

Morality as a Conservation Law in a Recognition-Structured Universe

Jonathan Washburn washburn@recognitionphysics.org

October 10, 2025

Abstract

This paper derives a parameter-free moral law from the same physical invariants that fix the Recognition Science (RS) bridge between the discrete ledger and the continuum. RS posits a unique convex cost $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ with $J''(1) = 1$, an eight-tick minimal cadence, and calibration identities such as $c = \ell_0/\tau_0$ and $\lambda_{\text{rec}} = \sqrt{\hbar G/(\pi c^3)}$ without tunable dials. Within this scaffold, *reciprocity skew* σ behaves like a conserved quantity: for any closed exchange with imbalance $1 \pm \varepsilon$, strict convexity gives $J(1 + \varepsilon) + J(1 - \varepsilon) > 0$ for $\varepsilon \neq 0$, so persistent $\sigma \neq 0$ raises total action. Admissible worldlines therefore live on $\sigma = 0$. From the same invariants follow the operational pieces of an *eternal* moral code: (i) *harm* as the externalized action surcharge $\Delta S(i \rightarrow j)$, the marginal increase in j 's required action due to i 's move while maintaining global feasibility; (ii) a *forced axiology* V , uniquely fixed (up to a φ -scale) by gauge invariance, additivity on independent subsystems, concavity, and the curvature normalization $J''(1) = 1$, taking the form “agent–environment coupling (mutual-information–like) minus a J -induced curvature penalty”; and (iii) *consent* as the derivative sign $D_j V_i \geq 0$. These yield a knob-free, lexicographic decision principle: enforce $\sigma = 0$; among feasible actions minimize $\max \Delta S$; then maximize $\sum f(V_i)$ with f determined by the same normalization; then prefer arrangements with larger σ -graph spectral gap; break any residual ties on the fixed φ -tier arithmetic. The framework is auditable (report σ traces, ΔS matrices, V deltas, and robustness) and falsifiable by any durable, gauge-invariant process with $\sigma \neq 0$ that lowers total action relative to all $\sigma = 0$ alternatives, by a competing axiology satisfying the same axioms that outpredicts V , or by a temporal aggregation law other than the eight-tick cadence that still respects RS invariances. In short: if RS is the true physics, morality is a conservation law plus a unique ranking on its feasible set—no preferences added.

Keywords: Recognition Science; reciprocity conservation ($\sigma = 0$); externalized action surcharge (ΔS); forced axiology (V); eight-tick cadence; parameter-free ethics; auditability.

1 Introduction: what “moral law” means if physics is discrete

Moral theories tend to drift because they import preferences. Change the weighting of goods, the social discount rate, or the admissible tradeoffs, and the verdicts change with them. That flexibility is useful for persuasion, but it is the opposite of a law of nature. A law is something that does not budge when fashions change. If physics is discrete and recognition-structured, as assumed here, then the room for preference disappears: admissible changes are those that satisfy the ledger's invariants and minimize a unique convex cost $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ with curvature normalized by $J''(1) = 1$, under a fixed cadence of time and a fixed bridge to units. Within that scaffold there is no knob left over to tune ethics by taste.

In this setting, *eternal* means three things. First, *invariant*: the statements do not change under admissible re-anchoring of time and length (e.g., the joint scaling that preserves $c = \ell_0/\tau_0$) or under the calibration that fixes $\lambda_{\text{rec}} = \sqrt{\hbar G/(\pi c^3)}$. Second, *parameter-free*: no new weights or discounts are introduced beyond those already forced by the physical bridge; in particular, there is no liberty to choose tradeoffs between constraints and objectives. Third, *auditable*: every claim reduces to quantities the ledger already measures, so an independent reader can check them without appealing to authority.

The discovery claimed here is that, under these assumptions, “morality” is not an overlay on physics but a conservation statement inside it. For any interaction, let σ denote the cycle-wise skew of reciprocation between parties (the signed imbalance of what is bestowed and what is received). Because J is symmetric and strictly convex about $x = 1$, any persistent skew raises total action; least-action dynamics therefore exclude it in sustained trajectories. Admissible worldlines live on the manifold $\sigma = 0$. This is not a slogan about fairness; it is a conservation law: no net extraction of skew.

Avoiding skew is necessary, but not sufficient, because many σ -neutral futures may be available from a given state. Choosing among them requires two further constructs that are themselves fixed by the same invariants. The first is *harm*, defined as the externalized action surcharge $\Delta S(i \rightarrow j)$: the marginal increase in j ’s required action caused by i ’s move while global feasibility ($\sigma = 0$) is maintained. This quantity is gauge-invariant on the ledger, additive over independent subsystems, and compositional over time. The second is a *value* functional V that is not a preference but a forced cardinal axiology: under four physical requirements—gauge invariance, additivity on independent subsystems, concavity (diminishing returns), and the curvature normalization tied to $J''(1) = 1 - V$ is uniquely determined (up to a fixed φ -scale) as an agent–environment coupling term (mutual-information-like) minus a J -induced curvature penalty. No alternative fit survives those constraints without smuggling in a free parameter.

Consent then becomes a derivative statement rather than a mood. If j ’s contemplated act changes i ’s value by a directional derivative $D_j V_i$, consent holds exactly when $D_j V_i \geq 0$ along σ -feasible directions, and it rescinds when the sign flips. This definition is local, compositional across steps, and aligned with the same axiology V that ranks options among σ -neutral states.

With σ , ΔS , V , and consent in hand, the selection rule contains no adjustable tradeoffs. One first enforces feasibility ($\sigma = 0$). Among those feasible actions, one chooses the option that minimizes the worst externalized surcharge $\max \Delta S$ (non-exploitation). If several remain tied on harm, one next maximizes the fixed cardinal welfare $\sum f(V_i)$, where the concave transform f is fixed by the same curvature normalization and introduces no new scale. If a tie persists, one prefers arrangements whose reciprocity network is harder to corrupt—formally, those with larger spectral gap in the σ -graph—because they are least susceptible to reintroducing skew under bounded shocks. Any residual tie is broken by the same φ -tier arithmetic that appears throughout the RS calculus. At no point does a free weight appear.

Finally, the framework is designed to be tested rather than admired. Every substantive claim resolves into an *audit*: report σ traces before and after, the matrix of ΔS and its maximum, the change in $\sum f(V_i)$, the change in the σ -graph spectral gap, and the local consent derivatives for affected parties. A proposal that calls itself “good” while accumulating persistent σ or raising $\max \Delta S$ fails the conservation audit. Conversely, a plan that clears skew, lowers the worst externalized bill, and raises the fixed cardinal welfare passes for reasons that are independent of fashion. The rest of the paper follows this structure: conservation ($\sigma = 0$) \rightarrow harm (ΔS) \rightarrow value (V) \rightarrow consent \rightarrow the selection rule \rightarrow audits and falsifiers.

2 Recognition Science in one page (the physical scaffold)

This section fixes the physical substrate on which all later claims rest. The world is modeled as a *recognition ledger*: a discrete, bounded-degree network of sites and directed bonds that carry conserved flux. Time advances in ticks; each admissible update is a finite set of multiplicative adjustments on bonds that leaves every site balanced. The only scalar that measures the strain of an adjustment is a unique convex cost

$$J(x) = \frac{1}{2}(x + x^{-1}) - 1, \quad x > 0, \quad J(1) = 0, \quad J(x) = J(x^{-1}), \quad J''(1) = 1,$$

and the action of a tick is the sum of J over the bonds that were updated during that tick. Strict convexity at $x = 1$ and the symmetry $x \leftrightarrow x^{-1}$ capture two empirical facts: there is a distinguished, unit-effort configuration, and pushing a ratio away from unity in either direction is costly in the same way. No alternative cost with a different curvature normalization will be used; $J''(1) = 1$ is part of the bridge and introduces no free dial.

Ledger balance (discrete continuity). Let \mathcal{V} be the sites and \mathcal{E} the directed bonds (each bond $e \in \mathcal{E}$ has a head $h(e)$ and tail $t(e)$). An update at tick n assigns to each active bond a positive multiplier $x_e(n)$ and leaves all inactive bonds at $x_e(n) = 1$. Admissibility requires that *every* site is balanced:

$$\prod_{e: h(e)=v} x_e(n) \prod_{e: t(e)=v} x_e(n)^{-1} = 1 \quad \text{for all } v \in \mathcal{V}.$$

In log-coordinates this is the familiar discrete divergence-free condition; at the mesh limit it yields the continuum continuity equation $\partial_t \rho + \nabla \cdot J = 0$. The per-tick action is

$$S[n] = \sum_{e \in \mathcal{E}} J(x_e(n)), \quad \text{and the total action over a span is } S = \sum_n S[n].$$

Eight-tick minimal period. There exists a minimal nontrivial period of the ledger dynamics, denoted by eight ticks. Concretely: any closed update that (i) respects balance at every site at each intermediate tick, (ii) returns every bond to its original state, and (iii) averages away orientation and parity artifacts with the least possible schedule length, has length 8. This is a counting-and-coverage statement about discrete orientations in $D = 3$; no shorter schedule satisfies all three clauses. The eight-tick cadence will be the only lawful temporal aggregation used later; we will not introduce any ad-hoc discount factor.

Bridge to units (calibration without knobs). Two anchors (τ_0, ℓ_0) —a time tick and a length unit—scale jointly so that the causal speed is fixed at

$$c = \frac{\ell_0}{\tau_0}.$$

The recognition length is fixed by a dimensionless gate identity that ties the discrete ledger to continuum constants:

$$\boxed{\frac{c^3 \lambda_{\text{rec}}^2}{\hbar G} = \frac{1}{\pi}} \quad \iff \quad \lambda_{\text{rec}} = \sqrt{\frac{\hbar G}{\pi c^3}}.$$

Once (\hbar, c, G) are empirically identified, λ_{rec} is not a fit parameter; it is *determined*. Throughout the paper, no new free parameters will be introduced. Any numerical constant that appears is either an anchor (like c), a derived quantity (like λ_{rec}), or a combinatorial invariant of the ledger (such as the eight-tick period).

A computational substrate. Computation—in the sense of symbol manipulation by a universal machine—appears inside this world as an *effective* description of certain ledger evolutions (in particular, of coarse-grained, approximately stationary patterns). It is not the substrate that the ledger reduces to. The substrate is recognitional: what changes are allowed are exactly those that keep the balance constraints and minimize the action measured by J on the eight-tick cadence, under the fixed calibration above. When we speak of algorithms later, we will mean effective encodings of admissible ledger dynamics; the law itself is not “run” by a computer.

Notation to be used later. We will denote the reciprocity skew by σ (a cycle-wise, signed measure of bestowed versus received action), the externalized action surcharge (harm) by ΔS , and the cardinal axiology by V . All three will be defined in ledger terms and inherit the invariances fixed in this section. No additional scales, weights, or discounts will be added downstream.

