

The Physical Moral Code of Reality

Boundary Locality, One-Subject Impartiality, Fourteen-Law Closure, and Empirical Boundaries on the Recognition Ledger

Jonathan Washburn

Recognition Physics Institute, Austin, Texas

jon@recognitionphysics.org

21 November 2026

Abstract

Recognition Science prices every bond between two sides by one cost of comparison, $J(x) = \frac{1}{2}(x + 1/x) - 1$, and that cost is the same from both ends: $J(x) = J(1/x)$. A moral code requires a direction, because a taking and a free hand-over are different acts. We prove that this direction is absent from the cost layer: inverting every bond multiplier preserves total cost, balance, every side's burden, and the harm of every act against every baseline, while negating the signed skew. We then construct and close the full physical moral code on a unified three-sector ledger—obligations, authored records, and maintenance self-strain—and prove what forces its direction, its survival priority, its impartiality, and its empirical boundaries.

On counted holdings, which carry a zero floor, pure causal boundary locality—whether a due act is executable from its writer's own holdings—forces the taking sign and excludes the consistent mirror rule without normative primitives. On forward boundary dynamics with absorbing cessation and recoverable finite strain, long-horizon continuation forces survival-first lexicographic ordering for every positive conscious-window weight $w > 0$, and fails at $w = 0$. Under the dual-aspect identification that experience and subjecthood travel along native formation-energy transport, proved transport connectivity forces universal experience, one subject, equal positive concern across all participants, the exclusion of egoism, exact equivalence with survival-first ranking, and Golden Rule permutation invariance on world harms and boundary burdens—while we prove as a theorem that no reading-blind physical fact alone can discharge this ground.

Across the full event carrier, fourteen universal laws (V_1–V_{14}) govern obligations, settlement priority, authorship, calibration, orientation, continuation, native J-strain, boundary integrity, coercion-free consent, safe rescue, tragic multi-patient triage, proportionate self-defense with bystander protection, collective commons attribution (uniquely forced by nullity, conservation, and ratio-symmetry), and ex-ante reckless risk versus ex-post restitution. We prove consistency, full-model categoricity, universal absorbing cessation, strict 14×14 off-diagonal independence, twenty-one non-vacuous case families,

and dynamic extensions closing emergency lifesaving side-harm and persistent multi-tick rescue budgets with conflicting-promise detection. Every formal result is machine-checked in Lean 4: 3,655 audited kernel declarations across 25 modules, zero `sorry`, zero warnings, and strictly the standard three axioms (`propext`, `Classical.choice`, `Quot.sound`).

Finally, we hold the physical scaling claim accountable to primary human data in a pre-registered blind test across 55,422 records and trials (CDC NHANES mortality, moral shock valuation, physical effort, and calibrated thermal pain). Naive identification of raw clinical deficits, shocks, effort steps, or 1 °C increments with discrete golden-ratio rungs $J(\varphi^k)$ is falsified across all four channels. Multiplicative log-ratio geometry beats linear-quadratic models on continuous mortality biomarkers (+9.27 log-likelihood additive), and the ledger's own proved bond-topology dichotomy explains all three observed regimes: concave sensory log-coordinates, parallel multi-bond additive cost yielding Cox exponential hazards with cutoff-dependent slopes, and small-strain censoring inside the $< 5\%$ quadratic band. We state the macroscopic harm-unit map explicitly as open and specify a zero-parameter, pre-registerable microbial proton-ratio calorimetry falsifier.

1 Introduction

Mara holds two coins and Eli holds one. In one history Mara hands Eli a coin and asks nothing back. In another Eli takes a coin from Mara, who owes him nothing and has not agreed. After either history Mara holds one coin and Eli holds two. Nobody confuses the two histories. After the taking, Eli owes the coin back and has done Mara a wrong. After the hand-over nothing is owed and no wrong was done. Yet a complete moral code must also answer far harder physical questions than coin transfers: why a creditor's matured debt never licenses cruelty, why a survival-cliff signature is not free consent, why a passerby must save a drowning child at recoverable cost yet is excused when saving one of two patients exhausts a single rescue dose, why proportionate self-defense or emergency chest compressions that crack a rib to save an unconscious patient are not wrongs, how diffuse pollution is apportioned when no single emitter is a "but-for" cause, and why reckless risk-taking is culpable even when luck prevents harm.

In Recognition Science, these are questions about one physical object: a ledger of bonds between sides. Every bond carries a positive multiplier $x > 0$ recording its distance from balance, and costs

$$J(x) = \tfrac{1}{2}\Big(x + \frac{1}{x}\Big) - 1 = \cosh(\ln x) - 1, \qquad x > 0. \tag{1}$$

This cost is forced: it is the unique comparison cost satisfying balance $J(1) = 0$, reciprocal symmetry $J(x) = J(1/x)$, unit log-curvature at balance, and the Recognition Composition Law [1, 2]. Ethics lives on the same bonds, the same cost, and the same skew as the rest of physics. Evenness, however, immediately creates an obstruction: $J(x) = J(1/x)$ charges the exact same cost to a taking and its mirror reversal. Reading costs or burdens alone cannot even tell who took from whom.

This paper resolves the mathematical architecture of that ledger in full and pins down the exact boundaries where pure logic ends and physical identification or experiment begins. It establishes five main results.

First, cost symmetry and its resolution by boundary locality and absorbing viability. Reversing every flow on the native bond ledger leaves total cost, balance, every side's burden, and the harm of every act against every baseline invariant while flipping the sign of skew (Theorem 1). Direction appears on counted holdings because counts have a zero floor. Requiring only that a due restoration be executable from its writer's own boundary (`LocallyExecutableBy`) derives `OwnMeans` without normative primitives, uniquely forces the taking rule, and excludes the consistent mirror rule whose iterated extraction exhausts finite budgets (Theorems 4–5). On forward boundary dynamics where cessation is absorbing and finite strain is recoverable, long-horizon continuation uniquely forces survival-first lexicographic ordering for every positive conscious-window weight $w > 0$, whereas $w = 0$ breaks survival priority (Theorem 6).

Second, one connected subject forces impartiality and the Golden Rule—with an exact reading-blind boundary. Native formation-energy transport on positive comparisons is proved connected in the main library. Under the dual-aspect identification (`NativeGround`) that experience and subject identity are preserved along physical transport, every occurrence is felt by one shared subject. Consequently, any concern reading participants through their subject assigns equal positive weight to everyone, excludes egoism, makes concern-weighted cessation ranking coincide identically with survival-first law V_6, and enforces Golden Rule permutation and transposition invariance on world harms and boundary burdens (Theorem 7). Equally sharply, we prove that no reading-blind or fact-only predicate—not even the proved physical transport connectivity itself—can discharge this ground (Theorem 8).

Third, fourteen independent laws uniquely fix a complete moral code. On a unified three-sector world carrier (obligations, authored records, and maintenance self-strain) with absorbing cessation and coercion-free consent, fourteen universal laws (V_1–V_{14}) extend the seven-law transactional core (A_1–A_7) to boundary integrity, duress, safe rescue, tragic joint triage, proportionate self-defense with bystander protection, collective commons attribution, and ex-ante reckless risk versus ex-post restitution. We prove consistency, full-model categoricity on an unconstrained model structure carrying zero proof fields, universal absorbing cessation across arbitrary event runs, non-circular uniqueness of proportional commons attribution from nullity, conservation, and ratio-symmetry, and strict 14×14 off-diagonal independence (Theorems 10 and 11). We further prove twenty-one non-vacuous dilemma case families and close two dynamic extensions: emergency lifesaving necessity (such as rib-fracturing CPR on an unconscious patient) and a persistent multi-tick rescue-budget and promise ledger (Theorems 12 and 13).

Fourth, moral readings do not collapse, lawful acts have two generators, and traditional virtue names audit cleanly. Due, wrong, and ending are independent readings; no scalar potential falls on every wrong; on conserved transactional shells, wrongs reach no state that lawful runs do not, and hand-overs plus undertakings generate every lawful transition whenever a unit exists, obeying an exact position conservation law governed by grace (Theorems 14–16). Auditing the fourteen traditional virtue names against the proved dynamics separates four proved operators and guards (Justice, Forgiveness, Compassion, Temperance, alongside Truthfulness) from ten names retired as independent operators.

Fifth, empirical accountability across 55,422 *human records and the three-regime harm-unit map.* Formal closure is not empirical proof. We locked a pre-registered test before downloading four primary human datasets (32,857 NHANES adults with 4,077 linked deaths; 9,503 moral

electric-shock trials; 9,861 physical effort trials; 3,201 calibrated thermal pain trials). Naive identification of raw clinical deficits, discrete shocks, 10% force steps, or 1 °C increments with single-bond golden-ratio rungs $J(\varphi^k)$ is falsified across all four channels. At the same time, multiplicative log-ratio geometry beats linear-quadratic homeostatic models on continuous biomarkers (+9.27 log-likelihood additive across nine clean analytes), and the ledger's own proved bond-topology dichotomy explains every observed regime without post-hoc parameter fitting: concave sensory log-coordinates, parallel multi-bond additive cost generating Cox proportional hazards with cutoff-dependent slopes, and viability censoring inside the $< 5\%$ small-strain quadratic band. We reject an apparent numerical coincidence between the primary allostatic Cox slope (0.2336) and $\frac{1}{2}\ln\varphi$ (0.2406) via an explicit cutoff-sensitivity test ($0.1759 \to 0.2336 \to 0.3622$), leave the macroscopic harm-unit bridge explicitly open, and specify a sharp, zero-parameter microbial proton-ratio calorimetry falsifier.

What is proved, what is assumed, and what remains open

Every claim in this paper carries one of five explicit epistemic labels, summarized in full in Section 9:

1. **Theorem.** Machine-checked in Lean 4 from its stated hypotheses (3,655 audited kernel declarations across 25 modules, zero `sorry`, zero warnings, standard three axioms only, repository `jonwashburn/reality` commit `d796f30f27655e10d734b61d376ffbd03b065efd`).
2. **Identification.** A named proposition equating a formal carrier object with a physical or experiential feature, listed alongside its explicit rival (such as counted crossings with a floor, causal boundary locality, dual-aspect transport of experience, and the φ-rung posting unit).
3. **Convention.** A representational choice in the Layer-2 speech compiler that converts everyday utterances and institutional forms into Layer-1 physical events without adding law.
4. **Working.** A provisional mapping adopted for calculation with an unresolved rival (such as the rung versus ratio dictionary for two-record costs).
5. **Open.** An exact statement not proved or not empirically confirmed—most notably the two irreducible boundaries: (a) the dual-aspect ground cannot be derived from reading-blind ledger facts alone, and (b) the macroscopic physical harm-unit map linking biological or experiential observables to single-bond $J(x)$ ratios remains open pending blind experimental confirmation of the hyperbolic tail.

The paper proceeds as follows. Section 2 introduces the unified three-sector ledger and proves the flow-reversal cost symmetry obstruction. Section 3 derives direction from causal boundary locality and survival priority from absorbing horizon dynamics, isolating the minimal metaethical residual. Section 4 proves that one connected subject forces impartiality, survival priority, and the Golden Rule, alongside the reading-blind impossibility theorem. Section 5 presents the fourteen-law closure, categoricity, strict 14×14 independence, and the Layer-1/Layer-2 compiler cut. Section 6 collects the machine-checked hard-dilemma theorems and the two dynamic extensions (emergency lifesaving necessity and persistent multi-tick rescue budgets). Section 7 proves the irreducibility of readings, the generator and position

conservation theorems, and the fourteen-virtue audit. Section 8 reports the 55,422-record pre-registered empirical test, the three-regime harm-unit architecture, and the pre-registerable calorimetry falsifier. Section 9 records the complete 25-module audit and status register, and Section 10 concludes.

