-window backlog slope; a run is classified *runaway* if the slope exceeds 1 op/epoch over epochs 70–100. The primary relief observable is the stitch count k . The dimensionless load coordinate used throughout is

$$x = \frac{m}{\Gamma - D_{\text{vac}}}, \quad (1)$$

with Γ and D_{vac} measured per run (late-window means). Threshold fits are maximum-likelihood logistic; confidence intervals are 2000-resample bootstraps (fixed rng seed), resampled at the run level where events within a run are correlated.

4 Chronic load: the gravitational-regime economy

4.1 Gravity without relief: a graded starvation continuum

With the valve disabled, chronic load m produces neither stability nor a binary collapse but a continuum of partial starvation. The normalized starvation fraction $s = \text{slope}/m$ occupies the middle band $0.2 < s < 0.8$ in 43–70% of runs at every load, with an atom of total starvation ($s = 1$; 7–27%) and a persistent survivor tail (3–13% fully balanced even at $m = 42$). An earlier

m	P	oracle	P	collapse valve
		$\bar{k}$		$\bar{k}$
4	0.07	13.3	0.00	14.0
8	0.10	21.1	0.00	22.9
12	0.13	26.4	0.20	29.0
16	0.33	30.9	0.23	34.7
20	0.53	35.7	0.23	37.8
26	0.63	39.3	0.53	43.0
34	0.80	44.5	0.70	46.4
fits	$m_{50} = 19.5; k \propto m^{0.57}$		$m_{50} = 26.2; k \propto m^{0.57}$	

Table 1: Valve adjudication sweep, $r_c = 0.06$, 30 seeds per cell. P = probability of runaway; $\bar{k}$ = mean relief events.

“bistable death spiral” reading from $n = 8$ pilot statistics is retracted; the $n = 30$ distributions are unimodal with heavy endpoint structure. Within cells at fixed load, outcome correlates strongly with each run’s own native capacity: the seed lottery is a capacity lottery. (Section 5.3 quantifies this: the per-run coordinate of Eq. 1 predicts individual outcomes at AUC 0.98–0.999 on three independent seed cohorts.)

4.2 The dimensionless collapse constant (no relief)

Raw collapse thresholds differ across window geometries ($m_{50} = 4.7$ and 5.6 at $r_c = 0.06$ and 0.12). Normalizing by gross capacity fails to reconcile them (ratio 1.21); normalizing by free capacity (Eq. 1) collapses both onto $x_{50} = 1.30$ exactly. A pre-registered falsification test on a third geometry ($r_c = 0.20$), run blind with the fit procedure frozen, gave

$$x_{50} = 1.29, \quad 95\% \text{ CI } [1.17, 1.45] \quad (\text{prediction: } 1.30). \quad (2)$$

The no-relief economy tolerates a chronic overload of $\approx 30\%$ of its free capacity before collapse odds reach one half, independently of window geometry.

4.3 Relief: two valves, one law, unequal margins

A matched dual-arm sweep (30 seeds, $m \in \{4, \dots, 34\}$, identical trigger) adjudicates the collapse valve against the oracle (Table 1, Fig. 1). Both arms obey the same stitch-rate law, $k \propto m^{0.57}$, with absolute counts agreeing within $\sim 5\%$ at every load. Their stability margins differ: $m_{50} = 26.2$ (collapse valve) vs. 19.5 (oracle). Relative to the no-relief economy ($m_{50} \approx 5$), relief extends the stable regime by a factor of ~ 4 . The valve maintains a genuine limit cycle with a persistent defect: tagged population oscillating between 59 and 177 faces over fifty epochs while backlog repeatedly returns to zero—relief, not annihilation. (We note for Section 6 that the 0.57 exponent fits were taken at loads where the per-epoch relief cap begins to censor k ; a below-ceiling refit is registered future work.)

4.4 The relieved constant and its universality

The relieved economy has its own wall. On a fresh seed cohort (141–160) under pure chronic load, the relieved threshold in free-capacity units is $x_{50} = 3.52$, 95% CI $[3.37, 3.71]$, consistent with the

value obtained independently from pulse-train schedules on a different cohort (3.72, [3.60, 3.87]; Section 5.3). A pre-registered three-arm test then varied the exterior capacity exponent, $p \in \{0.6, 0.75, 0.9\}$, same seeds and loads. Raw fragility differs wildly across arms (at $m = 32$, $P(\text{runaway}) = 0.85$ at $p = 0.6$ vs. 0.50 at baseline) because the exterior exponent changes the substrate the window grows in (mean free capacity 6.4/10.5/8.0); the dimensionless wall does not move: $x_{50} = 3.53$ [3.41, 3.62], 3.52 [3.37, 3.71], and 3.45 [3.37, 3.56] (Fig. 2). The registered caveat that the starvation channel might invert at $p \approx 0.9$ did not fire. **The relieved constant (≈ 3.5 in free-capacity units) is a constant of the mechanism class, not of the exterior scheduling exponent.** The small spread between the pulse-train estimate (3.72) and the chronic-arm cluster (3.45–3.53), with overlapping CIs throughout, is a possible schedule effect and is flagged as a consolidation item, not claimed either way.




Figure 1: Valve adjudication. Left: $P(\text{runaway})$ vs. m , both arms. Right: mean stitch count, log-log, with shared exponent 0.57.

the relieved constant: schedules, cohorts, capacity expone




Figure 2: The relieved constant across schedules, independent seed cohorts, and exterior capacity exponents $p \in \{0.6, 0.75, 0.9\}$.