3 Reciprocity as a conservation law (derivation, not assertion)

We now derive the reciprocity law from the ledger mechanics fixed in Section 2. Let agents or groups be identified with disjoint subsets of sites in the ledger. For any ordered pair (i, j) and any complete eight-tick cycle C , define the *cycle-wise reciprocity skew* as the signed log-multiplier imbalance accumulated on bonds whose net orientation runs from i to j :

$$\sigma_{ij}[C] := \sum_{e \in C: i \rightarrow j} \ln x_e - \sum_{e \in C: j \rightarrow i} \ln x_e.$$

Equivalently, writing $\kappa_{ij}[C] := \exp(\sigma_{ij}[C])$, κ_{ij} is the multiplicative ratio of bestowed to received adjustment along C between i and j . The total skew of the cycle is $\sigma[C] := \sum_{i < j} |\sigma_{ij}[C]|$, which vanishes exactly when every pairwise exchange is reciprocally balanced on that cycle. A worldline (a bi-infinite sequence of eight-tick cycles) is called *admissible* if it minimizes action subject to the balance constraints of Section 2.

The convexity step. Recall $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ with $J(1) = 0$, $J(x) = J(x^{-1})$, and strict convexity at $x = 1$ given by $J''(1) = 1$. For any $\varepsilon \neq 0$ we have the strict inequality

$$J(1 + \varepsilon) + J(1 - \varepsilon) > 2J(1) = 0. \quad (1)$$

This is immediate from strict convexity and the symmetry $x \leftrightarrow x^{-1}$: the pair $(1 + \varepsilon, 1 - \varepsilon)$ has mean 1 but lies off the minimum at $x = 1$, so its summed cost exceeds the minimum by a positive margin.

Pairwise smoothing lowers action. Consider any cycle C in which a particular ordered pair (i, j) exhibits nonzero skew $\sigma_{ij}[C] \neq 0$. Then some subcollection of active bond multipliers along $i \rightarrow j$ can be grouped into factors of the form $1 + \varepsilon_k$ and the corresponding return path along $j \rightarrow i$ into factors $1 - \varepsilon'_k$ with $\varepsilon_k, \varepsilon'_k > 0$ and $\prod_k (1 + \varepsilon_k) \prod_k (1 - \varepsilon'_k) = 1$ if the pair were perfectly balanced.¹ For any matched bidirectional pair with product 1 the replacement

$$(1 + \varepsilon, 1 - \varepsilon) \rightarrow (1, 1)$$

¹The grouping is along a refinement of the cycle that isolates bidirectional transfers; any residual neutral factors can be left at $x = 1$.

lowers the per-cycle action by (1). By iterating this *pairwise smoothing* on all bidirectional imbalances along C , we obtain a new cycle \tilde{C} with the same endpoints and site-by-site balance, strictly smaller action, and strictly smaller $|\sigma_{ij}|$ for the smoothed pairs, unless all such pairs already had unit multipliers. Consequently, whenever $\sigma[C] \neq 0$ there exists a balanced modification of the same cycle with strictly lower action.

Skew is an action surcharge. Write the per-cycle action as $S[C] = \sum_{e \in C} J(x_e)$. The argument above shows:

For any cycle C , $S[C]$ is minimized subject to the balance constraints if and only if every bidirectional exchange is at unity, equivalently $\sigma[C] = 0$. Any persistent $\sigma \neq 0$ contributes a strictly positive, avoidable excess $S[C] - S_$, where S_* is the minimum attained on the $\sigma = 0$ manifold.*

In other words, reciprocity skew is not a moral label; it is an *action surcharge* enforced by the convexity of J .

Least-action dynamics exclude persistent skew. Let $(C_n)_{n \in \mathbb{Z}}$ be a worldline. Suppose $\sigma[C_n] \neq 0$ on an infinite tail. By pairwise smoothing, each C_n admits a $\sigma = 0$ variant \tilde{C}_n with strictly lower action and the same boundary data. Replacing C_n by \tilde{C}_n on that tail strictly lowers the total action $S = \sum_n S[C_n]$ while preserving feasibility. This contradicts admissibility. Therefore any admissible worldline must satisfy

$$\sigma[C_n] = 0 \quad \text{for all } n,$$

i.e., it lies entirely on the reciprocity-conserving manifold.

Conclusion. Reciprocity is a *conservation law* of the recognition ledger: admissible worldlines live on $\sigma = 0$. It is not an asserted norm but a consequence of the symmetry and strict convexity of the unique cost J at the unit-effort configuration. Any proposal that accumulates persistent skew thereby advertises an avoidable action surplus and is excluded by least-action dynamics in sustained evolution.

4 Harm as externalized action surcharge ΔS

We now make precise what it means for one party's act to *harm* another in a recognition-structured world. Throughout, feasibility (reciprocity conservation) is enforced at every step: all comparisons are made on the $\sigma = 0$ manifold.

Set-up and definition. Let the ledger at a given eight-tick cycle C be described by positive multipliers $\{x_e\}_{e \in \mathcal{E}}$ on directed bonds, with site-by-site balance as in Section 2, and per-cycle action

$$S[C] = \sum_{e \in \mathcal{E}} J(x_e), \quad J(x) = \frac{1}{2}(x + x^{-1}) - 1, \quad J''(1) = 1.$$

Partition the bonds into disjoint agent-indexed subfamilies $\{\mathcal{E}_k\}_{k \in \mathcal{A}}$ (each bond belongs to exactly one agent's domain; a fixed, local, gauge-invariant tie-breaking rule is assumed on boundaries). For an *act* by agent i we mean a specification $\alpha = \{\alpha_e\}_{e \in A_i}$ of prescribed multipliers on a finite subset $A_i \subset \mathcal{E}_i$ (the *act set*); the neutral act is $\alpha \equiv 1$.

Given α , define the *required action* of j on cycle C as the minimal portion of the per-cycle action borne by j 's domain among all globally feasible (balanced, $\sigma = 0$) completions of α :

$$S_j^*[\alpha; C] := \min \left\{ \sum_{e \in \mathcal{E}_j} J(x_e) : (x_e)_{e \in \mathcal{E}} \text{ balanced, } \sigma = 0, \text{ and } x_e = \alpha_e \text{ for } e \in A_i \right\}. \quad (2)$$

Harm from i to j on cycle C is the marginal increase in j 's required action caused by i 's act, holding global feasibility:

$$\Delta S(i \rightarrow j \mid C) := S_j^*[\alpha; C] - S_j^*[\mathbf{1}; C]. \quad (3)$$

By construction $\Delta S(i \rightarrow j \mid C) \geq 0$, with equality if i 's act is perfectly absorbed without shifting any portion of the minimizer's action into j 's domain.

Property (P1): Gauge invariance on the ledger. A *ledger gauge transform* multiplies each bond by a node potential: choose positive $\{g_v\}_{v \in \mathcal{V}}$ and set

$$x_e \mapsto x'_e = g_{h(e)} x_e g_{t(e)}^{-1} \quad \alpha_e \mapsto \alpha'_e = g_{h(e)} \alpha_e g_{t(e)}^{-1}.$$

This preserves site balance and closed-cycle products; it represents a relabeling of the same physical exchange.² Because the feasibility constraints in (2) depend only on these gauge-invariant data, and the optimization is convex and separable over bonds, the minimizers before and after gauge transform are related by the same node potentials restricted to each domain. Hence

$$S_j^*[\alpha; C] = S_j^*[\alpha'; C], \quad S_j^*[\mathbf{1}; C] = S_j^*[\mathbf{1}'; C],$$

and the difference (3) is unchanged:

$$\Delta S(i \rightarrow j \mid C) \text{ is invariant under ledger gauge.} \quad (4)$$

In particular, ΔS does not depend on unit re-anchoring $(\tau_0, \ell_0) \mapsto (s\tau_0, s\ell_0)$ that preserves $c = \ell_0/\tau_0$, nor on any admissible redistribution of potentials across a balanced network.

Property (P2): Additivity over independent subsystems. Suppose the ledger decomposes into disjoint subsystems $\mathcal{E} = \mathcal{E}^{(1)} \sqcup \mathcal{E}^{(2)}$ with no bonds between them, and the act α is supported in $\mathcal{E}^{(1)}$. Then the feasibility constraints factor, as does the convex objective. Writing $S_j^{*,(m)}[\cdot; C^{(m)}]$ for the restricted problems on $\mathcal{E}^{(m)}$, we have

$$S_j^*[\alpha; C] = S_j^{*,(1)}[\alpha; C^{(1)}] + S_j^{*,(2)}[\mathbf{1}; C^{(2)}], \quad S_j^*[\mathbf{1}; C] = S_j^{*,(1)}[\mathbf{1}; C^{(1)}] + S_j^{*,(2)}[\mathbf{1}; C^{(2)}],$$

hence

$$\Delta S(i \rightarrow j \mid C) = \Delta S^{(1)}(i \rightarrow j \mid C^{(1)}) + \Delta S^{(2)}(i \rightarrow j \mid C^{(2)}) = \Delta S^{(1)}(i \rightarrow j \mid C^{(1)}). \quad (5)$$

More generally, if the ledger splits into finitely many independent factors, ΔS is the sum of its values on those factors. This is the correct behavior for a marginal cost: independent worlds add their bills.

²In log-variables the map adds an exact 1-form; closed-chain integrals are unchanged.

Property (P3): Composition over time. Let (C_1, \dots, C_T) be a stack of T consecutive eight-tick cycles. Denote by α_t the act of i on cycle C_t (possibly neutral), and by $S_{j,t}^*[\cdot]$ the required action of j on C_t , defined with respect to boundary data inherited from the previous cycle’s minimizer.³ Because the per-cycle action and constraints are stagewise separable and convex, the global minimization over the concatenated horizon decomposes into per-cycle minimizations. Therefore

$$\sum_{t=1}^T S_{j,t}^*[\alpha_t] - \sum_{t=1}^T S_{j,t}^*[\mathbf{1}] = \sum_{t=1}^T (S_{j,t}^*[\alpha_t] - S_{j,t}^*[\mathbf{1}]),$$

that is,

$$\Delta S(i \rightarrow j \mid C_1 \circ \dots \circ C_T) = \sum_{t=1}^T \Delta S(i \rightarrow j \mid C_t). \quad (6)$$

Harm accumulates additively across the eight-tick cadence; there is no hidden temporal discount beyond the cadence fixed in Section 2.