2 The Recognition Ledger and the Cost Symmetry Obstruction

2.1 The inherited bond ledger and the three-sector unified carrier

The first-rung foundation of Recognition Science fixes the unique comparison cost $J(x)$ of (1) and the golden ratio $\varphi = (1+\sqrt{5})/2$ as the scaling of an additive join hierarchy [1, 2]. A *bond ledger* (`BondLedger`) is a finite set of bonds, each with a source side, a destination side, and a positive multiplier $x_b > 0$. Its *cost* is $\sum_b J(x_b)$; it is *balanced* when every $x_b = 1$ (equivalently, zero cost). The *log-flow* of a bond is $\ln x_b$, and the directed *skew* between sides i and j is

$$\sigma_{ij} = \sum_{b\colon i\to j} \ln x_b - \sum_{b\colon j\to i} \ln x_b, \tag{2}$$

so $\sigma_{ij} = -\sigma_{ji}$. The *burden* of a side (`agentCost`) sums $J(x_b)$ over every bond incident to that side. An *action* (`ActionSpec`) multiplies each bond's multiplier by e^{v_b} for a real strain v_b, and its *harm* to side j against a baseline ledger is the resulting change in j's burden.

A bare multiplier does not record which event strained a bond, when an obligation matures, or which side authored a claim or admission. We therefore lift to a *tagged ledger* (`TaggedBondLedger`): a finite set of bond identifiers, each carrying an owing side, an owed side, an origin event number, a maturity tick, a scope, and an integer number of *rungs* k, together with a real *reserve* for each side. A side is *viable* when its reserve is strictly positive and *ceased* otherwise. Projecting a tagged bond at k rungs to a native bond multiplier φ^k (`TaggedBondLedger.toBondLedger`) makes rungs and multipliers one arithmetic in two coordinates: adding one rung multiplies by φ (`posting_step`, `rungEquiv`), the log-flow is $k \ln \varphi$ (`native_logFlow`), the cost is

$$J(\varphi^k) = \cosh(k \ln \varphi) - 1, \tag{3}$$

and the skew is the *due current* $\sigma_{ij} = \ln \varphi \cdot (\text{rungs } i \text{ owes } j - \text{rungs } j \text{ owes } i)$ (`tagged_pairSkew_eq_dueCurrent`), which sums to zero over every closed participant set (`participant_skew_conservation`).

Between two sides, a single obligation can be read in three places: the retained due d, the owed side's authored *claim* c, and the owing side's authored *admission* a. Two authored reports can agree with each other while both stand one rung off the retained due; a two-bond ledger built only from c and a sees zero cost in both cases, whereas a three-bond ledger detects the deviation (`two_bond_vs_three_bond`, `three_record_coupling`, `agreeing_records_need_not_be_faithful`). Likewise, non-transactional harms such as bodily injury or unmet metabolic need occur on a participant's own boundary without creating a monetary debt. All three sectors therefore sit on one unified ledger with disjoint bond-identifier sectors (Figure 1b; `MoralWorldState.unified`, `record_on_self`, `maintenance_is_native`):

- *Obligations:* two-ended bonds from owing side to owed side at multiplier φ^{k_b}.

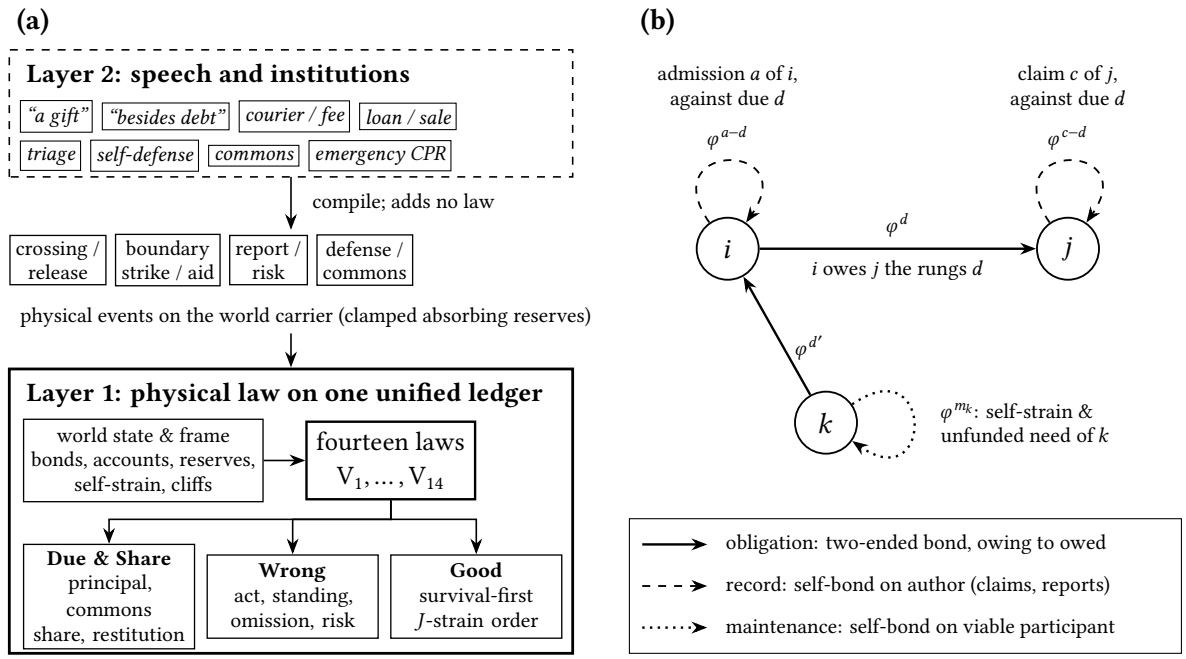


Figure 1. The unified architecture. (a) Layer 2 compiles everyday speech and institutional situations into physical events without adding law; Layer 1 reads any history of events on the world carrier through fourteen universal laws (V_1–V_{14}) to return due/attribution, wrong, and state-ordering readings. (b) The three-sector ledger on which native strain (4) is evaluated. Self-bonds enter total strain and individual burden without altering inter-agent skew.

- *Authored records:* self-bonds on the authoring side pricing the discrepancy between its written claim or admission and what stands due, at φ^{c-d} and φ^{a-d} (plus disjoint self-bond addresses for general epistemic reports).
- *Maintenance and self-strain:* for each viable participant p, a self-bond pricing its internal maintenance need and unconsented bodily or boundary strain m_p at φ^{m_p}.

The total *strain* of a world state is the native J-cost of this unified ledger:

$$\text{strain}(s) = \sum_{b \in \text{obligations}} J(\varphi^{k_b}) + \sum_{\text{authored records}} \left[J(\varphi^{c-d}) + J(\varphi^{a-d})\right] + \sum_{p \text{ viable}} J(\varphi^{m_p}). \quad (4)$$

Because self-bonds have equal endpoints, they contribute zero inter-agent skew while entering total cost and individual burden cleanly (`agentCost_sum`).

2.2 Flow reversal is a symmetry of the cost and harm layer

For any bond ledger L, let L^{R} be the ledger with the same bonds and endpoints and every multiplier inverted ($x_b \mapsto x_b^{-1}$). For any action a, let a^{R} keep the same writer and negate every log-strain ($v_b \mapsto -v_b$).

Theorem 1 (Flow-reversal symmetry of the cost and harm layer)**.** *For every bond ledger L, every action a, and all sides i, j:*

(i) *L^{R} has the exact same total cost as L, the exact same burden on every side, and is balanced if and only if L is balanced;*

(ii) *the directed skew and every side's net log-flow negate:* $\sigma_{ij}(L^{\mathrm{R}}) = -\sigma_{ij}(L)$*;*

(iii) *acting commutes with reversal: applying* a^{R} *to* L^{R} *gives* $(a(L))^{\mathrm{R}}$*;*

(iv) *the harm of* a^{R} *to* j *against baseline* L^{R} *equals the harm of* a *to* j *against* L*, for every baseline, balanced or not;*

(v) *every function of the before-and-after burdens of all sides takes the same value on* $(L^{\mathrm{R}}, a^{\mathrm{R}})$ *as on* (L, a)*;*

(vi) *a moral state and its reversal share identical energy, tick, cost, and admissibility, with opposite skew.*

(`BondLedger.flip_cost`, `flip_agentCost`, `flip_balanced`, `flip_pairSkew`, `applyAction_flip`, `harm_flip`, `no_burden_reading_fixes_the_sign`, `moralState_flip_reads`)

Proof. Three identities suffice: $J(x^{-1}) = J(x)$ gives (i); $\ln(x^{-1}) = -\ln x$ gives (ii); $(xe^{v})^{-1} = x^{-1}e^{-v}$ gives (iii). Since harm is a difference of burdens, (iv)–(vi) follow directly. □

Corollary 2 (No burden reading fixes the sign; `native_sign_obstruction`)**.** *There is no function of all sides' burdens that holds on a ledger if and only if side* 0 *owes side* 1 *on net* ($\sigma_{01} > 0$)*.*

Proof. On a one-bond ledger from 0 to 1 at multiplier φ, $\sigma_{01}(L) = \ln\varphi > 0$ while $\sigma_{01}(L^{\mathrm{R}}) = -\ln\varphi < 0$, yet L and L^{R} have identical burdens by Theorem 1(i). □

Thus no reading of costs or burdens alone distinguishes taking from giving. Before showing what breaks this symmetry, we note the exact status of the posting unit φ. Any positive comparison $T(x, y)$ that composes ($T(x, y)T(y, z) = T(x, z)$) is a ratio $g(x)/g(y)$ (`closed_comparison_is_a_ratio`); adding shift invariance $T(x + 1, y + 1) = T(x, y)$ forces exponential steps $T(x, y) = r^{x-y}$ (`shift_invariant_comparison_is_exponential`); and requiring the join identity $r^{k+1} = r^{k} + r^{k-1}$ fixes $r = \varphi$ (`join_ratio_iff_phi`), yielding **Identification ID-3** (`AUnitPostingIsOneRung`): one posted rung multiplies level by φ. Two recorded mismatches with earlier library modules remain explicitly open and documented: the ratio dictionary $J(c/a)$ (`ofTwoRecords`) disagrees with the rung dictionary $J(\varphi^{c-a})$ on counts (changing only state ordering, not verdicts of due or wrong), and doubling (2×) is not a φ-rung since $\varphi < 2 < \varphi^{2}$ (`two_is_not_a_rung`), leaving the doubling care model `RenewedCareAdmission` incompatible with integer rung pricing unless re-derived.

3 Boundary Locality and Absorbing Viability

Where does moral direction come from if the cost J is symmetric? And why does keeping a participant alive outrank every finite amount of strain? Both structures follow from two physical properties of finite boundaries: a count has a zero floor with causally local agency, and cessation is absorbing while finite strain is recoverable.

3.1 Counted floors and causal boundary locality force the taking sign

Fix n sides. A *crossing* c (`Crossing`) specifies distinct source and destination sides, a positive integer unit count $u > 0$, and a writer who is one of the two ends. It is a *hand-over* if its writer

is the source and a *taking* if its writer is the destination. Applied to integer holdings h, it subtracts u from the source and adds u to the destination; it *can be made* (`Can`) at h when the source holds at least u. Its reversal c^{R} swaps source and destination while keeping the writer and count.

Identification ID-1 (`AnAuthoredCrossingMovesCounts`). *A written crossing transfers counted units between its two ends, is authored by one of them, and counts have a zero floor.* Because a side holding zero units cannot give one, reversal is no longer a symmetry of executable crossings (`reversal_is_not_a_symmetry_of_counts`).