5 Acute and structured load: the collapse-regime economy

The engine, valve, trigger, and all constants are frozen from Section 4; only the exogenous load schedule changes.

5.1 The render gate

Single-epoch spikes of total action A produce no topological response below a threshold and a graded burst above it (Fig. 3). Pre-registered onset $A^* \approx s + \Gamma$ predicts $A^* \approx 23$; measured onset 15–20, within seed scatter. Above onset, $\bar{k}$ rises from 0.6 ($A = 20$) to 7.0 ($A = 150$) while clearing time grows only from 2 to 8 epochs: denser stitch bursts, not longer sieges. A spike of $A = 150$ always clears; the equivalent chronic load kills a quarter of runs.

5.2 Zeno-pattern suppression

Fixing total action $A_{\text{tot}} = 60$ and subdividing into N pulses:

schedule	1×60	2×30	4×15	6×10	10×6	15×4	30×2
$\bar{k}$	3.4	2.4	0.8	0.6	0.0	0.0	0.0

The suppression is monotone and complete (Fig. 4), exactly as pre-registered: once each pulse can be serviced below the trigger line before the next arrives, no sector is ever dropped. The temporal structure of action delivery alone determines whether topology changes. We name it Zeno-pattern suppression from bandwidth economics and claim nothing further.

5.3 One wall across the waveform plane, and a knee

A pre-registered waveform-plane campaign (loads $\bar{m} \in \{16, \dots, 40\}$, periods $\Delta t \in \{1, 2, 4, 8\}$, 20 seeds; extended to $\Delta t \in \{5, 6, 7\}$ on a fresh cohort) tested whether the chronic wall is waveform-invariant. Three results (Fig. 5):

(a) *The reduction generalizes.* The per-run coordinate $x = \bar{m}/(\Gamma - D_{\text{vac}})$ classifies individual runaways at AUC 0.979–0.999 on all cohorts, versus 0.59–0.66 for raw load. The chronic outcome of an individual run is essentially determined by its own free capacity—the capacity lottery, quantified.

(b) *Waveform invariance holds through $\Delta t = 5$.* Thresholds at $\Delta t = 1, 2, 4$ are statistically one number (pooled $x_{50} = 3.72$, [3.60, 3.87]); matched-pair tests against the chronic arm are null at $\Delta t = 2, 4$. A fresh-cohort point at $\Delta t = 5$ (3.69 [3.50, 4.03]) sits on the same band.

(c) *Above $\Delta t = 5$ the wall lifts along a graded ramp:* 4.16 [3.94, 4.38] at $\Delta t = 6$, 4.01 [3.77, 4.23] at $\Delta t = 7$, 5.30 [4.91, 5.71] at $\Delta t = 8$, where the effect is unambiguous (21 matched runs rescued, zero newly-runaway; $p < 10^{-5}$). Long-period packaging of equal mean load is genuinely stabilizing, and not through the capacity lottery: the lift survives the free-capacity reduction.

5.4 A registered failure: the knee is not drain arithmetic

The natural mechanism—a pulse is harmless if it can be drained before the next arrives—was pre-registered quantitatively: per-seed clear-time laws $T_{\text{clear}}(A)$ were measured on the fresh cohort by single-spike calibration, the prediction table was frozen *before* the knee grid ran, and the criterion (80% per-cell accuracy) was fixed. The prediction failed catastrophically and informatively: 30.6% accuracy, with 1.4% of cells predicted stable against 70.8% observed.

Pulse *trains* clear far faster than matched isolated spikes. The failure direction indicated a relief mechanism whose output grows with the standing defect population—a proportional controller—which Section 6 confirms and derives. This registered failure, reported here with its registration intact, is the first link in the chain that produced the paper’s central result.




Figure 3: The render gate: mean stitches vs. single-epoch spike size A . Dead zone below A^* , graded burst response above.

6 The derived two-queue law

6.1 The valve transfer function is exact

To measure the conjectured controller we instrumented the engine (Section 2.4) and pre-registered a statistical test: relief output per event vs. tagged population, predicted slope $1/6$ from the quota arithmetic. *The test failed as frozen* ($\beta \approx 0.193$, tightly clustered, excluding $1/6$; invariance formally rejected). The diagnosis is a methods specimen we retain deliberately: the registered predictor was population *at trigger*; the instrument logged population at epoch end—after the relief event removed its faces—turning n/pop into $n/(\text{pop} - n) \approx 1/5$. This is the declared-measure error class of Section 3, committed by the analysts, caught because the prediction was frozen. The amendment (corrected instrument, certified bit-exact; deterministic event-level prediction) was registered with the failure on the record first.

The corrected result is stronger than the hypothesis it replaced. The transfer function is not statistical at all; it is exact (Fig. 6):

$$n_{\text{voided}} = 2 \left\lceil \frac{1}{2} \max(1, \lfloor \text{pop}_{\text{trig}}/6 \rfloor) \right\rceil, \quad (3)$$

matching 4227 of 4230 relief events across chronic, pulsed, and burst schedules, all loads, and three seed cohorts; the three exceptions are all shortfalls (candidate exhaustion), and the exhaustion channel is otherwise never binding up to populations of 1227. The gain $1/6$ is derived from the quota arithmetic; the quantum of 2 is derived from the topology of the collapse move




Figure 4: Zeno-pattern suppression: mean stitches for equal total action ($A_{\text{tot}} = 60$) under increasing temporal subdivision.