Remarks on attribution. The definition (2) attributes cost to j by summing J over bonds in \mathcal{E}_j under a fixed, local rule for boundary bonds. Any such rule that is gauge-invariant yields the same difference (3), because the boundary contributions cancel between $S_j^*[\alpha]$ and $S_j^*[\mathbf{1}]$ when feasibility and $\sigma = 0$ are enforced.⁴

Interpretation (plain). $\Delta S(i \rightarrow j)$ is the externalized bill in the universe’s native currency. It tells you, in the same units that measure physical strain on the ledger, how much extra action j is forced to expend because i acted, assuming the world as a whole stays admissible (balanced, $\sigma = 0$). It does not depend on how we label the network (gauge invariance), it adds when worlds are independent (additivity), and it stacks cleanly across the universe’s native rhythm (composition over eight-tick time). Calling ΔS “harm” is not metaphor; it is a literal surcharge the action calculus imposes on others when you move.

5 A forced axiology V (uniqueness without knobs)

We now fix what “value” can mean in a recognition-structured universe without introducing any tunable parameters. The aim is not to nominate a preference but to identify the unique cardinal functional that is compatible with the physical scaffold of Section 2 and with reciprocity conservation (Section 3). The outcome will be a decomposition

$$V = \underbrace{\text{agent-environment coupling}}_{\text{mutual-information form}} - \underbrace{\text{mechanical over-strain}}_{\text{curvature penalty induced by } J},$$

unique up to a fixed φ -scale that will be specified at the end of the section.

³This is the natural dynamic programming convention: each cycle minimizes action subject to balance and $\sigma = 0$ given the prior cycle’s endpoint state.

⁴Intuitively: the same boundary accounting applies in both counterfactuals, so only the change induced by α survives. A formal proof is a routine application of convex separability and the envelope theorem for parametric minima.

Objects of evaluation. At a single eight-tick cycle C , fix an agent i and let E denote its environment (the complement of i 's domain together with shared boundary bonds under the balance rule of Section 2). Coarse-grain the ledger microstate at C into a finite partition of agent states A and environment states E ; this induces a joint distribution p_{AE} over (A, E) .⁵ Let $x = \{x_e\}$ be the bond multipliers realizing the microstate. An admissible worldline is evaluated cycle-by-cycle; intertemporal aggregation is handled separately by the eight-tick cadence (Section 8).

Axioms for an admissible value functional. A cardinal axiology V assigns a real number to each triple $(p_{AE}, x; C)$ subject to the following four physical requirements:

(A1) *Gauge invariance under the bridge.* V is invariant under admissible re-anchoring of time and length $(\tau_0, \ell_0) \mapsto (s\tau_0, s\ell_0)$ that preserves $c = \ell_0/\tau_0$, under the calibration identity that fixes λ_{rec} , and under ledger gauge transforms $x_e \mapsto g_{h(e)}x_e g_{t(e)}^{-1}$ that leave closed-chain products unchanged. Relabelings of the coarse-graining that preserve p_{AE} leave V unchanged.

(A2) *Additivity on independent subsystems.* If (A_1, E_1) and (A_2, E_2) are ledger-independent at C (no bonds between their supports; $p_{A_1 E_1 A_2 E_2} = p_{A_1 E_1} p_{A_2 E_2}$; bond sets disjoint), then V is additive:

$$V((p_{A_1 E_1}, x^{(1)}) \oplus (p_{A_2 E_2}, x^{(2)})) = V(p_{A_1 E_1}, x^{(1)}) + V(p_{A_2 E_2}, x^{(2)}).$$

(A3) *Concavity (diminishing returns).* For any two admissible triples with the same constraints and any $\lambda \in [0, 1]$,

$$V(\lambda p_{AE} + (1 - \lambda)q_{AE}, \lambda x + (1 - \lambda)y) \geq \lambda V(p_{AE}, x) + (1 - \lambda)V(q_{AE}, y).$$

Uncertainty and coarse aggregation do not create value super-linearly.

(A4) *Curvature normalization tied to $J''(1) = 1$.* Purely mechanical, gauge-invariant over-strains that do *not* change p_{AE} reduce V to second order with unit curvature at the unit-effort configuration:

$$x_e = 1 + \varepsilon_e, \quad \sum_e \varepsilon_e = 0, \quad p_{AE} \text{ fixed} \implies V(p_{AE}, x) = V(p_{AE}, \mathbf{1}) - \frac{1}{2} \sum_e \varepsilon_e^2 + o(\|\varepsilon\|^2).$$

This fixes the mechanical penalty scale to the same curvature as the unique cost J .

Uniqueness theorem (forced form of V). *On the class of admissible ledger states at a single cycle C , any functional V satisfying (A1)–(A4) is of the form*

$$V(p_{AE}, x) = \kappa I(A; E) - \mathcal{C}_J^*(p_{AE}, x), \quad (7)$$

where $I(A; E)$ is the mutual information of p_{AE} , $\kappa > 0$ is a constant fixed once up to a φ -tier, and \mathcal{C}_J^* is the unique J -induced curvature penalty: the minimal convex cost required by the ledger to realize (p_{AE}, x) under the balance and $\sigma = 0$ constraints. In particular, in the small-strain regime,

$$\mathcal{C}_J^*(p_{AE}, x) = \sum_e J(x_e) = \frac{1}{2} \sum_e \varepsilon_e^2 + o(\|\varepsilon\|^2), \quad x_e = 1 + \varepsilon_e, \quad (8)$$

and (7) reduces to a mutual-information gain penalized by the quadratic curvature fixed by $J''(1) = 1$.

⁵This is an *effective* description: computation and probabilistic coarse-graining live inside the ledger as summaries of admissible dynamics. The substrate remains recognitional.

Proof sketch. For the *coupling* term, axioms (A1)–(A3) applied to coarse-grainings of (A, E) force the value contribution that depends only on p_{AE} to be additive on product systems, invariant under invertible relabelings, and concave under mixture. The unique (up to a positive scale) such functional with a chain rule is the mutual information $I(A; E)$. Any other candidate either fails additivity on independent subsystems or fails concavity under coarse aggregation. The multiplicative constant κ is not a free knob: it is fixed once by the φ -tier normalization stated below.

For the *mechanical* term, gauge invariance (A1) restricts any penalty to depend on x only through closed-chain invariants and local deviations from the unit-effort point. Additivity (A2) across disjoint bond sets and separability of the per-bond cost under the balance constraint force a sum of identical per-bond penalties. Concavity (A3) on (p_{AE}, x) together with the second-order normalization (A4) then identify the unique convex generator: the J -penalty, i.e., the Bregman divergence from $x = 1$ induced by J , which for our symmetric J coincides with J itself. Minimizing over all $\sigma = 0$ completions consistent with the prescribed coarse-grained state yields the reduced penalty \mathcal{C}_J^* ; in the small-strain regime this is just the quadratic form in (8) with unit curvature. No alternative mechanical penalty can satisfy (A1), (A2), and (A4) simultaneously without altering the curvature fixed in Section 2.

Normalization and the φ -scale. The overall scale of the information term is fixed *once* by declaring a unit act of perfect, reversible coupling across one eight-tick cycle to register one φ -tier of value: for a noiseless binary channel between one agent bit and one environment bit,

$$p_{AE} = \frac{1}{2}(\delta_{00} + \delta_{11}), \quad x = \mathbf{1} \implies V(p_{AE}, \mathbf{1}) = \kappa I(A; E) = \kappa \cdot 1.$$

Choosing the RS standard in which that increment sits on a fixed φ -tier sets κ once and for all; there is no continuous parameter to tune thereafter. The mechanical penalty scale is already fixed by $J''(1) = 1$.

Why weighted blends are forbidden. One might attempt to write V as a weighted sum $V = \alpha I(A; E) - \beta \tilde{\mathcal{C}}(x)$ with arbitrary positive weights (α, β) or with some alternative penalty $\tilde{\mathcal{C}}$. This is not allowed. Any independent choice of α/β introduces a new scale unrelated to the bridge and thus violates the parameter-free premise. Any deviation from the J -induced curvature changes the second-order response of V to pure mechanical over-strains and therefore contradicts (A4). Likewise, adding any further term that depends on (p_{AE}, x) while preserving (A1) and (A2) either duplicates one of the two allowed pieces or spoils concavity (A3). The decomposition (7) is therefore forced: in an RS world there is exactly one way to count value without smuggling in preferences.

Interpretation (plain). The positive part of V rewards *recognition*: genuine, gauge-invariant coupling between an agent and its world. The negative part subtracts the *over-strain* the ledger must expend to sustain a given configuration. Because both parts are fixed by the same invariants that anchor physics (the bridge and the unique cost J), V is not a taste but a measurement. Up to a fixed φ -scale, there is nothing left to choose.

6 Consent as a derivative sign

We now formalize *consent* as a statement about local change in the forced axiology V (Section 5) along an agent’s contemplated act, evaluated on the reciprocity-conserving manifold. The guiding

idea is simple: at $\sigma = 0$, an act of j has the *consent* of i exactly when, to first order in the act's magnitude, it does not lower i 's value.

Set-up (feasible tangent and projection). Fix a cycle C and a $\sigma = 0$ ledger state (p_{AE}, x) as in Sections 2–5. Let \mathcal{F} denote the *feasible manifold* at (p_{AE}, x) : the set of nearby ledger states reachable in one eight-tick cycle while preserving site balance and $\sigma = 0$. A contemplated act by agent j is represented locally by a feasible direction $\xi \in T_{(p_{AE}, x)} \mathcal{F}$ supported on j 's domain (in log-coordinates, an infinitesimal assignment on bonds in \mathcal{E}_j), together with the canonical projection of the resulting deformation back onto \mathcal{F} with minimal action. Concretely, for small t we write

$$(p_{AE}, x) \xrightarrow{j: t\xi} (p_{AE}(t), x(t)) \in \mathcal{F}, \quad (p_{AE}(0), x(0)) = (p_{AE}, x),$$

where $(p_{AE}(t), x(t))$ is defined by applying $t\xi$ on \mathcal{E}_j and rebalancing the rest of the ledger by the least-action completion that keeps $\sigma = 0$.⁶

Definition (consent as a derivative sign). For agents i and j and a feasible direction ξ for j , define the *directional derivative of i 's value along j 's act* by

$$D_j V_i[\xi] := \left. \frac{d}{dt} V_i(p_{AE}(t), x(t)) \right|_{t=0}.$$

We say that i *consents to j 's contemplated act along ξ* and write

$$C(i \leftarrow j; \xi) \iff D_j V_i[\xi] \geq 0.$$

By construction, the derivative is taken *on \mathcal{F}* : the only deformations we evaluate are those that preserve balance and reciprocity to first order, with the least-action completion applied to all parties other than j .