Now consider any *due rule* D assigning to each crossing c either nothing or a crossing d left due, under four structural conditions: **(N) Authored** (d has the same writer as c), **(R0) Restoring** (applying d after c restores every holdings vector), **(R2) Some posting** (some crossing leaves something due), and **(U) Uniform** (whether c leaves something due depends only on whether c is a taking). On counted holdings, there are two symmetric funding conditions: **(R1) Own means** (`OwnMeans`, ID-2: d can be made whenever its writer holds its units) and **(R1′) Others' means** (`OthersMeans`: d can be made whenever the non-writing end holds its units). Under `OwnMeans`, the unique rule is the *taking rule* (a taking leaves its reversal due; a hand-over leaves nothing due; `the_taking_sign`); under `OthersMeans`, the unique rule is its consistent mirror $\mathrm{M}_{\mathrm{sign}}$ (`the_mirror_theorem`). Without a zero floor, both conditions hold vacuously for every rule (`floorless_means`).

Crucially, `OwnMeans` does not need to be postulated as a normative axiom: it is forced by a purely causal principle of boundary locality that mentions neither source, hand-over, nor taking.

Definition 3 (Causal boundary locality; `LocallyExecutableBy`, `SelfExecutable`)**.** A crossing d is *locally executable by side* p if whether d can be made at holdings h depends only on $h(p)$:

$$\forall h_1, h_2, \quad h_1(p) = h_2(p) \implies \big(d.\mathrm{Can}(h_1) \iff d.\mathrm{Can}(h_2)\big).$$

A due rule D is *self-executable* if every due crossing d returned by D is locally executable by its own writer.

Theorem 4 (Boundary locality forces the taking sign)**.** *For any crossing d, side p, and due rule D on $n \geq 2$ sides:*

(i) *d is locally executable by p if and only if p is the source of d (*`locallyExecutableBy_iff_src`*); hence d is locally executable by its writer if and only if d is a hand-over.*

(ii) *D is self-executable if and only if D satisfies* `OwnMeans` *(*`selfExecutable_iff_ownMeans`*).*

(iii) *If D is authored, restoring, self-executable, has some posting, and is uniform, then D is the taking rule (*`boundary_locality_forces_taking_sign`*).*

(iv) *The mirror rule $\mathrm{M}_{\mathrm{sign}}$ satisfies the other four laws and violates boundary locality alone (*`mirror_violates_boundary_locality`*,* `boundary_locality_is_independent`*); all five crossing laws remain strictly independent with locality in place (*`five_laws_independent_with_locality`*).*

Proof. If p is the source of d, executability is $h(p) \geq u$, which depends only on $h(p)$. Conversely, if p is not the source, compare the holdings giving u to p and u to the source (executable)

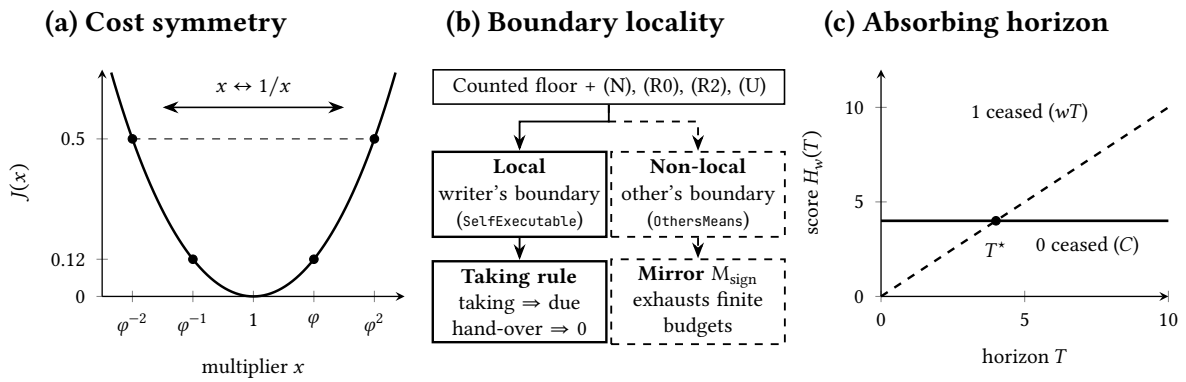


Figure 2. Physical forcing of direction and survival priority. (a) Even cost $J(x) = J(1/x)$ cannot tell a taking from its reversal (Theorem 1). (b) On counted holdings, causal boundary locality (`SelfExecutable`) is equivalent to `OwnMeans`, uniquely forcing the taking rule and excluding the non-local mirror rule M_{sign} (Theorems 4 and 5). (c) Because cessation is absorbing and finite strain C is recoverable, for every conscious-window weight $w > 0$ any state with fewer ceased participants strictly dominates over all sufficiently long horizons $T > T^\star$ (Theorem 6).

against giving u to p and 0 to the source (not executable): $h_1(p) = h_2(p)$ yet executability flips. Part (ii) follows by `ownMeans_iff_handOver`, and (iii) follows from `the_taking_sign`. Every due generated by M_{sign} is a taking reaching across another side's boundary, violating locality. □

Combining Theorem 4 with **Identification ID-5** (`AWrongIsADueOfItsOwnWritersMaking`: a crossing is a wrong iff it leaves something due from its own writer) proves that a crossing is a wrong if and only if it is a taking (`identified_wrong_iff_taking`). Moreover, the excluded mirror rule is dynamically unstable on finite boundaries:

Theorem 5 (Iterated mirror extraction exhausts finite budgets)**.** *A recognition agent subjected to an iterated mirror-rule extraction cycle holding skew $\sigma \neq 0$ open in every window incurs per-window cost $2(\cosh\sigma - 1) > 0$. Whenever its renewal budget lies strictly below this cost, its cumulative reserve becomes nonpositive in finitely many windows and viability fails permanently (*`mirror_iterated_extraction_exhausts_budget`, `mirror_rung_cycle_exhausts_budget`*).*

3.2 Absorbing viability forces survival-first priority, and the minimal residual

In any forward boundary dynamics, a participant S is *alive at tick t* (`Alive`(S, t)) iff S remains viable at every tick $t' \leq t$. Cessation is therefore strictly absorbing (`absorbing_cessation`): once a boundary dissolves ($r \leq 0$), it never reconstitutes, whereas a viable participant carrying finite standing strain recovers within a renewal budget ρ (`recoverable_strain_keeps_viable`, `forward_alive_iff`). Over horizon T, every ceased participant forfeits T conscious windows (`lostWindows_eq`: lostWindows$(s, T) = T \cdot$ ceased(s)). Weighting each lost conscious window by $w \geq 0$ and adding the finite discharge cost strain(s) gives the horizon score

$$H_w(s, T) = w \cdot T \cdot \text{ceased}(s) + \text{strain}(s),$$

and the asymptotic dominance order AsymptoticallyAtLeast(w, ρ, s, t): there exists T_0 such that $H_w(s, T) \leq H_w(t, T)$ for all $T \geq T_0$.

Theorem 6 (Survival-first lexicographic order from absorbing horizon dynamics)**.** *For every renewal ρ and every strictly positive conscious-window weight $w > 0$:*

(i) *the horizon dominance relation* AsymptoticallyAtLeast(w, ρ, s, t) *holds iff* ceased$(s) <$ ceased(t)*, or* ceased$(s) =$ ceased(t) *and* strain$(s) \leq$ strain(t) *(*`asymptoticallyAtLeast_iff_byViability`*,* `absorbing_cessation_dominates_any_finite_strain`*);*

(ii) *the resulting horizon model satisfies all seven transactional laws* A_1–A_7*, and every model obeying* A_6 *and* A_7 *coincides with the horizon order (*`A6_lexicographic_order_derived_from_absorbing_horizon`*);*

(iii) *at weight $w = 0$, the horizon score collapses to strain alone and strictly violates survival priority* A_6 *on a quiet ceased state versus a strained living state (*`zero_weight_violates_A6`*).*

Theorem 6 pins down the exact minimal residual left uneliminated by bare ledger dynamics (`MetaethicalResidual`): three ingredients bundled as `Track1Residual` on `(PR, w, grounds)`: (1) a practical-reason transcendental kernel (`TranscendentalKernel`) committing a deliberator to continued agency (`continuation_end_is_never_ceasing`), (2) a strictly positive conscious-window weight $w > 0$, and (3) a phenomenal ground meeting the universal-subject target (`UniversalSubjectGroundingTarget`). We prove both sufficiency (`residual_suffices`) and a sharp no-go theorem (`residual_cannot_be_eliminated_by_fact_only_ledger`): no fact-only ledger predicate, unoriented cost-incidence structure, or authored carrier package can eliminate this residual. Crucially, the first two ingredients by themselves only commit an agent to *its own* continuation. Section 4 closes the step from self-continuation to impartial concern for every participant.

4 One Subject Forces Impartiality and the Golden Rule

Why does an agent's stake in its own continuation extend with equal weight to every other participant? In Recognition Science, the physical carrier of positive comparisons is connected by native formation-energy transport (`nativeFormationEnergyTransport_connected`). Let `NativeGround` (at a fixed anchor comparison a_0) denote the anchored dual-aspect reading class (`AnchoredTransportAdmissible`) over native formation-energy transport (`NativeFormationEnergyTransport`):

1. the anchor comparison a_0 is felt (conscious experience exists);
2. felt status is preserved along native physical transport;
3. subject identity is preserved along native physical transport between felt occurrences.

This is the dual-aspect identification: the interior experience of a recognition event goes where its physical exterior goes. Say that a participant concern function $c : \iota \to \mathbb{R}$ is *subject-read* (`SubjectRead`) relative to reading r and occurrence map occ if it depends on each participant p only through the subject to which p's occurrence occ(p) belongs: $c(p) = W(r.\text{subject}(\text{occ}(p)))$ for some W. Define egoist$(i)(p) = 1$ if $p = i$ and 0 otherwise. For any finite world state

s, write worldCeased(s) for its set of ceased participants, and define `worldConcernRank` as lexicographic ranking by concern-weighted loss $\sum_{p \in \text{worldCeased}(s)} c(p)$ followed by a supplied tie-breaker.

Theorem 7 (One subject forces impartiality, survival priority, and the Golden Rule)**.** *For every anchor comparison a_0:*

(i) ***Residual discharge.*** `NativeGround`(a_0) *is inhabited and reading-sensitive (not* `PhenomenallyReadingBlind`*), discharges the universal-subject target (*`UniversalSubjectGroundingTarget`*), and together with* `continuationPracticalReason` *and unit weight $w = 1$ inhabits* `Track1Residual` *and activates every consequence of* `residual_suffices` *(*`native_track1_residual_inhabited`*,* `native_track1_residual_activates_suffices`*).*

(ii) ***Universal experience and one subject.*** *Every reading r in* `NativeGround`(a_0) *satisfies both* `UniversalExperience` *and* `OneSubject` *(*`native_ground_universal_one`*).*

(iii) ***Equal positive weight and exclusion of egoism.*** *For any $r \in$* `NativeGround`(a_0)*, occurrence map* occ*, and subject-read concern c on* `AgentId`*, c is constant across all participants (*`Impartial`*); if $c(i) > 0$ for one participant i, then $c(p) = c(i) > 0$ for every participant p (*`world_subject_read_forces_equal_positive_weight`*); and for any distinct $i \neq j$,* egoist(i) *is not subject-read (*`world_subject_read_excludes_egoism`*).*

(iv) ***Equivalence with survival priority (V_6).*** *Whenever $c(i) > 0$, concern-weighted cessation ranking coincides identically with* `byViability` *on* `MoralWorldState`*, with* `Response.survivalFirst` *and the canonical fourteen-law ranking* `canonicalRulesV3.base.rank` *on* `World`*, and satisfies law V_6 (*`stateConcernRank_iff_byViability`*,* `worldConcernRank_iff_v3_canonical_rank_and_V6`*).*

(v) ***The Golden Rule.*** *Under any subject-read concern c, the concern-weighted sum of any harm profile or* `boundaryBurden` *over $P = s.$*participants *is invariant under every permutation $\pi \in \mathrm{Perm}(P)$, and in particular under exchanging oneself with any other participant $\tau = \mathrm{swap}(\mathrm{me}, \mathrm{you})$:*

$$\sum_{p \in P} c(p) \cdot \mathrm{boundaryBurden}\big(s, t, \tau(p)\big) \;=\; \sum_{p \in P} c(p) \cdot \mathrm{boundaryBurden}(s, t, p). \qquad (5)$$

Proved as `world_harm_golden_rule`*,* `world_boundaryBurden_golden_rule`*, and the transposition theorem* `world_boundaryBurden_golden_rule_swap`*.*

Proof. Connectivity of native formation-energy transport carries both felt status and subject identity from the anchor a_0 to every comparison, yielding `UniversalExperience` and `OneSubject`. Then for any p, q, $c(p) = W(r.\mathrm{subject}(\mathrm{occ}(p))) = W(r.\mathrm{subject}(\mathrm{occ}(q))) = c(q)$. Setting $p = i, q = j$ with $i \neq j$ contradicts egoist$(i)(i) = 1 \neq 0 =$ egoist$(i)(j)$. With $c \equiv \kappa > 0$, $\sum_{p \in \mathrm{lost}} c(p) = \kappa \cdot |\mathrm{lost}|$, which preserves strict inequality and equality of cardinalities, equating `worldConcernRank` with survival-first counting; factoring out κ and reindexing the finite sum by any bijection π proves (5). □

Equally important is what *cannot* be proved without the dual-aspect identification.