Figure 5: Waveform invariance and the knee ramp: x_{50} vs. pulse period, two seed cohorts, with the pooled $\Delta t \leq 4$ band. Invariance holds through $\Delta t = 5$; the wall lifts along a graded ramp to $\Delta t = 8$.

(each collapse removes a face pair). Equation 3 is the first constant of this economy that is *derived* rather than measured.

6.2 A five-rule reduction and a chain of registered nulls

We then asked whether the engine’s entire load response follows from Eq. 3 plus measured capacities. The reduced model was frozen before adjudication: per epoch, (R1) delivery per schedule; (R2) forced requests $R_f = a_t + B$ compete with vacuum demand for capacity Γ (hypergeometric share, the engine’s literal randomized-order semantics); (R3) backlog $B \leftarrow R_f - \text{served}$; (R4) population $\text{pop} \leftarrow \text{pop} + 2 \text{served}$; (R5) if $B \geq \Gamma$ and $\text{pop} \geq 6$, relieve per Eq. 3. Constants: Γ and vacuum terms measured per run; gain derived. Nothing is fitted.

Adjudicated against 500 archived engine runs, the one-queue form (vacuum represented by its served rate) classified individual outcomes at 90.8–95.6% but placed every wall uniformly $\sim 10\%$ high—too stable. Two registered amendments then returned nulls: service shot noise moved nothing, and substituting each run’s full measured capacity trajectory $\Gamma(t)$ moved nothing (yielding, as a by-product, the robustness result that wall location is insensitive to capacity fluctuations—only means matter).

6.3 The standing pool: the economy has two queues

An instrumented diagnostic then found the structural omission. Core vacuum demand is not a flow near its served rate; it is a *standing pool*, $10\text{--}125\times$ the served rate, with $\Gamma - D < 0$ throughout the near-wall regime (stable chronic runs carry ~ 65 frustrated faces against $\Gamma \approx 17$; runaways carry ~ 160). The economy is a two-queue system: the tracked forced backlog and an untracked vacuum-demand pool compete for capacity, and near the wall, stability is purchased by relief, not by service. The one-queue reduction had been silently omitting one of the two queues—and still classifying 95% of runs, because the wall’s *location*, not its existence, is what the pool moves.

6.4 The bracket, the derivation, and a declared stopping point

Substituting the measured demand trace $D(t)$ for the served-rate proxy—the final registered amendment, still with zero fitted parameters—moved every wall by the size of the offset (Fig. 7). The two-queue law *derives*, within engine confidence intervals: the chronic wall ($3.45 \in [3.37, 3.71]$, at 98.3% per-run agreement), the $\Delta t = 1$ wall ($3.44 \in [3.35, 3.81]$), the $\Delta t = 7$ point ($3.96 \in [3.77, 4.23]$), and the burst-stabilization endpoint ($4.93 \in [4.91, 5.71]$ at $\Delta t = 8$). Two points sit marginally low (3.38 vs. $[3.50, \cdot]$ at $\Delta t = 5$; 3.84 vs. $[3.94, \cdot]$ at $\Delta t = 6$). Every measured wall in the study lies between the served-rate variant (uniformly high) and the demand-trace variant (marginally low): the bracket bounds the residue at $\pm 3\%$, and its sources are enumerated—token-level service microstructure (capacity redistribution when a drawn request is unservable)—rather than fitted.

We declare this the stopping point of the reduction program. Modeling the $\pm 3\%$ means reimplementing the engine inside the “reduced” model, at which point the derivation claim voids itself. The residue’s epistemic status is: bounded by the bracket, attributed by enumeration, permanently, unless a later tier requires otherwise.

6.5 The unification statement

Within the DEU: a two-queue law—forced backlog and standing vacuum demand competing hypergeometrically for boundary capacity, relieved by the derived $1/6$ even-quantized controller of Eq. 3—reproduces individual run outcomes at 95–98% and derives the chronic wall, the waveform-invariant band, and the burst-stabilization endpoint within measurement confidence intervals, with zero fitted parameters. The chronic (gravitational-regime) and acute (collapse-regime) phenomenologies of Sections 4–5, including the render gate and Zeno-pattern suppression as limiting cases of temporal structure, are two operating regimes of this one law. The conjecture of the earlier draft is retired; what replaces it was reached through five pre-registered adjudications of which three failed, each failure diagnosed on the record before its amendment.




Figure 6: The valve transfer function, event level: obligations voided per relief event vs. tagged population at trigger, with the exact law of Eq. 3 (4227/4230 events exact; the three exceptions are exhaustion-side).

7 Negative results and retractions

We report these with the same prominence as positive results.

Retractions from earlier sessions and drafts (unchanged from the prior draft): (i) a claimed “13/13 staircase verification” tested code against its own update rule and is void; (ii) a claimed curvature-saturation law ($k_{\max} \approx 3$) died under its pre-registered static control; (iii) a claimed geometry collapse under the native economy was a measurement-measure artifact; (iv) a bistable-collapse reading died under adequate statistics; (v) all geometric measurements on the pre-regulation substrate are unreliable.