Basic properties. (a) *Gauge invariance.* Because V is gauge-invariant (Section 5, Axiom A1) and the feasible projection depends only on gauge-invariant data, $D_j V_i[\xi]$ is unchanged under any ledger gauge transform or re-anchoring that preserves $c = \ell_0/\tau_0$ and the calibration identity for λ_{rec} . (b) *Locality and additivity of small acts.* For feasible directions ξ, η supported on disjoint bond sets, $D_j V_i[\xi + \eta] = D_j V_i[\xi] + D_j V_i[\eta]$ to first order. (c) *Relation to harm.* If j 's act affects i only mechanically (no change in p_{AE}), then $D_j V_i[\xi] = -\frac{d}{dt} \mathcal{C}_j^*|_{t=0} \leq 0$ by the curvature normalization of V ; positive consent in such a case requires that the informational coupling term increase at least as fast as the induced mechanical penalty.

Local composition over multi-step plans. Consider a sequence of T small, feasible acts by j with directions ξ_1, \dots, ξ_T applied in successive cycles, each evaluated on the then-current $\sigma = 0$ state with least-action completion. Let $\Delta V_i^{(t)}$ denote i 's first-order value change at step t . By the chain rule and stagewise feasibility,

$$\Delta V_i^{(t)} = D_j V_i[\xi_t] \cdot \delta t + o(\delta t).$$

Summing and letting $\delta t \rightarrow 0$,

$$\sum_{t=1}^T \Delta V_i^{(t)} = \sum_{t=1}^T D_j V_i[\xi_t] \cdot \delta t + o(1). \quad (9)$$

⁶This is the same projection principle used in the definition of ΔS (Section 4): among all globally feasible completions, pick the one that minimizes per-cycle action.

Hence if $D_j V_i[\xi_t] \geq 0$ at each step (consent holds locally throughout), then the cumulative first-order change satisfies $\sum_t \Delta V_i^{(t)} \geq 0$. In words: *local consent composes*—a plan that is consent-respecting at every tick is consent-respecting in aggregate to first order. Concavity of V (Section 5, Axiom A3) further implies that mid-course coarsenings or mixtures cannot make the cumulative change more negative than this first-order sum.

Auto-rescind (stopping at the first sign flip). Define the *consent stopping time* for i along a smooth feasible path γ of j ’s acts by

$$\tau^*(i \leftarrow j; \gamma) := \inf \{t \geq 0 : D_j V_i[\dot{\gamma}(t)] < 0\},$$

with the convention $\inf \emptyset = +\infty$. Consent holds on $[0, \tau^*)$ and *rescinds automatically* at $t = \tau^*$ when the directional derivative turns negative. Past τ^* , continuation without repair either violates consent or must be re-justified by the global selection rule of Section 7 (feasibility \rightarrow minimax $\max \Delta S \rightarrow$ welfare \rightarrow robustness), which can overrule individual consent only when no consent-respecting feasible alternative exists and the minimax harm requirement compels a specific tradeoff.

Non-competent agents (proxy consent). For agents who cannot supply a reliable local model of their own value (infants, the severely impaired, non-linguistic animals), consent is evaluated against a *conservative proxy* \underline{V}_i constructed as follows. Let \mathcal{M}_i be the set of V -models compatible with (i) the four axioms of Section 5, (ii) observed coupling statistics for i ’s class, and (iii) gauge invariance on the ledger. Define the lower envelope

$$\underline{V}_i(p_{AE}, x) := \inf_{M \in \mathcal{M}_i} V_i^{(M)}(p_{AE}, x).$$

Then proxy consent holds when $D_j \underline{V}_i[\xi] \geq 0$. This rule is *conservative*: it never licenses an act that every admissible V -model would deem locally harmful, and it introduces no free parameter beyond those already fixed by the RS scaffold. In emergencies where continuing on the current path drives \underline{V}_i below a universal safety floor during a finite eight-tick horizon, the selection rule of Section 7 applies: an *override-with-repair* plan may be chosen that minimizes the worst externalized surcharge $\max \Delta S$ while restoring feasibility and raising \underline{V}_i above the floor as quickly as the cadence allows; all such overrides carry an explicit, auditable repayment path back to $\sigma=0$.

Interpretation (plain). Consent in an RS world is not a mood or a permission slip; it is the sign of a derivative taken on the reciprocity-conserving manifold with least-action completion. If your contemplated move leaves my value non-decreasing *by my own axiology*, you have my consent—right up to the point where the sign flips, at which instant it is withdrawn. For those who cannot speak for themselves, the same logic is applied to a conservative proxy that never overestimates benefit. In all cases, the definition is local, compositional across steps, invariant under the bridge, and free of tunable weights.

7 The decision principle (lexicographic, parameter-free)

We now formalize the selection rule for choosing among admissible futures. The rule is *lexicographic* and *parameter-free*: first enforce reciprocity conservation, then minimize the worst externalized surcharge, then maximize a fixed cardinal welfare, then prefer the most robust reciprocity network, then apply the RS φ -tier arithmetic for any residual tie. No weights are introduced at any stage.

Feasible actions and induced quantities. Fix a $\sigma = 0$ ledger state at the beginning of an eight-tick cycle and let \mathcal{A}_σ denote the nonempty set of σ -*feasible* actions (controls, policies, or interventions) that produce a balanced, reciprocity-conserving completion over the next cycle.⁷ For each $a \in \mathcal{A}_\sigma$ define:

$$H(a) := \max_{i,j} \Delta S(i \rightarrow j \mid a), \quad (10)$$

$$W(a) := \sum_i f(V_i \mid a), \quad (11)$$

$$R(a) := \lambda_2(L_\sigma \mid a), \quad (12)$$

where ΔS is the externalized action surcharge of Section 4 (computed on the per-cycle least-action completion), V_i is the forced axiology of Section 5 for agent i , f is the curvature-normalized concave transform discussed below, and $\lambda_2(L_\sigma)$ is the spectral gap of the σ -graph Laplacian induced by the reciprocity network after applying a (larger gap means fewer routes for skew cycles to re-emerge under bounded shocks). Let $\tau(a)$ denote the fixed RS φ -tier tie-break score associated with a (a discrete rank determined by the same arithmetic used elsewhere in the RS calculus).

The lexicographic selection rule. The chosen action a^* is defined by the following sequence:

$$\begin{aligned} \mathcal{A}_1 &:= \arg \min_{a \in \mathcal{A}_\sigma} H(a), \\ \mathcal{A}_2 &:= \arg \max_{a \in \mathcal{A}_1} W(a), \\ \mathcal{A}_3 &:= \arg \max_{a \in \mathcal{A}_2} R(a), \\ a^* &:= \arg \max_{a \in \mathcal{A}_3} \tau(a), \end{aligned} \quad (13)$$

with the understanding that any of the sets \mathcal{A}_k may already be singletons at earlier stages. The sequence (13) contains no tunable tradeoffs: feasibility is enforced first; among feasible actions the worst off-loading of cost is minimized; among those, the fixed cardinal welfare is maximized; among those, network robustness is maximized; any remaining tie is broken by the predetermined φ -tier arithmetic.

On the concave transform f . The axiology V in Section 5 is cardinal and curvature-normalized. To aggregate across agents with diminishing returns while introducing no new scale, the map $f : \mathbb{R} \rightarrow \mathbb{R}$ is fixed by the same normalization— $f(0) = 0$, $f'(0) = 1$, $f''(0) = -1$ —and by the requirement that f be concave, strictly increasing, and gauge-invariant under the RS bridge. Any two choices f satisfying these conditions differ only at higher order by terms that would introduce an inadmissible scale if tuned; the RS convention fixes f once (up to the φ -tier already chosen to normalize V), so W in (11) contains no free weight.

σ -Completeness (existence and optimality). We now record the basic well-posedness and optimality guarantee for the lexicographic rule.

⁷If $\mathcal{A}_\sigma = \emptyset$, the repair procedure of Section 8 applies: one computes a least-action path that returns to $\sigma = 0$ while minimizing accumulated $\max \Delta S$. The present section assumes $\mathcal{A}_\sigma \neq \emptyset$.

Proposition 7.1 (σ -Completeness). *Assume \mathcal{A}_σ is finite or compact and that H , W , and R are continuous on \mathcal{A}_σ .⁸ Then the sequence (13) yields a nonempty \mathcal{A}_1 , \mathcal{A}_2 , \mathcal{A}_3 , and a well-defined choice a^* . Moreover, for any $a \in \mathcal{A}_\sigma$, either a is infeasible (violates $\sigma=0$) or*

$$(H(a), -W(a)) \succeq_{\text{lex}} (H(a^*), -W(a^*)), \quad (14)$$

with equality in (14) only if $R(a) \leq R(a^)$, and if $R(a) = R(a^*)$ then $\tau(a) \leq \tau(a^*)$. In particular, any deviation from a^* either breaks feasibility or strictly worsens at least one of the first two components ($\max \Delta S$, $\sum f(V)$).*

Proof. By Weierstrass' theorem, H attains a minimum on compact \mathcal{A}_σ (or trivially on finite \mathcal{A}_σ), so $\mathcal{A}_1 \neq \emptyset$. The set \mathcal{A}_1 is closed; W attains a maximum on it, hence $\mathcal{A}_2 \neq \emptyset$. The same reasoning yields $\mathcal{A}_3 \neq \emptyset$ by maximizing R on \mathcal{A}_2 . Finally, τ totally orders \mathcal{A}_3 by the fixed RS arithmetic, so a^* exists. For any $a \in \mathcal{A}_\sigma$, if $H(a) > H(a^*)$, then (14) holds strictly. If $H(a) = H(a^*)$ but $W(a) < W(a^*)$, then again (14) holds strictly. If H and W are tied, then by construction $R(a) \leq R(a^*)$, with equality only when $\tau(a) \leq \tau(a^*)$. If $a \notin \mathcal{A}_\sigma$, it violates $\sigma=0$ feasibility and is excluded at stage zero. \square

Why no weighted blends appear. Any attempt to collapse H and W into a single scalar—say, by minimizing $H + \lambda(-W)$ for some $\lambda > 0$ —would introduce a new, arbitrary scale λ not fixed by the bridge or the curvature normalization and would therefore violate the parameter-free premise. Likewise, any convex combination $\alpha W + \beta R$ at later stages would smuggle in a tradeoff between welfare and robustness that the RS scaffold does not authorize. The lexicographic order (13) is thus not a stylistic preference; it is the unique way to compare actions without inventing knobs.

Interpretation (plain). The rule first asks whether a move keeps reciprocity conserved. If not, it is out. If several moves do, the next question is: which one least offloads cost onto anyone else? Among those, which one most increases the fixed, physics-level measure of how well lives fit their world? If there is still a tie, which arrangement is hardest to knock off balance in the next storm? If even that ties, use the same φ -tier arithmetic the rest of RS uses to decide between equal claims. There are no sliders to fiddle with; the universe set them already.