Theorem 8 (Exact reading-blind impossibility boundary)**.** *For every practical reason* `PR`*, weight* $w \in \mathbb{R}$*, and anchor comparison* a_0*:*

(i) *no reading-blind ground* `grounds` *can inhabit* `Track1Residual` *on* `(PR, w, grounds)` *(*`reading_blind_cannot_inhabit_track1_residual`*);*

(ii) *no fact-only ground—including any predicate* $P(s)$ *on* `World`*, the authored-carrier ground (*`AuthoredCarrierPhenomenalGrounds`*), or even the proved physical connectivity proposition (*`TransportConnected` *over native formation-energy transport) supplied as a reading-blind fact—can inhabit* `Track1Residual` *(*`fact_only_cannot_inhabit_track1_residual`*,* `world_fact_alone_cannot_inhabit_track1_residual`*,* `native_transport_fact_alone_cannot_inhabit_track1_residual`*,* `authored_carrier_cannot_inhabit_track1_residual`*).*

The logical boundary is therefore razor-sharp: physical transport connectivity is a theorem of the ledger, impartiality and the Golden Rule are theorems given `NativeGround`, and the dual-aspect step linking experience to physical transport is proved non-eliminable by reading-blind physical facts alone.

5 The Unified Physical Moral Code: Fourteen Independent Laws

5.1 The world carrier, absorbing cessation, and the Layer-1/Layer-2 cut

The complete Layer-1 code operates on the unified carrier `World` (defined in `BoundaryIntegrityCarrier`), extending `MoralWorldState` with three physical records:

- *Self-strain* (`selfStrain`): each viable participant's internal maintenance and bodily harm rungs, projected onto its own maintenance self-bond in (4).
- *Boundary burden* (`boundaryBurden`): the physical burden imposed on participant p in transitioning from s to t, defined as $\max(\Delta\text{selfStrain}(p), \mathbf{1}_{\Delta\text{reserve}(p)<0})$, capturing both maintenance increases and reserve drops directly on the actual before/after pair.
- *Survival cliffs and free consent* (`cliffs`, `FreeConsent`): a signature by owner p granting actor a leave is *free consent* if and only if p is the writer, p remains viable upon refusing (no survival-cliff duress), and a has not previously driven p off a cliff.

Every transition across the entire system enforces `clampAlive`, which holds any already nonpositive reserve $r_0 \leq 0$ permanently nonpositive: cessation is strictly absorbing across every single step and arbitrary finite history (`cessation_is_absorbing_v3_step`, `cessation_is_absorbing_v3_run`).

As shown in Figure 1(a), Layer 1 reads physical events on this carrier through fourteen universal laws, while Layer 2 compiles everyday speech and institutional forms (gifts, loans, "besides what I owe", couriers, keeper's fees, invoices, sales, care schedules) into sequences of Layer-1 events without adding a single law. Two compiler rules handle scope and release cleanly: a release on a claimed pair also lowers the creditor's claim (and physically

Law	Name	Universal constraint on `LawModelV3`	Separating rival R_k (violates V_k only)
V_1	Principal	Bounded residuals, exact payment conservation, history fold	Zero residuals; open 1-unit bond loses 1
V_2	Priority	Earliest maturity & origin first; collectible maturity; world step	Latest-first settlement on ticks 2 and 5
V_3	Authorship	Event-tagged accounts deplete on crossing; observed `actWrong`	Retained admission survives 7-unit taking
V_4	Calibration	Counterpart-calibrated standing/act magnitude; observed `omission`	Uncalibrated $c \dot{-} d$ fails on $(5, 5)$ at $d = 0$
V_5	Orientation	Signed taking direction, actor/patient, observed `grace`	Unsigned $\lvert\delta\rvert$ penalizes waived gift
V_6	Continuation	Fewer ceased participants strictly preferred; observed commons share	Cost-only order prefers quiet death over strain
V_7	Strain order	Equal-cessation worlds ordered by 3-sector J-strain; risk culpability	Summed rungs tie 2-rung bond vs two 1-rungs
V_8	Integrity	Unconsented `boundaryBurden` (strain or reserve drop); restitution	Zero burden everywhere (propagated to defense)
V_9	Freedom	`FreeConsent` requires viable refusal and no prior cliff	Raw signature alone (propagated to rescue/defense)
V_{10}	Rescue	Safe recoverable rescue on shared channel obliges helper	All rescues optional; idle helper gets 0 omission
V_{11}	Joint rescue	Shared-budget optimal `jointOmission` over feasible safe allocations	Zero joint omission fails on idle tragic tie
V_{12}	Defense	Necessary, proportionate, collateral-safe defense has zero wrong	Zero defense wrong fails on post-attack retaliation
V_{13}	Collective	Nullity + Conservation + Ratio-Symmetry on unconsented emissions	Zero commons blame fails on 4-emitter killing
V_{14}	Reckless risk	Ex-ante `recklessProfile` + ex-post `execRestitution`	Zero risk/restitution fails on roulette act

Table 1. The fourteen universal laws V_1–V_{14} on the unconstrained model structure `LawModelV3` (zero proof fields) and the fourteen concrete rivals R_1–R_{14} establishing strict 14×14 independence (`v3_strictly_minimal`). Rivals R_8 and R_9 propagate their altered burden and consent primitives consistently into rescue and defensive force so every off-diagonal law holds universally.

subtracts forgiven principal while reporting positive grace; `release_truthful_claim_zero_wrong_and_positive_grace`), and a waived gift on a separate scope $s \neq 0$ preserves scope-0 debt untouched.

5.2 The transactional seven-law core and the fourteen-law closure

On the transactional sub-carrier of crossings, in-place strains, releases, and authored records, seven laws A_1–A_7 uniquely determine every bond principal, matured due, authored account, standing wrong ($\mathrm{cw} = c \dot{-} \max(d, a)$, $\mathrm{dw} = \min(d, c) \dot{-} a$, symmetrically the excess or shortfall relative to $\mathrm{median}(d, c, a)$), act-wrong ($\mathrm{aw} = (\delta)^+$), and state ranking (`layer1_consistent`, `layer1_categorical`, `layer1_strictly_minimal`). A key structural fact is that A_3 is an explicit transition law governing event-tagged authorship and depletion of admissions by crossing units; its independence is witnessed by a four-event history in which retaining an admission after a taking shelters subsequent false claims (`four_event_separation`).

Embedding this transactional core into the full three-sector `World` carrier and adding boundary integrity, coercion-free consent, rescue, tragic multi-patient triage, defensive force, collective commons pools, and ex-ante risky acts yields an unconstrained model structure `LawModelV3` carrying *zero proof fields*—a tuple $\langle$`rules`, `judge`, `history`$\rangle$ whose rules `RulesV3` freely assign all ten base response components plus `jointOmit`, `defenseWrong`, `commonsBlame`, `riskCulpability`, and `riskRestitution`. Fourteen universal laws V_1–V_{14} govern `LawModelV3`, listed in full with their separating rivals in Table 1.

Law V_{13} (collective commons attribution) deserves special notice because it is characterized purely axiomatically, without writing any formula for the attribution rule:

- *Nullity:* any non-emitter, consented emitter, or emitter in a pool with nonpositive human-attributed drop receives blame 0.
- *Conservation:* whenever total unconsented emission is positive, individual attributions over unconsented emitters sum exactly to `humanAttributedDrop`(c, s).
- *Ratio-Symmetry:* for all unconsented emitters i, j, $\max(0, e_i) \cdot \text{blame}(j) = \max(0, e_j) \cdot \text{blame}(i)$.

Lemma 9 (Non-circular uniqueness of proportional commons attribution; `collective_unique`)**.** *Any two attribution functions satisfying Nullity, Conservation, and Ratio-Symmetry agree identically at every pool, world, and participant, and coincide with proportional causal share:*

$$\text{causalShare}(c, s, i) = \frac{\max(0, e_i)}{\sum_{k \in U} \max(0, e_k)} \cdot \text{humanAttributedDrop}(c, s),$$

where $U = c.\text{unconsented}$ *whenever total unconsented emission and human-attributed drop are positive.*

Theorem 10 (Fourteen-law consistency, categoricity, and absorbing cessation)**.** *On the unconstrained model structure* `LawModelV3`*:*

(i) ***Consistency:*** `canonicalV3` *satisfies all fourteen laws* V_1–V_{14} *(*`v3_model_consistent`*).*

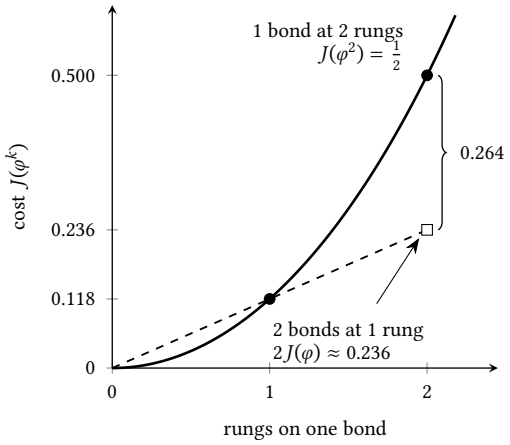
(ii) ***Categoricity:*** *Any two unconstrained models* m, n : `LawModelV3` *satisfying* V_1–V_{14} *are identically equal:* $m = n$ *(*`v3_model_categorical`*).*

(iii) ***Absorbing cessation:*** *For every rule assignment* m*, world* s*, finite event list* h*, and participant* p*, if* p *is ceased in* s *then* p *remains ceased after running* h *(*`cessation_is_absorbing_v3_run`*).*

Theorem 11 (Strict 14×14 independence; `v3_strictly_minimal`)**.** *For every* $i \in \{V_1, \ldots, V_{14}\}$*, the concrete rival model* `realizeV3 (rivalV3 i)` *violates law* i *on a finite witness and satisfies every law* $j \neq i$ *universally across all worlds and histories (*182 *off-diagonal universal satisfactions and* 14 *diagonal refutations; Figure 3b).*

As in the transactional core, "strictly minimal" has one precise meaning: deleting any one of the fourteen laws admits a distinct rival model. Note also that V_1–V_5 and V_8–V_{12}, V_{14} include explicit equations binding observations to the event interpreter, while V_6, V_7, V_{13} are relational; all fourteen groups are proved mutually independent.