Registered failures of this campaign are not retractions and remain in the results sections with their registrations: the drain-arithmetic knee prediction (Section 5.4); the statistical transfer-function test and its instrument-timing diagnosis (Section 6.1); and the shot-noise and capacity-trajectory nulls (Section 6). These five adjudications are the discovery path of the two-queue law.




Figure 7: Derived vs. measured walls. Bars: engine 95% CIs. Filled points: the two-queue law (zero fitted parameters). Open points: the one-queue variant. Every measured wall lies inside the two-variant bracket.

The witness-loop null. Servo-contracting material circles around the defect by up to 50% produces zero deficit on an untouched witness loop: circumference reduction without area excision creates no far-field curvature.

The open geometric problem. Across nineteen engine versions we attempted controlled macroscopic wedge excision-with-identification on the adaptive multi-scale foam and assembled a complete failure taxonomy: corridor floods leak at vertices; geodesic-pair sewing manufactures edge fans; boundary zipping fails on non-simple scars; collapse-based excision has uncontrollable quanta across four orders of scale; ring-sewing is curvature-null. A far-field deficit operator ($\Xi = k/6$ at all radii) remains unbuilt. We pose it, with the taxonomy as specification, as an open problem in computational geometry.

8 External bridge evidence and falsifiable diagnostics

The derivation above is deliberately internal: it shows that one frozen DEU engine and one derived two-queue ledger law generate both chronic (gravitational-regime) and acute (collapse-regime) phenomenology. It does not, by itself, establish that the DEU is the ontology of physical spacetime. The broader program nevertheless contains several bridge diagnostics: quantities or signals that are not fitted inside the two-queue campaign, but which connect the same registry, carrier, or finite-ledger architecture to external observables or to standard textbook structures. Table 2 records the ones mature enough to cite without making them load-bearing for the present paper.

The two-queue result changes how these bridges should be read. Earlier drafts described gravity as one starved queue and measurement as one acute queue. The present campaign shows that the local ledger is richer: forced defect backlog and standing vacuum demand compete for the same boundary capacity, and relief is an even-quantized proportional controller. The bridge evidence should therefore be interpreted through finite-ledger structure, not through a generic

Bridge object	DEU-side construction	Evidence reported in companion audits	Status here
Derived two-queue ledger	Forced backlog and standing vacuum demand compete for boundary capacity Γ ; relief uses the derived 1/6 even-quantized controller.	Reproduces individual run outcomes at 95–98%; derives the chronic wall, the $\Delta t = 1$ wall, the $\Delta t = 7$ wall, and the $\Delta t = 8$ burst endpoint within engine CIs; brackets all six measured walls with a $\pm 3\%$ residue.	Main internal mechanism; not an external physics claim.
DAG7 registry and render gate	The distinction engine yields a $137 = 81 + 40 + 16$ registry; the unresolved ternary spatial slice has $81/3 = 27$ nodes, whose complete conflict with the 16-node scaffold gives $27 \times 16 = 432$.	The registry audit derives the sector partition from the distinction axiom; the closure-audit papers use 432 as the render/screening boundary, with interface screening at $I_{\text{crit}} = 3/2$.	Structural bridge; not required for the two-queue proof.
Graph-index amplitude	A continuum 40/40 carrier readout is replaced by a boundary-completed graph readout $(40 + 1)/40 = 41/40$.	Exact $A_G = 41/40 = 1.025$ occupies the same CMB amplitude basin as the measured graph branch $A_{\text{Planck}} = 1.02505572438$ ($\log BF \simeq 0.07$), while the continuum branch $A = 1$ fails internally ($\log BF \simeq 511.46$, mean CMB χ^2 gap $\sim 10^3$).	Strongest external bridge diagnostic; not a claim of DEU-over- Λ CDM model selection.
Finite-carrier cosmological closure	The reduced carrier foam has a unique self-reproducing carrier $C = (0, 1, 2)$ with measured density $\rho_C \simeq 0.01909246$; the open-to-mature registry factor is $124/120$.	The locked product $(124/120)(1 + \rho_C) = 1.053062209$ nearly closes the finite-count de Sitter residual $F_{\text{req}} = 1.052718776$.	Falsifiable bridge target; rejected if longer carrier-foam runs move ρ_C away from the locked value.
Ringdown compression	Finite acute ledger response predicts mode-population compression or clipping: denser early relief, not arbitrarily many longer-lived tones.	A GW150914 toy damped-sinusoid audit found the nonlinear DEU compression template strongest at three tones ($\Delta \text{BIC}_{\text{DEU-GR}} = 14.15$, $\Delta \text{BIC}_{\text{DEU-uniform}} = 12.11$, add-one $p = 1/49$), with the fourth tone degrading the fit.	Candidate external signal; Appendix A; requires Kerr-QNM hardening and multi-event replication.
Textbook-equivalence audits	Closure amplitude supplies route selection after the render gate: $\mathcal{A}_i = \sum_{p \in \mathcal{P}_i} \nu_p e^{-\alpha \tilde{C}(p)/2} e^{i\theta(p)}$, $P_i \propto \mathcal{A}_i ^2$.	With fixed $\alpha = 0.050$, audits recover Born cost log-odds, Mach–Zehnder interference, subthreshold Zeno freezing, WKB-like tunneling, Bell correlations $S_{\text{CHSH}} = 2\sqrt{2}$, and no-signaling marginals.	Formal equivalence, not empirical proof; Appendix B.