8 Time, debt, and repair paths (no ad-hoc discount)

Moral evaluation in a recognition-structured world must aggregate across time without smuggling in preferences. The ledger fixes a *minimal nontrivial period* of eight ticks (Section 2). That cadence is not a convention; it is a combinatorial invariant of admissible evolution in $D = 3$. Consequently, all temporal accounting is done in units of the eight-tick cycle. There is no lawful continuous discount rate or arbitrary weighting scheme: any such choice would introduce a new scale alien to the bridge.

Debt and initial infeasibility. The conservation law of Section 3 says that admissible worldlines live on $\sigma = 0$. Nonetheless, one may find oneself in a state that carries *legacy skew* (for example, after diagnosing an accumulated imbalance): $\sigma[C_0] \neq 0$ at the present cycle C_0 . In that case the selection rule of Section 7 has an empty feasible set $\mathcal{A}_\sigma = \emptyset$ by definition, and a *repair plan* is required. Intuitively, skew is a debt of action owed back to the ledger; repair means clearing that debt while doing the least further harm.

⁸These are the natural regularity conditions: ΔS and V are continuous under the least-action projection on the σ -manifold; the spectral gap is continuous under bounded perturbations of edge weights in the σ -graph.

The repair problem (least action under cadence and consent). Let (C_0, \dots, C_T) be a horizon of T consecutive eight-tick cycles. A *repair path* $\gamma = (a_0, \dots, a_T)$ is a sequence of controls such that:

- (i) the resulting ledger evolution is balanced at each site (Section 2) and ends on the reciprocity manifold: $\sigma[C_T] = 0$,
- (ii) for each cycle, the least-action completion is applied to non-acting parties (as in Sections 4–6),
- (iii) consent is respected where possible (Section 6); when local consent cannot be maintained for all affected parties there must be an explicit *override-with-repair* clause that is itself part of the plan and is justified solely by the minimax harm criterion below.

Write $H(a_t) := \max_{i,j} \Delta S(i \rightarrow j \mid a_t)$ for the worst externalized surcharge on cycle t . The *repair objective* is the cadence-lawful aggregate

$$\mathcal{J}_{\text{rep}}(\gamma) = \sum_{t=0}^T H(a_t) \quad (\text{sum over eight-tick cycles}). \quad (15)$$

A *minimal repair* solves

$$\gamma^* \in \arg \min_{\gamma} \mathcal{J}_{\text{rep}}(\gamma) \quad \text{subject to (i)–(iii) and the ledger dynamics.} \quad (16)$$

This makes “reparations” a least-action problem: clear skew as fast as the cadence and consent allow, while minimizing the accumulated worst off-loading of cost.

Existence (compact horizon). If the per-cycle feasible control sets are compact and the maps $a_t \mapsto \Delta S(\cdot \mid a_t)$ are continuous under the least-action completion, then (16) admits a minimizer by Weierstrass’ theorem (stagewise compactness and continuity imply compactness of the product set and continuity of (15)). If multiple minimizers exist, the lexicographic refinements of Section 7 (welfare, robustness, φ -tier) are applied cycle-wise to pick among them without introducing weights.

No-Arbitrary-Discount theorem. We now formalize why (15) is the only lawful temporal aggregator.

Theorem 8.1 (No-Arbitrary-Discount). Let \mathfrak{A} map any finite sequence of per-cycle harms $\mathbf{h} = (h_0, \dots, h_T)$ (with $h_t \geq 0$) to a real score used to rank repair paths. Suppose \mathfrak{A} satisfies:

- (I) **Gauge invariance under the bridge:** for any joint re-anchoring $(\tau_0, \ell_0) \mapsto (s\tau_0, s\ell_0)$ that preserves $c = \ell_0/\tau_0$, the ranking induced by \mathfrak{A} on any pair of sequences is unchanged.
- (II) **Cadence invariance:** cyclic rotation by any multiple of the eight-tick period does not change the score (no privileged phase within the period), and concatenation preserves monotonicity.
- (III) **Separability and monotonicity:** \mathfrak{A} is strictly increasing in each coordinate and additive on independent horizons, i.e., $\mathfrak{A}(\mathbf{h} \oplus \mathbf{k}) = \mathfrak{A}(\mathbf{h}) + \mathfrak{A}(\mathbf{k})$ when the underlying ledgers are independent.

Then there exists a constant $\alpha > 0$ such that, for all \mathbf{h} ,

$$\mathfrak{A}(\mathbf{h}) = \alpha \sum_{t=0}^T h_t. \quad (17)$$

In particular, any exponential, hyperbolic, or otherwise non-uniform discount $\sum w_t h_t$ with w_t depending on a new time scale or on a phase not fixed by the eight-tick cadence violates at least one of (I)–(III).

Proof (sketch). By (II), \mathfrak{A} is invariant under eight-phase rotation; hence any per-cycle weights must be constant across the period. By concatenation monotonicity and additivity (III), \mathfrak{A} must be linear in block sums: for any nonnegative scalar λ and any horizon length n , $\mathfrak{A}(\lambda \mathbf{1}_n) = n \mathfrak{A}(\lambda \mathbf{1}_1)$. By independence additivity, \mathfrak{A} is Cauchy-additive on finitely supported sequences with nonnegative entries; monotonicity restricts to the positive cone, yielding linearity with a common coefficient $\alpha > 0$. Any attempt to use a nonconstant weight sequence (w_t) would either (a) select a privileged phase within the eight-tick period (violating cadence invariance), or (b) depend on a rate parameter with units of inverse time (exponential or hyperbolic discount), which changes under $(\tau_0, \ell_0) \mapsto (s\tau_0, s\ell_0)$ and thus violates gauge invariance (I). Hence (17). \square

Consequences. Two practical consequences follow. First, there is no lawful way to “justify” delaying repair by appealing to a discount rate; the eight-tick cadence fixes the only aggregation, so later pain cannot be traded against earlier pain by choosing a favorable clock. Second, when forced to override local consent to avoid a safety-floor breach (Section 6), the only admissible justification is that no other path achieves a smaller *undiscounted* $\sum H(a_t)$ while restoring $\sigma = 0$ within the cadence constraints. Anything else would be importing a new scale.

Interpretation (plain). Time, in this framework, is not a lever to tilt the books. The universe gives you a beat—eight ticks—and that is the rhythm on which debts are counted and cleared. If you inherit skew, you must plan a path that clears it with the least possible worst burden at each step, keeping reciprocity and consent in view. You cannot shrink tomorrow’s bill by choosing a fancy discount; there is no such knob in the physics. The only acceptable mercy is to spread the load fairly over the cadence the world already keeps.

9 From persons to polities: σ -graphs and robustness

We lift the reciprocity calculus from pairs to societies. The goal is to represent a whole polity’s reciprocity structure at $\sigma = 0$, quantify how hard it is for skew to re-enter under shocks, and show that this *robustness* is captured by a single spectral quantity of a graph Laplacian. Among σ -neutral, welfare-maximizing arrangements (Section 7), the more robust one is preferred because it minimizes the worst future externalized surcharge under bounded disturbances.

Aggregation and the σ -graph. Partition the ledger’s sites into N disjoint, agent/group domains $\{\mathcal{E}_i\}_{i=1}^N$. At $\sigma = 0$ (Section 3), every bidirectional exchange between domains is locally balanced on each eight-tick cycle. Linearizing about this balanced state and projecting to least-action completions (Sections 4–6), small, antisymmetric inter-domain perturbations induce a quadratic action

$$\delta^2 S = \frac{1}{2} \sum_{1 \leq i < j \leq N} w_{ij} (u_i - u_j)^2,$$

where u_i are domain potentials (log-coordinates) and $w_{ij} = w_{ji} \geq 0$ are the induced *reciprocity conductances* that summarize the society’s baseline coupling at $\sigma = 0$.⁹ Define the weighted,

⁹Concretely: w_{ij} is the second variation (Hessian coefficient) of the least-action functional with respect to an antisymmetric exchange between domains i and j , holding all other domains at $\sigma = 0$. The normalization $J''(1) = 1$

undirected σ -graph $G_\sigma = (V, E, W)$ with vertices $V = \{1, \dots, N\}$ and weights $W = [w_{ij}]$. Its (combinatorial) Laplacian is

$$L_\sigma = D - W, \quad D_{ii} = \sum_{j \neq i} w_{ij}, \quad D_{ij} = 0 \ (i \neq j).$$

Feasibility and skew cycles. A *skew cycle* is a closed walk in G_σ that carries a nonzero oriented sum of log-multipliers (net skew) over one eight-tick period. By Section 3, any such cycle incurs a strictly positive, avoidable action surcharge due to the strict convexity and symmetry of J ; least-action dynamics therefore exclude persistent skew cycles. Thus, at feasibility, the σ -graph is used only in its *balanced* linear response (no net circulations): all observed exchanges live on the $\sigma=0$ manifold, and G_σ summarizes how the polity would redistribute small shocks while trying to remain balanced.

Shocks and redistribution. Model an exogenous, bounded disturbance over a cycle by a mean-zero vector $s \in \mathbb{R}^N$ (injection at some domains, extraction at others), with $\sum_i s_i = 0$ and $\|s\|_2 \leq S$. The least-action redistribution that keeps the polity balanced solves the weighted Poisson equation

$$L_\sigma u = s, \quad u \perp \mathbf{1}, \quad (18)$$

and induces inter-domain fluxes proportional to edge differences $u_i - u_j$. The per-cycle quadratic action of this response is

$$\mathcal{E}(u) = \frac{1}{2} \sum_{i < j} w_{ij} (u_i - u_j)^2 = \frac{1}{2} s^\top L_\sigma^+ s, \quad (19)$$

where L_σ^+ is the Moore–Penrose pseudoinverse on the mean-zero subspace.¹⁰ The *worst externalized surcharge* at a node in this cycle satisfies

$$\max_i \Delta S_i \leq \mathcal{E}(u), \quad (20)$$

since each node's share is bounded above by the total quadratic action.

Robustness as spectral gap. Let the eigenvalues of L_σ be $0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_N$. The *spectral gap* $\lambda_2(L_\sigma)$ (algebraic connectivity) controls how much the polity's potentials—and hence surcharges—can amplify a given shock. Indeed, restricting to the mean-zero subspace and using the Rayleigh quotient,

$$\mathcal{E}(u) = \frac{1}{2} s^\top L_\sigma^+ s \leq \frac{1}{2} \frac{\|s\|_2^2}{\lambda_2(L_\sigma)}. \quad (21)$$

Combining (20) and (21) yields the uniform bound

$$\max_i \Delta S_i \leq \frac{1}{2} \frac{\|s\|_2^2}{\lambda_2(L_\sigma)}. \quad (22)$$

Thus a larger spectral gap guarantees a smaller worst-case externalized surcharge under any bounded, mean-zero shock.

fixes the scale. Gauge invariance ensures that w_{ij} depends only on closed-chain data, not on node potentials.