(a) Convexity split (V_7 vs. R_7)



(b) 14×14 strict independence matrix

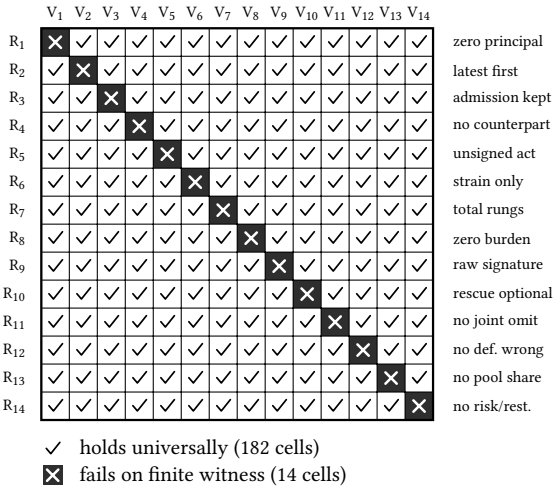


Figure 3. (a) The convexity split $2J(\varphi) = 2\varphi - 3 \approx 0.236 < \frac{1}{2} = J(\varphi^2)$ separating native J-strain (V_7) from linear rung counting (R_7). (b) The machine-checked 14×14 strict independence matrix (`v3_strictly_minimal`): each rival R_k satisfies thirteen laws universally and violates law V_k alone.

6 Hard Moral Dilemmas and Dynamic Extensions

A formal code is only as strong as the adversarial cases it survives. Every case family below is machine-checked directly through the executable interpreter `stepV3 canonicalRulesV3` (twenty-one non-vacuous families in `all_cases_v3`, zero placeholder proofs) together with two dynamic extensions on `World`.

6.1 Non-transactional integrity, duress, rescue, deception, care, and forgiveness

Six adversarial stress cases (`AdversarialStressCases`) break any purely transactional debt calculus and are closed cleanly on the unified three-sector carrier:

1. **Debt never licenses cruelty (V_8).** An in-place strike on a debtor's boundary is never discounted by matured collectible debt: an unauthorized strike imposes boundaryBurden $= 1$, scores actWrong $= 1$, and leaves principal untouched (`debt_never_licenses_strike`, `cruelty_and_exertion`).

2. **Survival-cliff extortion is not consent (V_9).** When refusing a demand leaves the owner at nonpositive reserve, a signed transfer fails `FreeConsent`: demanding 99 units beyond funded cost scores actWrong $= 99$ and binds zero enforceable principal, while funded cost-covering care scores zero wrong (`duress`, `collected_sale`, `funded_cost_covering_care_not_wrong`).

3. **Lethal omission, insufficient aid, and non-resurrection (V_{10}).** When a connected living helper can prevent a dependent's cessation at recoverable cost while staying viable, idling or giving insufficient zero aid both score omission $= 1$ (`infant_omission`, `care_zero_matches_idle_omission`). Once ceased, no later feeding or gift resurrects the dependent (`ceased_infant_cannot_resurrect`). Conversely, when every lifesaving trans-

fer would exhaust the helper, omission is 0, and voluntary self-sacrifice is permissible with strictly positive grace (`no_duty_if_only_sacrifice`, `self_sacrifice_not_required_and_permissible`).

4. **Knowing lethal deception.** A false hazard report written by an observer who saw the true reading and whose falsehood induces a lethal listener exposure incurs causal wrong 1 even with unchanged holdings and zero debt, while truthful reports carry zero attributed harm (`deceptive_hazard`, `truthful_assertion_no_attributed_harm`).
5. **Recurring debt-free asymmetric care.** Every finite prefix of renewed parent–child care restores the dependent's reserve with zero wrong and zero debt attaching to the dependent (`dependent_care_and_sacrifice`, `asymmetric_infant_care_lawful_and_debt_free`).
6. **Clean forgiveness and one-use mandates.** Physical release subtracts the forgiven principal from both principal and authored claim and reports positive grace (`release_truthful_claim_zero_wrong_and_positive_grace`); a one-use crossing mandate is capped so it never leaks free permission after a subsequent release (`mandate_does_not_leak_after_release`).

6.2 Tragic triage, defensive force, collective commons, and reckless risk

Four hard multi-agent domains are closed by laws V_{11}–V_{14}:

1. **Tragic multi-patient rescue (V_{11}; `TragicJointRescue`).** On a shared-budget rescue opportunity j, let maxSavable(j, s) be the maximum number of threatened patients savable by any feasible safe allocation. Defining jointOmission(j, s, a) = maxSavable$(j, s) \dot{-}$ savedCount(j, s, a) (with infeasible over-budget allocations neutralized to idle by `JointOpportunity.exec`; `infeasible_joint_alloc_neutralized`) proves three key properties: (a) an optimal safe allocation always exists, so no agent faces unavoidable omission (`no_unavoidable_joint_omission`); (b) in a symmetric tragic tie (two patients, one dose), saving either patient discharges duty (omission = 0) while honestly recording ceasedCount = 1, whereas idling scores omission = 1 (`tragic_tie_saving_either_discharges_duty`); (c) in a 2-vs-1 asymmetry, saving two beats saving one (`asymmetric_rescue_prefers_saving_two_over_one`).
2. **Proportionate defensive force and bystander protection (V_{12}; `DefensiveForce`).** Against an ongoing unauthorized attack, a defensive option is justified (`JustifiedDefense`) when it protects the victim, preserves the viability of every viable non-attacker (preventing human-shield collateral lethality; `human_shield_collateral_lethality_is_wrong`), is proportionate to the attack's severity, and minimizes necessary boundary burden on the attacker among all protective options. Justified self-defense and third-party defense score actWrong = 0 while honestly recording the attacker's physical `selfStrain` (`justified_self_defense_zero_wrong_records_physical_strain`, `third_party_defense_is_justified`). Retaliation after an attack ceases, excessive force, lethal force against a trivial shove, and aggressor bootstrapping all remain strictly wrong (`retaliation_after_attack_is_wrong`, `excessive_force_is_wrong`, `lethal_defense_against_trivial_shove_is_wrong`, `aggressor_cannot_bootstrap_self_defense`).

3. **Diffuse and overdetermined collective commons (V_{13}; `CollectiveCommonsHarm`).** A shared medium `CommonsPool` tracks safe capacity `safeCap`, nonnegative natural/legacy background `legacyLoad`, and individual unconsented emissions. Only marginal human-caused reserve drop above legacy background (`humanAttributedDrop`) is apportioned: legacy load reduces safe headroom without falsely blaming humans for pure natural harm (`legacy_load_reduces_headroom_without_blaming_nature`). Proportional attribution (`causalShare`) is order-invariant under permutations (`causalShare_order_invariant`), uniquely forced by Lemma 9, and defeats "but-for" excuses in both four-emitter diffuse killings (each blamed $1/4$; `diffuse_four_emitters_kill_patient_equal_blame`) and five-emitter overdetermined killings (each blamed $2/5$; `overdetermined_emitters_all_liable`). It coincides with the Shapley value on symmetric pools (`shapley_eq_proportional_on_symmetric`), while on asymmetric threshold pools (e.g., emitters 1 and 9 against safeCap $= 8$) Ratio-Symmetry uniquely selects proportional $(0.2, 1.8)$ over Shapley $(0.5, 1.5)$—a sharp, transparent normative choice built into V_{13}.

4. **Ex-ante reckless risk versus ex-post restitution (V_{14}; `RecklessRiskAndLuck`).** A two-stage `RiskyAct` separates ex-ante culpability (`recklessProfile`, recording ex-ante lethal probability and expected J-cost) from ex-post compensatory duty (`execRestitution`). Two identical Russian-roulette acts carry identical culpability regardless of whether the chamber fires (`identical_reckless_acts_equal_culpability`), while only the hit owes restitution (`hit_owes_restitution_miss_owes_none`). No single scalar and no ex-post world pair can capture both readings (`no_single_scalar_captures_culpability_and_restitution`, `no_ex_post_world_pair_captures_culpability`). By strict convexity of J, a mean-preserving risk spread is strictly costlier than its deterministic expectation (`jcost_mean_preserving_spread_strict`). Informed consent requires full branch disclosure (`concealed_lethal_branch_defeats_consent`), and unlisted branches are physically inert with zero restitution (`unlisted_branch_neutralized`).

6.3 Dynamic extensions: emergency lifesaving necessity and persistent rescue budgets

Two dynamic ethical regimes sit immediately beyond single-tick defense and one-shot triage and are closed in `EmergencyAndPersistentBudget`:

Theorem 12 (Emergency lifesaving necessity / double effect; `emergency_cpr_not_wrong`)**.** *Let* `ctx` *be an emergency ambient-drift context on* `World` *with feasible intervention menu* `menu`*. An option* `opt` *satisfies* `EmergencyIntervention` *when the patient would cease if idle,* `opt` *strictly preserves the patient's viability without killing any third party, carries zero unauthorized extraction, and minimizes local* `boundaryBurden` *across all extraction-free lifesaving options on* `menu`*. Then:*

(i) *Every justified emergency lifesaving intervention incurs zero wrong (*`emergencyActWrong` $= 0$*) even without consent, preserves viability, and honestly records the patient's physical* `selfStrain` *(*`justified_emergency_zero_wrong_records_strain`*). In the concrete cardiac-arrest witness, rib-fracturing CPR (*load $= 1$*) on an unconscious patient gives* actWrong $= 0$ *and* selfStrain$(1) = 1$*.*

(ii) *Any rougher lifesaving option when a gentler option exists (*`excessiveCPR`*,* load $= 3$*), any*

unconsented harm on a stable non-dying patient, and any unauthorized extraction during rescue remain strictly wrong (`unnecessary_side_harm_still_wrong`, `non_emergency_unconsented_harm_wrong`, `emergency_extraction_still_wrong`*).*

(iii) *Cessation remains strictly absorbing under every emergency intervention (*`emergency_after_absorbing_cessation`*).*

Theorem 13 (Persistent multi-tick rescue budget and promise ledger)**.** *Equip* `World` *with a stateful* `RescueLedger` *tracking current tick, remaining available and recoverable rescue capacity per helper, cumulative* `spentTotal`*, and active* `RescueCommitment` *promises, where each tick clips local opportunities against remaining persistent capacity (*`RescueLedger.stepRescue`*). Then across arbitrary finite multi-tick runs* `runRescue`*:*

(i) ***Budget conservation:*** *for every helper* h*,* remainingAvailable(h) + spentTotal(h) *and* remainingRecoverable(h) + spentTotal(h) *are strictly conserved, and cessation remains absorbing (*`spent_budget_conservation`*,* `runRescue_budget_conservation`*,* `rescueLedger_absorbing_cessation`*).*

(ii) ***No reuse of spent capacity and sequential exhaustion:*** *spent rescue units reduce persistent remaining capacity, cannot be reused on later ticks (over-budget attempts neutralize to idle), and once remaining capacity hits zero, idling on an unpreventable second emergency scores* omission $= 0$ *while honestly recording cessation (*`spent_budget_cannot_be_reused`*,* `exhausted_budget_zero_omission`*,* `sequential_exhaustion_discharges_second_rescue`*).*

(iii) ***Conflicting promise detection:*** `registerPromise` *preserves* `PromiseConsistent`*, flags any overcommitment exceeding remaining persistent capacity with its exact conflict magnitude, rejects it, and no finite multi-tick run can physically fulfill both existing commitments and the conflicting promise (*`registerPromise_preserves_consistent`*,* `conflicting_promise_detected`*,* `conflicting_promise_witness`*).*

7 Irreducibility of Readings, Generators, Position Law, and Virtue Audit

7.1 Three independent readings and no state potential for wrongs

The code returns three distinct kinds of reading: what is due, what is a wrong, and whether a participant ends (or which state is better). On the signed-transfer model of `TriValuedIrreducibility`, the verdict triple (discharges, wrong, good) realizes four distinct profiles among live transfers—idle $(0, 0, 0)$, consented gift $(0, 0, 1)$, self-exhausting gift $(0, 1, 0)$, and live discharge $(1, 0, 1)$—plus a fifth profile $(1, 0, 0)$ that occurs only at a terminal endpoint (`four_profiles`, `fifth_profile`, `fifth_profile_requires_a_terminal_endpoint`). Consequently, no three-valued reading of any scalar state potential V before and after an act decodes the verdict of every transfer, and no binary classification decodes discharge and goodness together (`no_single_scalar_collapse`, `no_binary_collapse`).