Table 2: Bridge evidence outside the main two-queue derivation [5, 4, 2]. The table separates the paper’s internal mechanism from external diagnostics and formal equivalence checks. None of these bridge rows is used to fit or rescue the two-queue law.

“simulation” analogy. In particular, the ringdown diagnostic is not a claim that black-hole quasinormal frequencies are wrong; it is a claim that a finite acute ledger may under-populate, compress, or clip higher mode families while leaving the main ringdown scale intact. Likewise, the textbook substitutions do not establish physical equivalence; they show that once a render event is timed by finite capacity, the route-selection side of the same registry can reproduce standard quantum probability structures under a declared closure calculus.

9 Conclusion and the frontier

Within the DEU, the economics of the generative ledger is now a derived body of law. The measured tier contains a graded starvation continuum; a no-relief collapse constant of 1.30 in free-capacity units, blind-confirmed; a manifold-legal relief valve exceeding fiat deletion’s stability; an acute render gate; Zeno-pattern suppression in the tested equal-action ladder; and a relieved-economy constant statistically consistent across the schedules and seed cohorts tested, and invariant under exterior capacity exponents $p \in \{0.6, 0.75, 0.9\}$. The derived tier contains an event-level valve transfer function whose gain and quantum follow from quota arithmetic and topology, and a five-rule two-queue law that reproduces the measured response surface with zero fitted parameters, four of six walls within confidence intervals and all six inside a two-variant bracket.

What this paper does not contain is geometry’s far field, and it says so: the deficit operator $\Xi = k/6$ is the program’s declared Track-B frontier, posed in Section 7 with its failure taxonomy as specification. Nor does this paper claim a physical derivation of general relativity or quantum mechanics. Appendices A–C record why the program is worth testing externally—a preliminary ringdown-compression signal and a suite of textbook-equivalence substitutions—but those bridges remain diagnostics, not premises. The model earned its core results the slow way: every inflated claim it once carried is retracted by name, and its central law was found by a chain of registered failures. The surviving result is narrower than a theory of nature and stronger for that narrowness: a finite, two-queue, topologically relieved ledger unifies the chronic and acute regimes of the DEU.

Acknowledgments

Simulation engineering, statistical analysis, and adversarial review were performed in collaboration with Anthropic’s Claude (Fable 5 model), whose contributions included the identification and retraction of several of the authors’ (and its own) errors documented in Sections 6.1 and 7. All claims remain the responsibility of the author.

Data and code availability

Engines (rung1 series including certified instrumented variants v21i/v21i2/ v21i3), sweep and reduction scripts, all per-run CSVs and event datasets, frozen pre-registration documents, the continuity ledger (v2), and the surgical failure-taxonomy archive accompany this paper. All runs are deterministic given the recorded seeds; a ten-row bit-exact certification harness is included as the resumption protocol. <https://github.com/jrmerwin/generative-ledger.git>

A Observational diagnostic: finite-ledger ringdown compression

This appendix preserves the preliminary gravitational-wave bridge from the earlier generative-relativity audit in the language of the present two-queue law. It is not part of the derivation in the main text. It is an external diagnostic: if strong gravitational mergers are acute finite-ledger events, then their ringdown should not necessarily populate an indefinitely smooth ladder of higher tones. The expected signature is *mode-population compression or clipping*: the main ringdown scale can remain intact while higher tones become under-excited, rapidly cleared, or less useful to the fit once the finite relief controller has saturated the acute event.

The phenomenological template used in the earlier audit was a throat-dispersion frequency map,

$$f_{\text{DEU}} = f_c \sin\left(\frac{40}{41} \frac{f_{\text{GR}}}{f_c}\right), \quad (4)$$

where f_c is an effective throat cutoff. Equation 4 is a toy map, not a Kerr-QNM calculation. The audit fitted damped-sinusoid ladders to open GW150914 H1/L1 data, with detector-specific linear amplitudes and phases, and compared four nested families: a fixed GR toy ladder, a uniform 40/41 redshift ladder, the nonlinear DEU compression ladder, and a one-parameter free frequency-scale ladder. Mode subsets were nested: $n = 0$, $n = 0, 1$, $n = 0, 1, 2$, and $n = 0, 1, 2, 3$.

Mode set	$\Delta\text{BIC}_{\text{DEU-GR}}$	$\Delta\text{BIC}_{\text{DEU-uniform}}$	$\Delta\text{BIC}_{\text{DEU-free}}$	$p_{\text{add-one}}$	Reading
$n = 0$	1.34	-0.02	-0.01	0.184	weak
$n = 0, 1$	13.55	10.95	3.52	0.020	strong
$n = 0, 1, 2$	14.15	12.11	1.99	0.020	strongest
$n = 0, 1, 2, 3$	8.87	2.50	-0.27	0.061	degraded

Table 3: Earlier GW150914 toy ringdown mode-ablation audit. Positive ΔBIC favors the DEU compression family over the listed comparator. The strongest candidate is the three-tone subset; adding the fourth tone degrades the compression preference.

The best event subset remained $n = 0, 1, 2$ under a family-level off-source correction, with $\Delta\text{BIC}_{\text{DEU-GR}} = 14.15$ and add-one $p = 1/49 = 0.0204$. Both detectors had the same sign and timing in the three-tone candidate: H1 gave $\Delta\text{BIC}_{\text{DEU-GR}} = 1.60$ with $f_c = 550$ Hz, while L1 gave $\Delta\text{BIC}_{\text{DEU-GR}} = 7.22$ with $f_c = 450$ Hz; both selected $t_0 = +0.002$ s and τ scale 1.0.