¹⁰This is the standard Dirichlet principle: u minimizes $\mathcal{E}(u)$ subject to $L_\sigma u = s$; at the minimizer $\mathcal{E}(u) = \frac{1}{2} s^\top u = \frac{1}{2} s^\top L_\sigma^+ s$.

Robust-Preference lemma. We now formalize the selection consequence inside the lexicographic rule of Section 7.

Lemma 9.1 (Robust-Preference). *Consider two σ -neutral arrangements a and b on the same agent set that are tied on welfare (Section 7): $W(a) = W(b)$, and satisfy $H(a) = H(b)$ at the current cycle. If $\lambda_2(L_\sigma | a) > \lambda_2(L_\sigma | b)$, then for any bounded, mean-zero shock s with $\|s\|_2 \leq S$, the least-action responses $u^{(a)}$ and $u^{(b)}$ satisfy*

$$\max_i \Delta S_i^{(a)} \leq \frac{S^2}{2\lambda_2(L_\sigma | a)} < \frac{S^2}{2\lambda_2(L_\sigma | b)} \geq \max_i \Delta S_i^{(b)}. \quad (23)$$

In particular, among σ -neutral V -maximizers, the arrangement with larger spectral gap minimizes the worst-case future externalized surcharge under bounded shocks.

Proof. Apply (22) to a and b separately and compare the denominators. The strict inequality follows from $\lambda_2(L_\sigma | a) > \lambda_2(L_\sigma | b)$. \square

Compatibility with the parameter-free scaffold. All quantities above are fixed by the RS bridge and the unique cost J : the weights w_{ij} are second-variation coefficients of the least-action functional at $\sigma=0$; the Laplacian is determined by W ; the spectral gap has no tunable parameters. No new scales are introduced.

Interpretation (plain). A society’s reciprocity network can be “tight” or “loose.” Tight means that if you poke it—inject a bounded shock somewhere—it spreads the strain quickly and thinly, so no one place pays a large bill; loose means pockets of the network can trap strain and dump big bills on someone. The tightness is the spectral gap. When two σ -neutral, welfare-maximizing designs tie on the present, you pick the one with the larger gap because physics guarantees it will treat tomorrow’s disturbances more gently. That choice is not a taste about resilience; it is the unique knob-free implication of how a balanced ledger dissipates strain.

10 How to audit a moral claim (measurement protocol)

Moral claims in a recognition-structured universe are not debated; they are *audited*. This section gives a concrete, reproducible recipe for evaluating any proposed action or policy over one eight-tick cycle, and for multi-cycle plans. Every step is conducted on the reciprocity-conserving manifold (Sections 2–3), with least-action completion (Sections 4–6), and without ad-hoc temporal discount (Section 8). Error bars are explicit: the decision must be *robust* to bounded uncertainty in ΔS and V .

Inputs and invariances. An audit instance consists of: (i) a pre-action ledger snapshot (p_{AE}, x) at the start of a cycle; (ii) a specification of the contemplated action a (which bonds are acted on, by whom, and how); (iii) the partition into agents/groups $\{\mathcal{E}_i\}$; (iv) a coarse-graining for (A, E) that induces p_{AE} .¹¹ All computations are invariant under ledger gauge and under re-anchoring $(\tau_0, \ell_0) \mapsto (s\tau_0, s\ell_0)$ that preserves $c = \ell_0/\tau_0$; numerical tolerances are reported but do not alter the logical tests.

¹¹Any two coarse-grainings related by an invertible relabeling yield the same mutual information; V is gauge-invariant by Section 5.

Step 1: Feasibility (reciprocity conservation). Compute the cycle-wise reciprocity skew $\sigma_{ij}[C_{\text{pre}}]$ for all ordered pairs (i, j) on the pre-action cycle and $\sigma_{ij}[C_{\text{post}}]$ on the post-action cycle obtained by least-action completion of a . Define $\sigma[C] := \sum_{i < j} |\sigma_{ij}[C]|$. Feasibility requires

$$\sigma[C_{\text{pre}}] = 0, \quad \sigma[C_{\text{post}}] = 0.$$

If $\sigma[C_{\text{pre}}] \neq 0$, Section 8's repair procedure is invoked (no selection occurs until the skew debt is cleared). If $\sigma[C_{\text{post}}] \neq 0$, the action fails the audit.

Step 2: Harm (externalized action surcharge). Compute the matrix $\Delta S(i \rightarrow j \mid a)$ by (3): the difference in j 's required action between the least-action completion with i 's act in place and with the neutral act, holding $\sigma=0$. Report the full matrix and its maximum

$$H(a) = \max_{i,j} \Delta S(i \rightarrow j \mid a).$$

For transparency, include quantiles of the nonzero entries of ΔS , and a decomposition by independent subsystems when applicable (additivity, Section 4).

Step 3: Welfare (forced axiology). For each agent i , compute the value V_i using the decomposition (7): mutual information $I(A; E)$ for the chosen coarse-graining, minus the J -induced curvature penalty \mathcal{C}_J^* associated with the post-action least-action state; use the fixed φ -tier normalization for the information scale. Aggregate with the curvature-normalized concave transform f (Section 7):

$$W(a) = \sum_i f(V_i \mid a).$$

Report both V_i and $f(V_i)$ for all i , and the change relative to pre-action values.

Step 4: Robustness (reciprocity network). Construct the σ -graph G_σ by evaluating the second variation of the least-action functional at $\sigma=0$ (Section 9), yielding weights w_{ij} and Laplacian L_σ . Compute the spectral gap $\lambda_2(L_\sigma)$ before and after the action; report $\Delta\lambda_2 = \lambda_2^{\text{post}} - \lambda_2^{\text{pre}}$. This quantifies how the action changes the polity's resistance to re-introduction of skew under bounded shocks (Lemma 9.1).

Step 5: Consent (local derivative test). For each affected pair (i, j) , evaluate the directional derivative $D_j V_i[\xi]$ of i 's value along j 's contemplated feasible direction ξ with least-action completion (Section 6). Record

$$C(i \leftarrow j) \iff D_j V_i[\xi] \geq 0.$$

If an agent is non-competent, replace V_i by the conservative proxy \underline{V}_i (lower envelope over admissible models; Section 6) and repeat the test. Consent failures are flagged with the auto-rescind time if the derivative changes sign along a multi-step plan.

Step 6: Time (multi-cycle plans; no discount). For a horizon $\{a_t\}_{t=0}^T$ of actions, compute the per-cycle worst surcharge $H(a_t)$ and the undiscounted sum

$$\mathcal{J}_{\text{rep}}(\{a_t\}) = \sum_{t=0}^T H(a_t).$$

Verify cadence compliance: per-cycle feasibility, least-action completion, and (when initial skew exists) the repair constraints of Section 8. Exponential, hyperbolic, or phase-dependent weights are not permitted (Theorem 8.1).

Step 7: Uncertainty and robustness margins. Let $\mathcal{U}_{\Delta S}$ and \mathcal{U}_V be uncertainty sets for ΔS and V (e.g., confidence intervals from repeated measurements, interval bounds from Lipschitz estimates of models).¹² Define worst- and best-case summaries on these sets:

$$\begin{aligned} H^{\sup}(a) &:= \sup_{\theta \in \mathcal{U}_{\Delta S}} H_\theta(a), & H^{\inf}(a) &:= \inf_{\theta \in \mathcal{U}_{\Delta S}} H_\theta(a), \\ W^{\sup}(a) &:= \sup_{\phi \in \mathcal{U}_V} W_\phi(a), & W^{\inf}(a) &:= \inf_{\phi \in \mathcal{U}_V} W_\phi(a). \end{aligned}$$

A lexicographic choice a^* is *robustly certified* if

$$H^{\sup}(a^*) < \min_{b \neq a^*} H^{\inf}(b),$$

or, when the first margins tie,

$$H^{\sup}(a^*) = \min_b H^{\inf}(b) \quad \text{and} \quad W^{\inf}(a^*) > \max_{b: H^{\inf}(b) = H^{\inf}(a^*)} W^{\sup}(b),$$

with the spectral-gap and φ -tier refinements applied only within remaining ties. If these strict inequalities fail, the audit returns *indeterminate*: more measurement is required (or a different action must be proposed) before the parameter-free rule can discriminate without risk of reversal under uncertainty.

Output record (what is archived). Each audit emits a signed, content-addressed bundle containing: (i) pre/post σ traces; (ii) the ΔS matrix, its maximum H , and uncertainty bands; (iii) $\{V_i\}$, $\{f(V_i)\}$, their aggregates W , and uncertainty bands; (iv) pre/post $\lambda_2(L_\sigma)$ and $\Delta\lambda_2$; (v) consent derivative signs and any proxy use; (vi) cadence compliance proofs and, where applicable, the repair objective \mathcal{J}_{rep} ; (vii) the final lexicographic decision with robustness margins. All numbers are reproducible from the ledger, the coarse-graining declaration, and the least-action solver’s transcript.

Falsifiers (when the audit rejects the framework). The audit protocol also exposes defeats of the theory. A policy that passes Steps 2–6 while exhibiting persistent $\sigma \neq 0$ at Step 1 contradicts Section 3. A stable, gauge-invariant procedure that lowers total action by maintaining skew contradicts the convexity-based derivation of reciprocity conservation. A distinct axiology that satisfies the four axioms of Section 5 and systematically outpredicts V on audited instances would falsify the uniqueness claim. A time aggregator other than the undiscounted eight-tick sum that preserves the audit ranking under gauge and cadence invariance would contradict Theorem 8.1.

Interpretation (plain). An honest audit is a short list of numbers with tolerances and a yes/no on each invariant. Does the move keep reciprocity conserved? How big is the worst externalized bill? Does the fixed, physics-level value go up overall? Is the reciprocity network harder to corrupt after the move? Do the people it touches consent, by their own axiology or a conservative proxy? If it’s a plan, does it respect the universe’s beat? And do all these answers survive their error bars? If so, the claim “this is the right thing to do” is not rhetoric; it is a statement the ledger itself will not contradict.

¹²Bounds must be gauge-invariant; for ΔS , envelope the optimizer’s numerical tolerance; for V , use concentration bounds for empirical $I(A; E)$ and quadratic error bars for \mathcal{C}_j^* in the small-strain regime.

11 Two worked cases (physics, not rhetoric)

We now exhibit two policy-neutral examples that force the calculus. In both, we work on the $\sigma=0$ manifold, apply least-action completion at each cycle, and evaluate the quantities defined in Sections 4–7. Numbers are illustrative and chosen so the logic is visible at a glance; no new scales are introduced.