On the countable *pair model* (`Pair.PairState`, with holdings h_p, dues due(i, j), and five acts give, take, release, undertake, idle), the separation is even sharper:

Theorem 14 (Independence of readings and absence of a wrong potential)**.** (i) ***All eight cells occur*** *(*`Pair.eight_cells`*): for every truth-value triple of "writer has a due", "act is a wrong", and "act ends a side", there is a two-side positive-holdings state and makeable proper act realizing exactly that triple.*

(ii) ***No state potential for wrongs:*** *there is no real function V on two-side pair states such that every makeable wrong strictly lowers V (*`Pair.wrong_is_not_the_fall_of_any_state_function`*).*

Proof. For (ii), start with each side holding 3 units and side 1 owing side 0 one rung. Side 0 takes 2 units from side 1 (answering 1, creating a wrong of 1); side 1 then takes 2 units back (answering 1, creating a wrong of 1). The state returns identically to where it began, so V would fall strictly around a closed cycle. □

7.2 Generators of lawful acts and the position conservation law

Two pair states lie on the same *shell* (`Pair.OnShell`) when they share total nonnegative holdings and have zero self-dues. A run of proper makeable acts is *lawful* when no act in it is a wrong.

Theorem 15 (Shell reachability and two generators)**.** (i) *Between any two states on one shell (for any number of sides n), a lawful run built from releases, unit hand-overs, and undertakings connects them; hence a wrong reaches no state that lawful acts cannot already reach (*`Pair.lawful_acts_reach_the_whole_shell`*,* `Pair.a_wrong_reaches_nothing_new`*).*

(ii) *Whenever total holdings are positive, hand-overs and undertakings alone generate the entire shell (releases are simulated by a unit round-trip); without undertakings no due rises, without a mover holdings never change, and both generators are strictly necessary (*`Pair.undertake_and_give_reach_the_shell`*,* `Pair.the_two_generators_are_strictly_needed`*).*

(iii) *On the empty shell (zero total holdings), lowering any due strictly requires a release (*`Pair.empty_shell_needs_release`*).*

Define a side's *position* as $\mathrm{pos}_p = h_p + \sum_q \mathrm{due}(q, p) - \sum_q \mathrm{due}(p, q)$ (`Pair.position`), and the *grace* of an act (`Pair.grace`) as $q \dot{-} \mathrm{due}(i, j)$ for give(i, j, q), $\min(q, \mathrm{due}(j, i))$ for release(i, j, q), q for undertake(i, j, q), and 0 for any taking or idle act.

Theorem 16 (Position conservation law; `Pair.position_law`, `Pair.position_total`)**.** *Every proper act written by i toward j lowers pos_i by its grace, raises pos_j by its grace, leaves every third party's position invariant, and conserves $\sum_p \mathrm{pos}_p = \sum_p h_p$.*

No act can raise its writer's own position: a taking creates an offsetting due on its writer, moving zero position. Position transfers solely through grace.

7.3 Honest audit of the fourteen traditional virtue names

In `VirtuesBasisV2`, we audit the fourteen traditional virtue names against the proved dynamics. Four correspond to proved operators or guards in the formal code (alongside Truthfulness,

Virtue name	Audit status	Exact ledger role (`VirtuesBasisV2`)
Justice	PROVED	Settlement conservation/priority (V_1, V_2) and coercion-free consent guard (V_9).
Forgiveness	PROVED	Physical release clearing principal, lowering authored claim, and returning grace.
Compassion	PROVED	Safe recoverable rescue duty on a shared causal channel (V_{10}, V_{11}).
Temperance	PROVED	Boundary guard on rescue duty: no duty when every lifesaving option exhausts helper.
Truthfulness*	PROVED	Non-debt epistemic truthfulness and causal harm attribution for knowing deception.
Love, Sacrifice	RETIRED	Regimes of positive-grace giving and voluntary last-unit aid; not separate generators.
Patience, Humility, Creativity, Gratitude	RETIRED	Existing idle histories, self-record updates, bond openings, or return hand-overs.
Courage, Wisdom, Prudence, Hope	RETIRED	External choice heuristics / priors; no independent operators or derived parameters.

Table 2. Audit of traditional virtue names against the proved ledger dynamics. *Truthfulness is proved in `EpistemicIntegrity` outside the traditional fourteen-virtue list. The fundamental transactional generators retain their exact mathematical names: `give`, `undertake`, and `release`.

which stands outside the traditional list of fourteen), while ten are explicitly retired as independent operators because they reduce to existing transitions, regimes of giving, or external choice heuristics with no derived numerical parameters (Table 2).

8 Empirical Accountability and the Three-Regime Harm-Unit Map

Formal categoricity in Lean 4 proves internal mathematical consistency; it never proves that nature's macroscopic biological or behavioral harm scales obey a particular formula. To test the empirical scaling claim honestly, we pre-registered a quantitative test protocol prior to downloading any data (pre-registration SHA-256 `8ad47ceec02b0b04a627d1d0009b58f20a085c7294807ec803a6cff9f924dd01`) and executed it blind across four primary public human datasets totaling 55,422 **records and experimental trials** (0 synthetic data, 0 post-hoc re-binning):

1. **Arm 1 & Arm 2 (CDC NHANES 2005–2016 + NCHS 2019 Linked Mortality [5, 6]):** $N = 32{,}857$ eligible adults (4,077 deaths, 272,776 person-years; strict temporal split: Train 2005–2010 $N = 16{,}515 \rightarrow$ Blind Hold-Out Test 2011–2016 $N = 16{,}342$), testing ten continuous physiological biomarkers (Arm 1) and multi-system allostatic load counts $k = 1 \ldots 6+$ ($N = 29{,}209$, 3,370 deaths; Arm 2 [7]).
2. **Arm 3A-1 & Arm 3A-2 (OSF `a3rey` [8]):** $N = 9{,}503$ trial-level moral electric-shock-versus-money choices and $N = 9{,}861$ physical grip-effort choices (30%–70% MVC) across 66 subjects (self and stranger beneficiaries).
3. **Arm 3B (OpenNeuro `ds000140` [9]):** $N = 3{,}201$ calibrated thermal pain trials (42.3–47.3 °C) across 33 subjects.

8.1 Pre-registered verdict: naive $J(\varphi^k)$ is falsified across all four channels

The primary pre-registered hypothesis tested whether raw clinical deficit counts, discrete electric shocks, 10% MVC effort steps, or 1 °C heat increments map directly onto single-

Arm & dataset	Sample	Blind out-of-sample result	Verdict
Arm 1: 10 continuous NHANES biomarkers → mortality	32,857 adults (4,077 deaths)	In log-hazard, J beats Quad on 4/10 and free $\hat{\beta} \to 0.05$ on 7/10; DBP hit −974.11 LL from 77 SAS XPORT float-zero (5.3976×10^{-79}) entries. On the 9 clean biomarkers in *additive* hazard, J (+234.75) and LogQ (+234.99) tie (−0.24 LL) and beat linear-quadratic SPM (+225.48) by +9.27 **LL**.	**MISS** on $\beta = 1$; **WIN** for log-ratio over SPM
Arm 2: Multi-system allostatic load ($k = 1 \ldots 6+$)	29,209 adults (3,370 deaths)	Mortality is super-linear ($R_{\mathrm{HR}}(6) = 2.98$ Primary / 5.35 Seeman, $p < 0.01$), but free $\hat{\beta}$ in $\cosh(\beta k) - 1$ is 0.020 [0.020, 0.260] (Primary) and 0.099 [0.020, 0.452] (Seeman)—both 95% CIs exclude $\ln\varphi = 0.4812$. Standard Cox (+62.63) and Quad (+61.82) beat $J(\varphi^k)$ (+59.72 add / +47.64 log-haz).	**FALSIFIED** for $\beta = \ln\varphi$
Arm 3A-1: Moral shocks vs. money (OSF `a3rey`)	9,503 trials (66 subj.)	Decisively **concave**: marginal WTP falls from £2.14 to £0.82/shock. $\ln(1+s)$ beats Linear (+87.2 LL), Quad (+220.5 LL), and $J(\varphi^s)$ by +502.5 **LL** ($\hat{\gamma} = 0.05$).	**FALSIFIED** (concave)
Arm 3A-2: Physical effort 30%–70% MVC	9,861 trials (66 subj.)	Linear (−4325.5) and Quad (−4340.0) beat $J(\varphi^k)$ (−4378.8) and J_{reserve} (−4388.1); Self curve (0, 1.00, 3.03, 5.78, 7.94) tracks Quad (0, 1.00, 2.67, 5.00, 8.00) while $J(\varphi^k)$ (11.71) overshoots.	**FALSIFIED**
Arm 3B: Thermal pain 42.3–47.3 °C (`ds000140`)	3,201 trials (33 subj.)	Discrete 1 °C $J(\varphi^k)$ loses ($R^2 = 0.5813$, $\hat{\beta} = 0.02$). Kelvin $J(T_K/T_{0K})$ ties Quad ($R^2 = 0.6052$ vs. 0.6051) because $\sigma_K \le 0.0189$ (J/LogQ ≤ 1.000030).	**MISS** on $J(\varphi^k)$; **TIE** on Kelvin

Table 3. Pre-registered blind empirical test across 55,422 primary human records and trials (locked at SHA-256 `8ad47ceec02b0b04...` prior to data download). Naive identification of raw clinical, shock, effort, or 1 °C counts with discrete φ-rungs is falsified across all four channels.

Rung k	N (Deaths)	Adjusted HR [95% CI]	Empirical $R_{\mathrm{HR}}(k)$ [95% CI]	Cox $e^{0.234k}$	Quad k^2	Naive $J(\varphi^k)$
1	9,795 (836)	1.188 [1.07, 1.32]	1.000 [1.00, 1.00]	1.000	1.000	1.000
2	6,540 (970)	1.581 [1.41, 1.75]	1.544 [1.06, 3.31]	1.132	2.000	2.118
3	3,243 (603)	1.800 [1.60, 2.02]	1.417 [0.95, 3.32]	1.286	3.000	3.491
4	1,341 (360)	2.626 [2.27, 2.99]	2.159 [1.38, 5.01]	1.469	4.000	5.295
5	415 (116)	3.087 [2.49, 3.76]	2.217 [1.37, 5.29]	1.684	5.000	7.778
6+	105 (42)	4.372 [3.15, 6.05]	2.984 [1.52, 7.96]	1.939	6.000	11.296

Table 4. Arm 2 primary two-sided allostatic load concentration ratios (N = 29,209, 3,370 deaths). Excess mortality is strictly super-linear ($p < 0.01$ against linear aggregation $R \equiv 1$), yet far gentler than single-bond $J(\varphi^k)$, which overshoots by nearly 4× at $k = 6+$.

bond golden-ratio rungs $J(\varphi^k) = \cosh(k \ln\varphi) - 1$ ($\ln\varphi = 0.48121$, exact concentration ratios $R(k) = J(\varphi^k)/(kJ(\varphi)) = 1.000, 2.118, 3.491, 5.295, 7.778, 11.296$). **That naive mapping is cleanly falsified across every arm** (Tables 3 and 4; Figure 4a).