Under the two-queue law, the bridge interpretation is:

large acute load $\longrightarrow$ dense early relief, finite mode population, and possible high-tone clipping. (5)

The diagnostic is therefore not that ringdown frequencies must depart from Kerr, nor that all overtones must be absent. The diagnostic is a structured population effect: after controlling for source parameters and detector response, higher tone families should show a finite-ledger residual—suppressed amplitude, shortened usefulness, or degradation when too many tones are added. A standard smooth ladder may fit the dominant tone, while a finite-ledger map may better capture which higher tones are populated.

B Formal quantum-equivalence checks

This appendix provides formal equivalence checks under a declared DEU dictionary. The main text derives a capacity law—when chronic or acute load forces the engine to resolve backlog.

B.1 Render gate and route selection

The registry-side render gate used in the earlier closure audit begins from the sector partition

$$|S| + |I| + |G| = 81 + 40 + 16. \quad (6)$$

An unresolved quantum subgraph occupies the ternary zero slice of the spatial sector. Since $|S| = 81$, that unresolved slice has

$$|S_0| = \frac{|S|}{3} = 27. \quad (7)$$

The maximal recognized conflict between this unresolved slice and the binary macroscopic scaffold is therefore the complete bipartite boundary

$$B_{\text{crit}} = |S_0| |G| = 27 \times 16 = 432. \quad (8)$$

Equivalently, one active interface channel screens

$$C_{\text{screen}} = \left(\frac{2}{9}|S|\right) |G| = 18 \times 16 = 288 \quad (9)$$

edges, so the complete unresolved boundary requires

$$I_{\text{crit}} = \frac{432}{288} = \frac{3}{2} \quad (10)$$

active channels. Thus one channel underscreens the conflict and two channels screen it. In the closure audit this yields the timing rule

$$B_\psi < 432 \Rightarrow \psi_0 \text{ remains deferred}, \quad B_\psi \geq 432 \Rightarrow \psi_0 \text{ is forced to render.} \quad (11)$$

This is the quantum-side analogue of the acute capacity gate in the main text. It decides when the unresolved sector can no longer remain off the shared plane. It does not decide which outcome is selected.

B.2 Closure amplitude and Born cost weighting

For a render event with possible outcomes i , let $\mathcal{P}_i$ be the set of legal closure paths realizing outcome i . Each path p carries forced closure cost $C(p)$, topological-clock phase $\theta(p)$, and orientation/parity coefficient ν_p . The closure amplitude is

$$\mathcal{A}_i = \sum_{p \in \mathcal{P}_i} \nu_p e^{-\alpha C(p)/2} e^{i\theta(p)}, \quad P_i = \frac{|\mathcal{A}_i|^2}{\sum_j |\mathcal{A}_j|^2}. \quad (12)$$

The audit used one fixed inverse generative temperature, $\alpha = 0.050$, across cost weighting and tunneling. No refit is made when moving from the cost audit to the tunneling audit.

The Born-style law follows immediately in the decohered single-path limit. If outcomes i and j each have one legal closure path and no relative phase interference, Eq. 12 gives

$$\frac{P_i}{P_j} = e^{-\alpha(C_i - C_j)}, \quad \log \frac{P_i}{P_j} = -\alpha(C_i - C_j). \quad (13)$$

The standalone audit recovered the log-odds slope 0.050000000000, matching the fixed value of α to numerical precision. In this form, probability is not installed as a separate postulate after rendering; it is the normalized square of the closure ledger.

The same equation also explains why decoherence is cross-term loss. In a coherent path family,

$$|\mathcal{A}_i|^2 = \sum_{p,q \in \mathcal{P}_i} \nu_p \nu_q e^{-\alpha(C(p)+C(q))/2} e^{i(\theta(p)-\theta(q))}. \quad (14)$$

If the environment records path identity, loops with $p \neq q$ are no longer legal recognitions. The expression reduces to the cost ensemble

$$|\mathcal{A}_i|^2 \longrightarrow \sum_{p \in \mathcal{P}_i} e^{-\alpha C(p)}. \quad (15)$$

The audit therefore distinguishes coherent path addition, which scales as g^2 for g identical aligned routes, from randomized or decohered intensity addition, which scales as g .

B.3 Mach–Zehnder interference from topological-clock phase

A two-route interferometer is represented by two unresolved closure routes with phase offset

$$\Delta\theta = \theta_A - \theta_B. \quad (16)$$

At the recombiner, the two output channels are the parity classes of the closing stitch. Up to normalization, their amplitudes are

$$\mathcal{A}_B = e^{i\theta_A} + e^{i\theta_B}, \quad \mathcal{A}_D = e^{i\theta_A} - e^{i\theta_B}. \quad (17)$$

Squaring and normalizing gives the textbook Mach–Zehnder probabilities

$$P_B = \cos^2\left(\frac{\Delta\theta}{2}\right), \quad P_D = \sin^2\left(\frac{\Delta\theta}{2}\right). \quad (18)$$

The audit matched Eq. 18 to machine precision. No physical wave field is added; the wave-like interference is the phase-sensitive part of the same closure sum that becomes Born cost weighting after cross terms are removed.