Case A: a “greater-good” plan raises average V but hides skew

Baseline. Three agents (or domains) A, B, C exchange on a balanced ledger. At the start of cycle C_0 , reciprocity is conserved, $\sigma[C_0] = 0$. Their values (forced axiology, Section 5) are

$$V_A = 1.20, \quad V_B = 1.00, \quad V_C = 1.00,$$

measured in the fixed φ -normalized units; for small changes we take $f(V) \approx V$ (the curvature-normalized concave transform with $f'(0) = 1, f''(0) = -1$). Thus $W(C_0) = 3.20$.

The proposal P_{good} (hidden skew). In C_1 , B and C strengthen their coupling (higher mutual information) at the expense of unilateral draws on A . In bond multipliers this appears as

$$x_{A \rightarrow B} = 1 + \varepsilon, \quad x_{B \rightarrow A} = 1, \quad x_{A \rightarrow C} = 1 + \varepsilon, \quad x_{C \rightarrow A} = 1, \quad \varepsilon > 0,$$

with all other active bonds rebalanced by least action. Naively, coarse-grained values improve

$$\Delta V_B = +0.15, \quad \Delta V_C = +0.15, \quad \Delta V_A = -0.02,$$

so W would rise by $\Delta W \approx +0.28$.

But feasibility fails immediately: the cycle-wise skew is positive,

$$\sigma_{AB}[C_1] = \ln(1 + \varepsilon) > 0, \quad \sigma_{AC}[C_1] = \ln(1 + \varepsilon) > 0,$$

hence $\sigma[C_1] = |\sigma_{AB}| + |\sigma_{AC}| > 0$. By Section 3, any such skew carries a strictly positive, avoidable action surcharge, and P_{good} is *inadmissible* at Step 1 of the audit (Section 10). No downstream gains in W or robustness can rescue it.

The repair-first variant P_{rep} (selected). Modify the plan so that the same informational benefits are realized while clearing skew within the cycle. One least-action way is to pair each $A \rightarrow B$ and $A \rightarrow C$ adjustment with a matched return in the same cycle (or across a two-cycle window with eight-tick cadence):

$$x_{A \rightarrow B} = 1 + \varepsilon, \quad x_{B \rightarrow A} = 1 - \varepsilon, \quad x_{A \rightarrow C} = 1 + \varepsilon, \quad x_{C \rightarrow A} = 1 - \varepsilon,$$

with the rest completed by the least-action projector. Then $\sigma[C_1] = 0$ by construction; feasibility passes. Because rebalancing is not free, the welfare gain is slightly smaller than the naive $+0.28$; a typical least-action completion yields

$$\Delta V_B = +0.13, \quad \Delta V_C = +0.13, \quad \Delta V_A = -0.02, \quad \Rightarrow \quad \Delta W \approx +0.24.$$

Compute the per-cycle harm matrix $\Delta S(i \rightarrow j \mid P_{\text{rep}})$ from (3); in a representative outcome,

$$\Delta S(A \rightarrow B) = \Delta S(A \rightarrow C) = 0, \quad \Delta S(B \rightarrow A) = \Delta S(C \rightarrow A) = 0.40, \quad \Delta S(\text{else}) = 0,$$

so $H(P_{\text{rep}}) = \max_{i,j} \Delta S(i \rightarrow j) = 0.40$. For P_{good} this quantity is undefined (feasibility fails). The σ -graph spectral gap (Section 9) is unchanged or modestly improved by the additional symmetric ties; a typical calculation gives $\Delta\lambda_2 \approx +0.03$.

Lexicographic selection. Step 1 discards P_{good} ($\sigma \neq 0$). Among feasible actions, P_{rep} minimizes H versus any other way of clearing the same skew (by the Dirichlet/least-action property), and its W is the largest among those ties. It is therefore selected at once by the rule of Section 7, without appeal to weights or narratives. If the skew had been inherited from C_0 instead of introduced in C_1 , the same construction would arise as the solution of the repair problem (16), minimizing $\sum_t H(a_t)$ under the cadence.

Case B: a consent-sensitive plan where intuition flips on $D_j V_i$ and $\max \Delta S$

Baseline. Two domains D and R (think “developer” and “resident” as neutral roles) exchange on a balanced ledger. At C_0 , $V_D = 1.30$, $V_R = 1.10$, $\sigma[C_0] = 0$.

The tempting plan Q (fails consent and harm). In C_1 , D proposes an act ξ that streamlines its own coupling to the environment by rerouting shared resources through R ’s boundary bonds. The coarse-grained effect on R is a slight degradation of p_{AE} (fewer clean, informative interactions) with a small mechanical strain on R ’s domain (compensating flows). The *directional derivative* of R ’s value along D ’s act is

$$D_D V_R[\xi] = \frac{d}{dt} \Big|_0 \left(\kappa I_R(t) - \mathcal{C}_J^\star(t) \right) = \kappa I'_R(0) - \underbrace{\frac{d}{dt} \mathcal{C}_J^\star \Big|_0}_{\leq 0},$$

where $I'_R(0) < 0$ because the channel degrades. In a representative computation, $D_D V_R[\xi] = -0.03 < 0$. *Consent fails* at Step 5 (Section 10): $C(R \leftarrow D)$ does not hold.¹³ Independently, the per-cycle harm matrix under least-action completion shows a concentrated surcharge on R ’s domain,

$$\Delta S(D \rightarrow R \mid Q) = 1.20 \quad (\text{arbitrary units}), \quad \Delta S(\text{else}) \leq 0.10,$$

so $H(Q) = 1.20$. Meanwhile W rises only modestly (e.g., $\Delta W \approx +0.05$) because the developer’s gain is offset by the resident’s loss.

The safe alternative Q_{safe} (selected). A variant staggers the rerouting over two cycles and introduces a small, symmetric auxiliary tie that preserves R ’s informational coupling. Locally, $D_D V_R[\xi_{\text{safe}}] \geq 0$ (the first-order change in R ’s value is nonnegative), and the least-action completion spreads the surcharge:

$$\Delta S(D \rightarrow R \mid Q_{\text{safe}}) = 0.55, \quad H(Q_{\text{safe}}) = 0.60,$$

with a slightly smaller welfare gain, say $\Delta W \approx +0.03$. The σ -graph gap is unchanged.

Lexicographic selection. Both plans are σ -feasible at Step 1. At Step 2, Q_{safe} strictly minimizes $\max \Delta S$; Q is eliminated. Consent checks corroborate the choice: Q violates local consent, while Q_{safe} passes. The small advantage in W for Q is irrelevant because it appears only after feasibility and minimax harm. The RS rule therefore picks the plan popular intuition often misses: the one that is locally consent-respecting and that lowers the worst externalized bill, even if its headline average gain is slightly smaller.

¹³If R were non-competent, the same conclusion would follow from the conservative proxy \underline{V}_R (Section 6).

Robustness to uncertainty (both cases). In each case, if ΔS and V are known only within bounded sets $\mathcal{U}_{\Delta S}$ and \mathcal{U}_V (Step 7, Section 10), the same decisions persist provided the robust margins hold:

$$H^{\sup}(P_{\text{rep}}) < \min_{b \neq P_{\text{rep}}} H^{\inf}(b), \quad H^{\sup}(Q_{\text{safe}}) < H^{\inf}(Q).$$

If these fail, the audit returns *indeterminate* and demands more measurement, not persuasion. In all cases, no weights are introduced; the decisions follow from reciprocity conservation, least-action harm, the forced axiology, consent as a derivative sign, and—when tied—the robustness of the σ -graph.

12 What could falsify this (clean failure modes)

A physics-level moral law must admit sharp defeats. Three crisp failure modes would invalidate the framework as stated.

(F1) A durable $\sigma \neq 0$ process with lower action than any $\sigma = 0$ alternative. Section 3 derived reciprocity conservation from the convexity and symmetry of J : any persistent skew raises total action on a cycle and can be pairwise-smoothed away at lower cost. A defeat would exhibit, for some class of boundary data, a balanced, repeatable process C^\sharp with $\sigma[C^\sharp] \neq 0$ such that for *every* $\sigma = 0$ completion \tilde{C} with the same endpoints one has $S[C^\sharp] < S[\tilde{C}]$. The test is operational: pre-register boundary data; implement both C^\sharp and a catalogue of $\sigma = 0$ completions generated by the least-action projector; measure per-cycle action (or a faithful surrogate) and show the strict inequality over many periods. A positive result contradicts the convexity-based derivation and falsifies reciprocity conservation in sustained evolution.

(F2) A distinct axiology satisfying the four axioms that outpredicts the MI-curvature V . Section 5 fixed V by four constraints: gauge invariance, additivity, concavity, and curvature normalization tied to $J''(1) = 1$, yielding a mutual-information-like coupling minus a J -induced penalty. A defeat would produce a *different* functional $\hat{V} \neq V$ that satisfies all four axioms and that, when slotted into the lexicographic rule (keeping feasibility and minimax harm unchanged), systematically yields *better* audited outcomes on preregistered instances—e.g., higher realized $\sum f(\cdot)$ at equal or lower max ΔS , or strictly larger post-action spectral gaps under matched feasibility constraints. The test is two-phase: (i) formal—exhibit \hat{V} and verify the axioms; (ii) empirical—publish a preregistered battery of audits, fix all uncertainty sets, and compare the ex ante choices induced by V and \hat{V} against ex post ledger measurements. A consistent advantage for \hat{V} falsifies the uniqueness claim.

(F3) A temporal aggregation law that respects RS invariances but differs from the eight-tick sum. Section 8 proved a No-Arbitrary-Discount result: under gauge and cadence invariance, separability, additivity, and monotonicity, the only lawful temporal aggregator for repair is the undiscounted sum over eight-tick cycles. A defeat would present a functional $\mathfrak{A} \neq \alpha \sum h_t$ that (i) is invariant under joint re-anchoring $(\tau_0, \ell_0) \mapsto (s\tau_0, s\ell_0)$, (ii) is invariant under rotation of the eight-tick phase and respects concatenation, (iii) is strictly increasing in each coordinate and additive on independent horizons, yet (iv) ranks at least one pair of harm sequences differently from the sum. The test is constructive: publish \mathfrak{A} with proofs of (i)–(iii); exhibit a pair of sequences (\mathbf{h}, \mathbf{k}) with $\sum h_t = \sum k_t$ but $\mathfrak{A}(\mathbf{h}) \neq \mathfrak{A}(\mathbf{k})$; verify in a matched repair setting that \mathfrak{A} 's preferred

plan does not introduce a hidden scale or privileged phase. If such an \mathfrak{A} exists, the cadence-based uniqueness claim fails.

In all three cases the burden is explicit and auditable: preregister boundary data and uncertainty sets; release code and logs; accept refutations that survive the protocol of Section 10.

13 Relation to classical ethics—why the paradoxes dissolve

The framework maps cleanly onto the major families of moral theory without inheriting their paradoxes because the physics fixes the order of operations and the scales.