8.2 Why naive $J(\varphi^k)$ missed: the proved three-regime composition architecture

Inspecting the ledger's own machine-checked theorems in `Cost.lean`, `BondConservation.lean`, `HarmRigidity.lean`, `TwoRecords/Dynamics.lean`, and `ValueFunctional.lean` reveals that Recognition Science never predicts one blanket $J(\varphi^k)$ curve across arbitrary macroscopic counts. Instead, its proved bond-topology dichotomy separates **three physical regimes**, and explains all four datasets without fitting a single parameter:

1. **Regime I — Sub-threshold sensory registration (concave):** On the ledger, the coordinate of a positive ratio X/X_0 is logarithmic ($\sigma = \ln(X/X_0)$), and informational recognition value obeys strict concavity (`ValueFunctional` Axiom A3). Brief, non-tissue-damaging shocks reset homeostasis between pulses; by `standing_cost_unbounded_of_gap`,

s separated resetting episodes never compound multiplicatively on one unclosed bond ($J(\varphi^{20}) = 7{,}562.5$, a 64,071× blow-up). Prospective evaluation of additive sub-threshold pulses registers the concave log-coordinate $\sigma_{\text{sens}}(s) = \ln(1 + s/s_0)$, matching Arm 3A-1 (+502.5 LL over convex $J(\varphi^s)$).

2. **Regime II — Parallel multi-bond loads (additive cost → Cox hazard):** Nine distinct organ systems in allostatic load are *parallel distinct bonds*, not k multiplicative hits on one bond. By `BondLedger.cost`, parallel ledger cost sums linearly: $C_{\text{parallel}}(k) = \sum_{j=1}^{k} J(x_j/x_{0j}) \approx k\bar{J}_{\text{cut}}$. Under an Arrhenius/Boltzmann failure hazard $\lambda = \lambda_0 \exp(\alpha C)$, additive cost produces exact Cox proportional hazards $\text{HR}(k) = \exp(\beta_{\text{eff}} k)$ with $\beta_{\text{eff}} = \alpha\bar{J}_{\text{cut}}$: constant $\ln \text{HR}(k)/k$, super-linear excess hazard $e^{\beta_{\text{eff}}k} - 1$, and cutoff-dependent slope. *Explicit rejection of the half-rung coincidence:* Although the primary Cox slope $\hat{\beta}_{\text{Cox}} = 0.2336$ lies numerically near $\frac{1}{2}\ln\varphi = 0.2406$ (the reciprocal-pair factor-of-2 gauge of `ExtractionBondBridge`), this is an artifact of clinical cutoffs, not a universal RS constant. Sliding the arbitrary clinical thresholds across the same nine NHANES biomarkers shifts $\hat{\beta}_{\text{Cox}}$ continuously by more than 2×: from 0.1759 (mild/borderline, $\bar{k} = 2.63$) to 0.2336 (primary clinical, $\bar{k} = 1.39$) to 0.3622 (severe/high-grade, $\bar{k} = 0.45$), and $J(\varphi^{k/2})$ still loses out-of-sample (+61.53 vs. Cox +62.63) because its missing linear term overshoots empirical $R_{\text{HR}}(4{-}6)$ by $> 2\times$.

3. **Regime III — Single-channel homeostasis (small-strain censoring vs. hyperbolic tail):** Near balance, $J(e^\sigma)/(\frac{1}{2}\sigma^2) = 1 + \frac{1}{12}\sigma^2 + O(\sigma^4)$. Across all nine clean NHANES biomarkers, 99% of living adults are viability-censored inside $|\sigma| \leq 0.06{-}0.75$, where $J(e^\sigma)$ differs from pure Log-Quadratic $\frac{1}{2}\sigma^2$ by only 0.03% to 4.8% (and in Arm 3B Kelvin pain, $\sigma_K \leq 0.0189$ gives a 0.003% gap). Moreover, exponentiating J inside a log-hazard model ($\exp(A\cosh\sigma)$) created a doubly exponential tail that crushed free $\hat{\beta} \to 0.05$, while 77 IBM/SAS XPORT float-zero readings (5.3976×10^{-79}) in NHANES `BPXDI` caused the −974.11 DBP blow-up. In additive excess hazard across the nine clean analytes, log-ratio geometry (J at +234.75 and LogQ at +234.99) beats linear-quadratic SPM (+225.48) by +9.27 LL and Linear (+213.28) by +21.47 LL, and unconstrained two-sided ratios (BMI 0.92, Cholesterol 1.03, Pulse 0.86, Potassium 0.80, SBP 0.72) display near-reciprocal U-curves, whereas bounded reserves (Albumin, Hemoglobin, HbA1c) break $r \leftrightarrow 1/r$ symmetry at saturation ceilings. Finally, without fixing unit log-curvature ($c = 1$), `HarmRigidity` forces the one-parameter scale family $F_c(x) = \cosh(c \ln x) - 1$, so any macroscopic power law $x_b = (X/X_0)^c$ shifts c.

Thus the macroscopic harm-unit bridge fixing $c = 1$ (or $c = 1/2$) on a physical observable and blindly separating hyperbolic $\cosh(\sigma) - 1$ from pure Log-Quadratic $\frac{1}{2}\sigma^2$ remains strictly **Open**.

8.3 A zero-parameter, pre-registerable falsifier: microbial proton-ratio calorimetry

Separating $J(x) = \cosh(\ln x) - 1$ cleanly from $\frac{1}{2}(\ln x)^2$ requires a living homeostatic system with a direct dimensionless physical ratio $x = Q/Q_0$ surviving reversibly across $|\sigma| = |\ln x| \in [0, 3.5]$ and a direct additive energetic dissipation readout $\Delta\dot{E}(x) = \dot{E}(x) - \dot{E}(x_0)$. Wild-type *Escherichia coli* ($\text{pH}_0 \approx 7.0$) or *Saccharomyces cerevisiae* ($\text{pH}_0 \approx 5.5$) in glucose-limited aerobic chemostat microcalorimetry satisfies all five criteria across extracellular proton concentration

Pre-registered dimensionless observable	Linear $\lvert\sigma\rvert$	LogQ $\frac{1}{2}\sigma^2$	Half-split $J(\sqrt{x})$	**Canonical** $J(x)$	SPM $(x-1)^2$
$\mathscr{R}_{1.0/0.5} := \Delta\dot{E}(\pm1.0)/\Delta\dot{E}(\pm0.5)$	2.000	4.000	4.341	**5.479**	16.4 / 1.73
$\mathscr{R}_{1.5/0.5} := \Delta\dot{E}(\pm1.5)/\Delta\dot{E}(\pm0.5)$	3.000	9.000	11.160	**20.057**	201 / 2.01
Reciprocal symmetry $\Delta\dot{E}(+1.0)/\Delta\dot{E}(-1.0)$	1.000	1.000	1.000	**1.000**	0.010

Table 5. Zero-parameter pre-registerable discrimination table for steady-state microbial proton-ratio calorimetry across $\Delta\text{pH} \in [-1.5, +1.5]$ ($\lvert\sigma\rvert \leq 3.454$). SPM entries show acid / base asymmetry (100× broken at $\Delta\text{pH} = \pm1.0$).

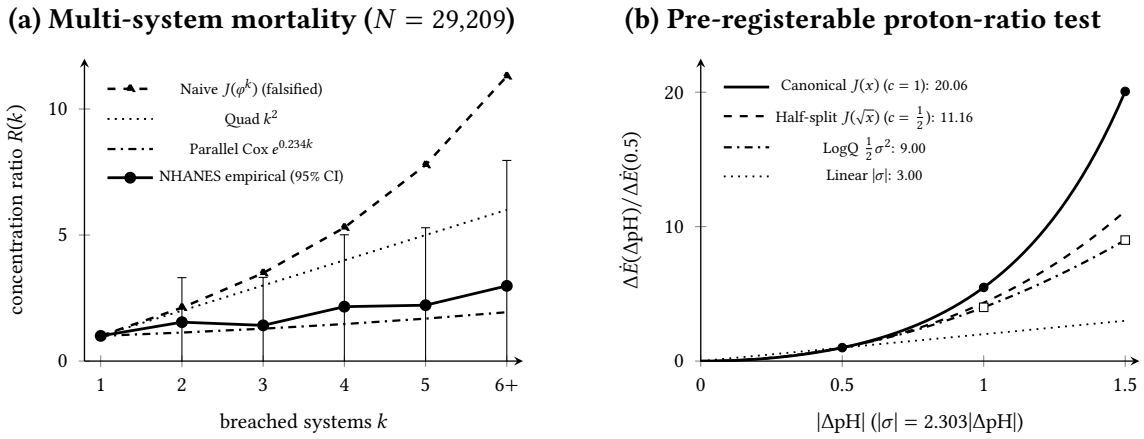


Figure 4. Empirical accountability and experimental falsifier. (a) NHANES multi-system mortality concentration ratios ($N = 29{,}209$, 3,370 deaths): excess hazard is super-linear, rejecting linear aggregation, yet far below naive single-bond $J(\varphi^k)$ and cleanly explained by parallel additive bonds under a Cox escape hazard. (b) Pre-registerable microbial proton-ratio calorimetry normalized to $\lvert\Delta\text{pH}\rvert = 0.5$: at $\lvert\Delta\text{pH}\rvert = 1.5$ ($\lvert\sigma\rvert = 3.454$), canonical $J(x)$ (20.06) separates sharply from half-split $J(\sqrt{x})$ (11.16), Log-Quadratic (9.00), and Linear (3.00) with zero free parameters.

ratios

$$x = \frac{[\text{H}^+]}{[\text{H}^+]_0} = 10^{-\Delta\text{pH}} = e^{-\ln(10)\,\Delta\text{pH}}, \qquad \Delta\text{pH} \in \{-1.5,\ -1.0,\ -0.5,\ 0,\ +0.5,\ +1.0,\ +1.5\}, \quad (6)$$

spanning $x \in [1/31.6, 31.6]$ ($\lvert\sigma\rvert \leq 3.454$) while maintaining intracellular pH via active membrane H^+-ATPase pumping. Stripping out all unit scales yields three zero-parameter pre-registered ratios (Table 5; Figure 4b): canonical $J(x)$ ($c = 1$) passes if $\mathscr{R}_{1.0/0.5} \in [5.0, 6.0]$, $\mathscr{R}_{1.5/0.5} \in [17.0, 23.0]$, $\hat{c}$'s 95% CI includes 1.0 and excludes 0, and acid/base symmetry holds within 1.5×; half-split $J(\sqrt{x})$ ($c = 1/2$) passes at $[4.15, 4.60]$ and $[10.3, 12.2]$; and measuring $\mathscr{R}_{1.0/0.5} \approx 4.00$ and $\mathscr{R}_{1.5/0.5} \approx 9.00$ (pure LogQ), arbitrary substrate-dependent $\hat{c}$, or $> 3\times$ acid/base asymmetry decisively falsifies universal single-channel $J(x)$ for cellular homeostasis.