B.4 Zeno freezing from subthreshold boundary screening

The closure audit contains a second Zeno-like mechanism, distinct from the equal-action pulse ladder in Section 5. Here a weak measurement creates boundary load that remains below the render threshold:

$$B(t) = |\partial\psi_0(t)| < 432. \quad (19)$$

Because the load is subthreshold, no render event occurs. However, the interface sector must screen the boundary, reducing the internal topological-clock rate:

$$\dot{\kappa}_{\text{eff}} = \dot{\kappa}_0 \left(1 - \frac{B(t)}{432}\right). \quad (20)$$

For a two-state internal oscillator, survival is then

$$P_{\text{surv}} = \cos^2\left(\frac{\theta_{\text{eff}}}{2}\right). \quad (21)$$

The audit drove the unmeasured oscillator through $\theta = \pi$, giving survival near zero, while rapid subthreshold interrogation capped the boundary at $431 < 432$, caused no stitch, and raised survival to 0.9998977384. The interpretation is not projection by an observer. It is subthreshold interface screening: repeated near-threshold contacts consume clock capacity without forcing the unresolved state to render.

The main text's engine-native Zeno ladder and Eq. 20 should be read together but not conflated. The main text shows that equal total action, when temporally subdivided, can produce zero topological relief events in the frozen load engine. Equation 20 shows how a deferred closure state can freeze internally under repeated subthreshold boundary screening. Both are finite-ledger suppression phenomena; they operate at different levels of the DEU dictionary.

B.5 Tunneling as shortcut economics

A classically forbidden barrier is represented as a high generative cost for rendering a step-by-step spatial path. The alternative is to defer spatial realization, traverse the unresolved graph, and restitch on the far side:

$$C_{\text{shortcut}} = C_{\text{defer}} + C_{\text{far stitch}}. \quad (22)$$

The route-selection criterion is the economic inequality

$$C_{\text{shortcut}} < C_{\text{barrier}}. \quad (23)$$

For the textbook square-barrier comparison, introduce the dimensionless WKB action proxy

$$S_{\text{WKB}} = \int_0^L \sqrt{V(x) - E} dx \quad (24)$$

and the fixed cost map

$$C_{\text{tunnel}} = \frac{2}{\alpha} S_{\text{WKB}}. \quad (25)$$

Then Eq. 12 gives

$$T_{\text{DEU}} \propto e^{-\alpha C_{\text{tunnel}}} = e^{-2S_{\text{WKB}}}. \quad (26)$$

The audit recovered the expected exponential slopes: -2 versus S_{WKB} , -4 in the width sweep with $V - E = 4$, and -3 versus $\sqrt{V - E}$ with $L = 1.5$. This is a formal substitution, not a calibrated map from DEU cost units to laboratory electron-volts. A physical tunneling claim would require that additional cost-to-action calibration.

B.6 Bell correlations from joint ancestral closure

For entangled pairs, the DEU dictionary treats the two particles as spatially separated in the rendered plane but adjacent through an unresolved shared ancestor. A measurement does not send a signal from Alice to Bob. Rather, the render event forces a joint consistency closure on the shared unresolved parent. The probability law is therefore joint from the start:

$$P(a, b|x, y), \quad a, b \in \{-1, +1\}. \quad (27)$$

For detector offset

$$\delta = \theta_x - \theta_y, \quad (28)$$

the closure audit uses two ancestral paths per joint outcome. Same-outcome branches are anti-oriented,

$$\mathcal{A}_{\text{same}} \propto 1 - e^{i\delta}, \quad (29)$$

while opposite-outcome branches are co-oriented,

$$\mathcal{A}_{\text{opp}} \propto 1 + e^{i\delta}. \quad (30)$$

Normalizing squared closure amplitudes gives

$$P(a, b|x, y) = \frac{1}{4} [1 - ab \cos(\theta_x - \theta_y)]. \quad (31)$$

The correlation function is therefore

$$E(x, y) = \sum_{a, b} ab P(a, b|x, y) = -\cos(\theta_x - \theta_y). \quad (32)$$

At the standard CHSH settings,

$$S_{\text{CHSH}} = 2\sqrt{2}. \quad (33)$$

At the same time, the local marginals remain uniform:

$$P(a = +1|x, y) = P(a = -1|x, y) = \frac{1}{2}, \quad P(b = +1|x, y) = P(b = -1|x, y) = \frac{1}{2}. \quad (34)$$

The audit matched the joint probability and correlation laws to numerical precision and found marginal deviations from 1/2 at the level of floating-point roundoff. The result is a formal no-signaling Bell substitution inside the closure harness. It is not a claim that the main two-queue engine has derived Hilbert-space quantum mechanics; it says that the same registry/closure architecture has a joint-ancestor route law with the textbook Bell shape.

B.7 Status of the quantum substitutions

The quantum-side bridge therefore has two layers. The first is timing: Eq. 8 and the main text's acute render gate say when unresolved or overloaded structure is forced to resolve. The second is route selection: Eq. 12 says how legal outcomes are weighted once resolution occurs. In compact form,

$$432 \text{ or the live capacity gate decides when; } \quad \mathcal{A}_i \text{ decides what.} \quad (35)$$

C Formal relativity-equivalence checks

This appendix adds the relativity-side textbook substitutions from the earlier DEU relativity audit. They should be read in the same evidential class as Appendix B: formal equivalence under a declared dictionary. Their role in the paper is to show that the same finite-ledger language that now derives chronic and acute DEU load response also has the right algebraic shape for the standard clock laws of special and Schwarzschild relativity.