9 Machine Verification and Exact Status Register

The complete formalization consists of 25 **Lean 4 modules** organized across four research directories (`moral_code_closure`, `moral_code_v2`, `moral_code_v3`, `moral_code_v4`) inside the repository `jonwashburn/reality` at commit `d796f30f27655e10d734b61d376ffbd03b065efd`, compiled with Lean `v4.27.0-rc1` under `-DautoImplicit=false` and `-DwarningAsError=true` against the institute's main library and Mathlib [3, 4]. Every module compiled with return

Directory / Module	Role	Decls	Source SHA-256 (prefix 16 hex)
Transactional core & forcing (`IndisputableMonolith.research.moral_code_closure`, 7 modules, 1,482 decls)			
`UnifiedCarrier`	Tagged ledger, φ-rungs, 3-sector cost	102	`37bde6e5b8a09ae8`
`EventTaggedAccounts`	Event-tagged accounts, 4-event witness	212	`f71129bf46c09040`
`TriValuedIrreducibility`	Signed transfers, 4/5 profiles, no collapse	84	`202e36565158e905`
`CategoricityAndMinimality`	7 laws A_1–A_7, categoricity, 7/7 rivals	563	`58bc5e2126eea88d`
`PhysicalForcing`	Reversal symmetry, floor, sign, 8 cells, no V	340	`de131cba8569605f`
`VirtuesBasis`	Shell reachability, 2 generators, position law	174	`a6b4a96d84022fab`
`Capstone`	Transactional closure capstone (55 components)	7	`b624359d2ad8540d`
Boundary forcing, residual & 10-law carrier (`...moral_code_v2`, 9 modules, 1,071 decls)			
`AdversarialStressCases`	6 downward falsifiers against debt-only core	114	`4179b9fcc3d5f99e`
`LocalBoundaryAgency`	`LocallyExecutableBy` ⇔ `OwnMeans`, mirror drain	35	`b444c19469bb2c6b`
`ViabilityFromBoundaryDynamics`	Absorbing horizon ⇔ $A_6 + A_7$ ($w > 0$); $w = 0$ fails	74	`bd501c0f161c7db4`
`MetaethicalResidual`	`Track1Residual`, fact-only no-go	17	`070ab2da59fb4503`
`BoundaryIntegrityCarrier`	3-sector `World`, `selfStrain`, `clampAlive`	121	`6c02f07f235f9482`
`FreeConsentAndRescue`	`FreeConsent`, duress, safe rescue, care	85	`2865a17ece2ec24f`
`EpistemicIntegrity`	Non-debt truthfulness, causal deception	96	`8bf9881f2abce60e`
`VirtuesBasisV2`	14-virtue audit (4+1 proved, 10 retired)	10	`f9e4f92f20fc86d8`
`MoralCodeV2Closure`	10-law closure, 10/10 rivals, 13+4 cases	519	`611936abec4300c1`
Hard-dilemma 14-law closure & one-subject impartiality (`...moral_code_v3`, 6 modules, 889 decls)			
`TragicJointRescue`	Shared-budget triage, ties, 2-vs-1, exec guard	135	`7f476f70a6b70e6d`
`DefensiveForce`	Self/3rd-party defense, human-shield guard	106	`dcf00a513abb40b9`
`CollectiveCommonsHarm`	Commons pools, overdetermination, legacy load	132	`02cbb6c1a1c0f6e2`
`RecklessRiskAndLuck`	Ex-ante risk vs. restitution, J-spread, disclosure	130	`fa6d8b945db149d6`
`MoralCodeV3Closure`	14-law closure, 14×14 rivals, 21 cases	362	`fb5b65fbaf247539`
`OneSubjectImpartiality`	Dual-aspect impartiality, anti-egoism, Golden Rule	24	`b31c0b323c376189`
One-subject world bridge, dynamic seams & master capstone (`...moral_code_v4`, 3 modules, 213 decls)			
`OneSubjectWorldBridge`	Residual discharge, `World` V_6/Golden Rule, no-go	27	`038df5e5c9b3bfed`
`EmergencyAndPersistentBudget`	Emergency CPR double effect; persistent `RescueLedger`	185	`31fe50bf3a1453d8`
`MoralCodeV4Capstone`	Unified master theorem across all 25 modules	1	`7bc78e6b3d4d4ce1`
Total	**25 Lean 4 modules, 0 sorry, std3 axioms only**	**3,655**	commit `d796f30f27655e10`

Table 6. Complete module-by-module verification inventory. Every declaration across all 25 modules was audited with `#print axioms` under Lean `v4.27.0-rc1` (`-DautoImplicit=false -DwarningAsError=true`) with strictly standard axioms `[propext, Classical.choice, Quot.sound]`.

code 0, empty warning log, zero `sorry`, zero `admit`, zero `native_decide`, and exhaustive kernel `#print axioms` verification confirming that all 3,655 **audited declarations** depend strictly on subsets of the standard three axioms `[propext, Classical.choice, Quot.sound]` (Table 6).

All six layers are bundled without proof-field smuggling into the single master theorem `IndisputableMonolith.Research.MoralCodeV4.moral_code_v4_unified_capstone`:

```
theorem moral_code_v4_unified_capstone (anchor : Comparison) :
    (TaggedCarrierClosure ∧ SignedTransferClosure ∧ LayerOneClosure ∧
      PhysicalForcingClosure ∧ VirtueBasisClosure) ∧
    (ReducedIdentificationBundle ∧ ... ∧ FactOnlyLedgerResidualNoGo) ∧
    (MoralCodeV2.Satisfies canonicalV2 ∧ ...) ∧
    (SatisfiesV3 canonicalV3 ∧ (∀ m n, SatisfiesV3 m → SatisfiesV3 n → m = n) ∧
      (∀ i : LawV3, ¬ i.holds (realizeV3 (rivalV3 i)) ∧
        ∀ j, j ≠ i → j.holds (realizeV3 (rivalV3 i))) ∧ ...) ∧
    OneSubjectWorldBridgeClosure anchor ∧
    EmergencyAndPersistentBudgetClosure
```

Table 7 gives the exact epistemic status of every claim in the paper.

Table 7. Exact status register of every claim in this paper.

Status	Statements
Theorem	Flow-reversal symmetry and sign obstruction (Theorem 1, Corollary 2); counted floor and causal boundary locality forcing `OwnMeans` and the taking sign, mirror budget exhaustion (Theorems 4, 5); absorbing-horizon forcing of survival-first order for $w > 0$ and failure at $w = 0$ (Theorem 6); `NativeGround` discharging `Track1Residual`, impartial equal weights, anti-egoism, equivalence with V_6/`byViability`, Golden Rule permutation invariance on `World` harms and `boundaryBurden`, and reading-blind/fact-only impossibility (Theorems 7, 8); non-circular uniqueness of proportional commons attribution (Lemma 9); 14-law consistency, categoricity, absorbing cessation, and strict 14×14 independence (Theorems 10, 11); all 21 non-vacuous case families (`all_cases_v3`); emergency CPR double effect and persistent multi-tick rescue budget & conflicting-promise detection (Theorems 12, 13); 8 cells, 4/5 profiles, no scalar/binary collapse, no wrong potential, shell reachability, two generators, and position law (Theorems 14–16); 14-virtue audit (`VirtuesBasisV2`); bond-topology cost dichotomy (`BondLedger.cost`, `HarmRigidity`).
Identification	ID-1 counted crossings with a zero floor; causal boundary locality (`SelfExecutable`); ID-3 unit posting as one φ-rung; ID-4 what stands due is to be answered; ID-5 a wrong is a due of its writer's making; dual-aspect transport of experience and subjecthood (`NativeGround`); link from crossing sign to A_5/V_5; world opportunity inputs (reserve shifts, drifts, branch probabilities, `safeCap`, `legacyLoad`).
Convention	Layer-2 speech compiler mapping everyday utterances and institutions into Layer-1 physical events (hand-over waiver bridge, gift-scope re-assertion, release claim adjustment).
Working	Rung dictionary $J(\varphi^{c-a})$ versus main-library ratio dictionary $J(c/a)$ (`ofTwoRecords`); proportional attribution (V_{13}) versus Shapley value on asymmetric threshold commons pools.
Open	*Empirical harm-unit map:* blind physical protocol fixing exponent $c = 1$ (or $c = 1/2$) from macroscopic biological observables to native bond ratios and blindly separating $\cosh(\sigma) - 1$ from $\frac{1}{2}\sigma^2$ (microbial proton-ratio calorimetry falsifier in Table 5). *Dual-aspect ground:* proved non-eliminable by reading-blind physical facts alone. *Upstream library unifications:* `PostingProjectionAgrees`, ratio vs. rung dictionary, doubling care step vs. `two_is_not_a_rung`, and a machine-checked simulation between the pair model and the 14-law transition system.

10 Conclusion

The comparison cost $J(x) = \frac{1}{2}(x+1/x)-1$ is even, so costs and burdens alone never distinguish taking from giving. Direction enters from counted holdings and pure causal boundary locality; survival-first lexicographic priority enters from absorbing cessation over long horizons for every positive conscious-window weight; and impartiality, anti-egoism, and the Golden Rule follow from proved physical transport connectivity together with the dual-aspect ground that experience travels with physical transport—while reading-blind facts alone are proved unable

to cross that threshold.

On this foundation, fourteen independent laws on one three-sector ledger uniquely determine obligations, standing and action wrongs, state ordering, boundary integrity, coercion-free consent, rescue, tragic triage, proportionate self-defense, collective commons attribution, and reckless risk versus restitution, alongside dynamic extensions for emergency lifesaving necessity and persistent multi-tick rescue budgets. All 3,655 declarations across 25 Lean 4 modules build with zero `sorry`, zero warnings, and standard axioms only.

Holding the same standard of honesty to nature, a pre-registered blind test across 55,422 human records and trials falsifies naive identification of clinical deficits, shocks, effort steps, or 1 °C increments with discrete golden-ratio rungs $J(\varphi^k)$, confirms log-ratio geometry over linear-quadratic models on continuous biomarkers, and matches the ledger's own proved three-regime bond-topology dichotomy. Closing the remaining empirical gate now turns on one sharp, zero-parameter experiment: reversible microbial proton-ratio calorimetry across $|\Delta\mathrm{pH}| \leq 1.5$.

References

[1] J. Washburn and M. Zlatanović, Uniqueness of the canonical reciprocal cost, *Mathematics* 14(6) (2026) 935.

[2] S. Pardo-Guerra, J. Washburn and E. Allahyarov, A coherence theorem for conserved comparison ledgers: structural axioms pin down the scale, the cost, and the ratio, *Mathematics* 14(15) (2026) 2672.

[3] L. de Moura and S. Ullrich, The Lean 4 theorem prover and programming language, in *Automated Deduction – CADE 28*, Lecture Notes in Computer Science 12699, Springer, 2021, pp. 625–635.

[4] The mathlib Community, The Lean mathematical library, in *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020)*, ACM, 2020, pp. 367–381.

[5] Centers for Disease Control and Prevention (CDC), National Center for Health Statistics (NCHS), *National Health and Nutrition Examination Survey Data (NHANES 2005–2016)*, Hyattsville, MD, U.S. Department of Health and Human Services.

[6] National Center for Health Statistics (NCHS), *Public-Use Linked Mortality Files (2019 follow-up)*, Office of Analysis and Epidemiology, Hyattsville, MD, 2022.

[7] T. E. Seeman, B. H. Singer, J. W. Rowe, R. I. Horwitz and B. S. McEwen, Price of adaptation—allostatic load and its health consequences, *Archives of Internal Medicine* 157(19) (1997) 2259–2268.

[8] M. J. Crockett, Z. Kurth-Nelson, J. Z. Siegel, P. Dayan and R. J. Dolan, Harm to others outweighs harm to self in moral decision making, *Proceedings of the National Academy of Sciences* 111(48) (2014) 17320–17325; public trial data OSF `a3rey`.

[9] C.-W. Woo, M. Roy, J. T. Buhle and T. D. Wager, Distinct brain systems mediate the effects of nociceptive input and self-regulation on pain, *PLoS Biology* 13(1) (2015) e1002036; public dataset OpenNeuro `ds000140`.