The substitutions use two graph-native rules. First, update operations are counted and counts add over a coarse epoch:

$$N_\tau + N_x + D_{\text{unres}} = B_{\text{vac}}, \quad (36)$$

where N_τ is spent on internal self-update, N_x on lateral translation, D_{unres} on unresolved ledger demand, and B_{vac} is the matched vacuum budget. In the present two-queue language, D_{unres} should be read as the total unresolved demand relevant to the clock: forced defect backlog plus standing vacuum demand when both are active. Second, coarse metric intervals are root-cost readouts of operation counts:

$$\Phi^2 = \frac{N_\tau}{B_{\text{vac}}}, \quad v^2 = \frac{N_x}{B_{\text{vac}}}. \quad (37)$$

The spatial instance of this rule is the audited metric adapter $\ell_f = 3^{-k_f/2} = A_f^{1/2}$; extending it to the temporal channel is the explicit dictionary assumption of this appendix.

C.1 Textbook substitution 1: inertial time dilation

In vacuum, $D_{\text{unres}} = 0$. Dividing Eq. 36 by B_{vac} and substituting Eq. 37 gives

$$\Phi^2 + v^2 = 1, \quad \Phi = \sqrt{1 - v^2}, \quad \gamma = \frac{1}{\Phi}. \quad (38)$$

Thus the Lorentz factor is the reciprocal of remaining internal update capacity. The important point is not that the finite graph has already derived continuum observer clocks in full; it is that, once a Lorentzian scaling regime exists, the frontier allocation law supplies the usual time-dilation form by operation accounting. Linear metrology, by contrast, would give a Galilean drag law $\Phi = 1 - v$, i.e. first-order time dilation. The square-root readout is therefore not a cosmetic normalization: it is the metric exponent needed for the observed quadratic clock law and is independently the exponent used by the audited spatial adapter.

C.2 Textbook substitution 2: Schwarzschild clock laws

With unresolved ledger demand present, define

$$\Omega_{\text{op}} = \frac{D_{\text{unres}}}{B_{\text{vac}}}. \quad (39)$$

Equations 36–37 give the lapse allocation law

$$\Phi_{\text{total}}^2 + v^2 = 1 - \Omega_{\text{op}}, \quad \frac{d\tau}{dt} = \sqrt{1 - \Omega_{\text{op}} - v^2}. \quad (40)$$

The Schwarzschild textbook formulas follow from a single dictionary entry,

$$\Omega_{\text{op}}(r) = \frac{r_s}{r}. \quad (41)$$

At rest,

$$\frac{d\tau}{dt} = \sqrt{1 - \frac{r_s}{r}}, \quad (42)$$

which is the Schwarzschild static clock law. For tangential coordinate motion,

$$\frac{d\tau}{dt} = \sqrt{1 - \frac{r_s}{r} - v_{\text{coord}}^2}. \quad (43)$$

Equivalently, writing the local speed as u so that $v_{\text{coord}}^2 = (1 - r_s/r)u^2$, Eq. 43 factorizes as

$$\frac{d\tau}{dt} = \sqrt{1 - \frac{r_s}{r}} \sqrt{1 - u^2}. \quad (44)$$

The weak-field expansion is also the standard one:

$$\frac{d\tau}{dt} \simeq 1 - \frac{\Omega_{\text{op}}}{2} - \frac{v^2}{2} = 1 + \Phi_N - \frac{v^2}{2}, \quad \Phi_N = -\frac{\Omega_{\text{op}}}{2} = -\frac{GM}{r}. \quad (45)$$

Finally, the horizon is the saturation surface of the ledger: at $\Omega_{\text{op}} = 1$ and $v = 0$, the rest clock has $d\tau/dt = 0$ because the unresolved demand consumes the entire matched vacuum frontier budget, $D_{\text{unres}} = B_{\text{vac}}$.

References

- [1] J. Merwin, *Temporal Necessity in Relational Mathematical Realism: A Godelian Argument Against the Block Universe*, [ai.viXra.org:2602.0071](https://arxiv.org/abs/2602.0071) (2026).
- [2] J. Merwin, *A Finite Carrier Vacuum and Discrete Collision Fractures in the Distinction Engine*, [ai.viXra.org:2604.0103](https://arxiv.org/abs/2604.0103) (2026).
- [3] J. Merwin, *Incompleteness as Expansion: A Self-Reference Mechanism for Shared Reality, the Arrow of Time, and the Open Future*, [ai.viXra.org:2606.0082](https://arxiv.org/abs/2606.0082) (2026).
- [4] J. Merwin, *The Distinction Engine: Recovery of the 137-Element Relational Registry from a Single Logical Axiom*, [ai.viXra.org:2604.0057](https://arxiv.org/abs/2604.0057) (2026).
- [5] J. Merwin, *The (41/40 Graph Index in the Distinction Engine Universe: Discrete-Substrate Resonance, CMB Amplitude Closure, and Cross-Observable Clock Projections*, [ai.viXra.org:2605.0023](https://arxiv.org/abs/2605.0023) (2026).
- [6] H. Hoppe et al., *Mesh optimization*, SIGGRAPH (1993).