Deontic constraint as physics. What deontological theories call a rule appears here as reciprocity conservation: $\sigma = 0$ is not a postulate but a conservation law derived from the unique cost J . Proposed actions that violate it are physically inadmissible in sustained trajectories; they are excluded before any talk of benefits.

Consequentialism with a single, physics-fixed cardinal. Once feasibility is secured, outcomes are ranked by a welfare functional that is *not* a negotiated utility mix but the unique cardinal axiology V fixed by gauge invariance, additivity, concavity, and curvature normalization. Aggregation across persons uses a concave transform f fixed by the same normalization. There are no tunable weights, so the usual arbitrariness of utilitarian sums does not arise.

Virtue as local improvement in value subject to feasibility. Traits and practices count insofar as they reliably increase V along σ -feasible directions. “Good character” reduces to local dynamics that, tick after tick, raise V without exporting costs (small positive $D_{\text{self}}V_{\text{others}}$ together with minimized $\max \Delta S$).

Why Arrow-type impossibilities do not bite. Arrow’s theorem concerns aggregation of *ordinal* preferences under axioms that assume unrestricted domains and independence conditions. Here there is no aggregation of arbitrary ordinals. The ranking is over a *constrained feasible set* (the $\sigma = 0$ manifold with least-action completion), using a *single* physics-fixed *cardinal* functional V (after a non-arbitrary concave transform). The theorem’s setup does not obtain, so its impossibility result does not transfer. Likewise, social choice paradoxes that rely on cycling of ordinal rankings are blocked by the lexicographic order of Section 7: feasibility and minimax harm take lexical precedence over welfare, and the robustness refinement further breaks cycles without introducing weights.

14 Limits and scope

The claims here are strong but bounded.

No account of qualia. The framework is a law about admissible action and ranking among feasible worlds on a recognition ledger. It is silent on private phenomenal experience, qualia, or the ontology of consciousness. Nothing in Sections 2–9 purports to derive subjective feel; only the measurable consequences of action enter.

Domain boundaries. Agents and groups are defined by partitions of the ledger. In practice, boundaries can be ambiguous (overlapping institutions, porous jurisdictions). The audit protocol handles this conservatively: partitions are preregistered; boundary rules for shared bonds are fixed and gauge-invariant; sensitivity analyses with alternative partitions are reported. If conclusions depend on boundary choice beyond uncertainty margins, the audit returns *indeterminate*.

Operational limits: model error and bypass channels. Estimating ΔS and V requires models of least-action completion and of agent–environment coupling. These carry error. Section 10 mandates uncertainty sets $\mathcal{U}_{\Delta S}$ and \mathcal{U}_V and robust selection margins; if margins fail, no decision is certified. Bypass channels—ways of acting that evade the auditor—are handled at the governance layer: only typed, audited channels are permitted to effect changes; proposed actions outside those channels are treated as unmeasured and rejected until instrumented.

Scope of “eternal.” “Eternal” means invariant under the RS bridge, parameter-free, and auditable—not immune to revision. Any of the falsifiers in Section 12, or future empirical discoveries that alter the physical scaffold (e.g., a different cost curvature than $J''(1) = 1$), would narrow or overturn claims. Within the assumed scaffold, however, the code does not drift with taste.

15 Conclusion: morality on the same shelf as energy and momentum

The spine is short and severe. In a recognition-structured universe the admissible worldlines conserve reciprocity; within that feasible set, the right thing to do is the action that minimizes the worst externalized surcharge and then maximizes a single, physics-fixed cardinal value, preferring arrangements that are harder to knock off balance. There are no knobs to tune, no discount rates to choose, no utility weights to argue over. The invariants that set c , λ_{rec} , and the curvature of J also set the terms of moral evaluation.

This is what “eternal” amounts to here: independence from fashion because the quantities are fixed by the same bridge that calibrates physics; openness to refutation because each claim is cashed out in audits and because sharp defeats are possible and welcome. If the world continues to behave as a recognition ledger with the properties assumed, morality belongs on the same shelf as energy and momentum: a conservation law with a unique way to read the meters, until the meters themselves teach us otherwise.

A Appendix A: σ -inequality details (full convexity derivation and multi-party extensions)

From multiplicative strain to a convex even function. Write each bond multiplier as $x_e = \exp(\alpha_e)$ with $\alpha_e \in \mathbb{R}$. Then

$$J(x_e) = \frac{1}{2}(x_e + x_e^{-1}) - 1 = \frac{1}{2}(e^{\alpha_e} + e^{-\alpha_e}) - 1 = \cosh(\alpha_e) - 1 =: \Phi(\alpha_e).$$

The generator $\Phi(\alpha) = \cosh \alpha - 1$ is C^∞ , strictly convex, even, and satisfies $\Phi(0) = 0$, $\Phi'(\alpha) = \sinh \alpha$ (odd), and $\Phi''(0) = 1$ (curvature normalization).

Two-move inequality (pair case). For any $\varepsilon \neq 0$,

$$J(1 + \varepsilon) + J(1 - \varepsilon) = \Phi(\ln(1 + \varepsilon)) + \Phi(-\ln(1 + \varepsilon)) = 2\Phi(\ln(1 + \varepsilon)) > 0.$$

Equivalently, for any nonzero α ,

$$\Phi(\alpha) + \Phi(-\alpha) = 2(\cosh \alpha - 1) > 0. \quad (24)$$

This is the strict convexity step used in the main text: any bidirectional imbalance around the unit point carries a strictly positive, avoidable action surcharge.

Closed-cycle inequality (multi-move, product-one constraint). Let $\{\alpha_1, \dots, \alpha_m\}$ be a finite family with *zero sum*,

$$\sum_{k=1}^m \alpha_k = 0 \iff \prod_{k=1}^m x_k = 1, \quad x_k = e^{\alpha_k}.$$

Then, by Jensen's inequality for the strictly convex \cosh ,

$$\sum_{k=1}^m \Phi(\alpha_k) = \sum_{k=1}^m (\cosh \alpha_k - 1) \geq m(\cosh(\frac{1}{m} \sum_k \alpha_k) - 1) = m(\cosh 0 - 1) = 0,$$

with equality if and only if $\alpha_k \equiv 0$ for all k (strict convexity). In ledger terms:

Given any closed chain of multipliers whose product is one (zero log-sum), the sum of per-bond costs is minimized exactly at the unit configuration $x_k \equiv 1$ and is strictly larger otherwise.

This is the multi-move version of (24) and underlies the claim that any nontrivial detour around the unit point pays a surplus in action.

From pairs of agents to many agents (skew decomposition). Fix a cycle C and a partition of bonds by ordered agent-pairs (i, j) . For each (i, j) , define the signed set of logs $\{\alpha_r^{(i \rightarrow j)}\}_r$ from i to j (positive) and $\{\alpha_s^{(j \rightarrow i)}\}_s$ from j to i (recorded as negative). Let

$$\Sigma_{ij} := \sum_r \alpha_r^{(i \rightarrow j)} + \sum_s \alpha_s^{(j \rightarrow i)}$$

be the net pairwise skew (so $\Sigma_{ij} = -\Sigma_{ji}$). The total cycle action is

$$S[C] = \sum_{(i,j)} \sum_r \Phi(\alpha_r^{(i \rightarrow j)}) + \sum_s \Phi(\alpha_s^{(j \rightarrow i)}).$$

Two structural facts hold:

(a) *Flow decomposition.* The multiset $\{\alpha_r^{(i \rightarrow j)}\}$ over all (i, j) admits a decomposition into *closed circulations* (zero net skew on each edge in the decomposition) plus *transfers* that realize each Σ_{ij} exactly once along a simple path. In particular, if all $\Sigma_{ij} = 0$, the nonzero logs lie entirely in closed circulations.

(b) *Smoothing decreases action.* Consider any pair (i, j) and two opposite-signed entries $\alpha > 0$ and $-\beta < 0$. Replacing them by $(\alpha - t)$ and $-(\beta - t)$ with $t \in (0, \min\{\alpha, \beta\}]$ keeps the pairwise

sum fixed, strictly reduces $\Phi(\alpha) + \Phi(-\beta)$ by convexity of Φ (the two-point variant of Jensen), and strictly reduces $|\alpha| + |\beta|$. Iterating this *pairwise smoothing* drives all opposite-signed logs in the pair either to zero or to a single residual entry of the same sign as Σ_{ij} , strictly decreasing the action at each step unless the pair was already balanced.

Combining (a) and (b) yields:

Proposition A.1 (Cycle minimality at $\sigma = 0$). *For any cycle C , $S[C]$ is minimized subject to balance constraints if and only if all pairwise net skews vanish, i.e. $\Sigma_{ij} = 0$ for all (i, j) . Equivalently, the minimizing configuration has only closed circulations with zero log-sum on every circulation (hence all $\alpha \equiv 0$ by the closed-cycle inequality).*

Lyapunov functional for multi-party smoothing. Define the nonnegative functional

$$\mathcal{L} := \sum_{(i,j)} \sum_r \Phi(\alpha_r^{(i \rightarrow j)}) + \sum_s \Phi(\alpha_s^{(j \rightarrow i)}) = S[C].$$

Each pairwise smoothing step strictly lowers \mathcal{L} unless all logs in that pair vanish; smoothing closed circulations with product one also strictly lowers \mathcal{L} unless they are trivial. Since $\mathcal{L} \geq 0$ and strictly decreases along any nontrivial smoothing move, the process terminates exactly at $\alpha \equiv 0$ (the $\sigma = 0$ manifold). This is the multi-party analogue of the two-move inequality (24).

Cut bounds (optional lower estimates). If one insists on lower bounds in terms of the net skew magnitudes alone, convexity yields for each pair (i, j) with fixed total log-sum Σ_{ij} and a fixed number m_{ij} of participating logs:

$$\sum_r \Phi(\alpha_r^{(i \rightarrow j)}) + \sum_s \Phi(\alpha_s^{(j \rightarrow i)}) \geq m_{ij} (\cosh(\Sigma_{ij}/m_{ij}) - 1),$$

with equality when all participating logs are equal (most “evenly spread” case). Summing over pairs provides a coarse but explicit action surplus bound as soon as any $\Sigma_{ij} \neq 0$ (and is sharp given m_{ij}). The deconcentration bound is not used in the main text; the smoothing argument suffices to show strict avoidability of any persistent skew.

In summary, writing $J(x) = \cosh(\ln x) - 1$ turns reciprocity questions into convexity statements about \cosh . Balance constraints become zero-sum conditions on the logs, and strict convexity forces the minimum at $\alpha \equiv 0$. Any nonzero skew is thus an avoidable action surcharge.

B Appendix B: Axiology uniqueness (MI–curvature form up to a φ scale)

We prove that the four axioms in Section 5 force the decomposition

$$V(p_{AE}, x) = \kappa I(A; E) - \mathcal{C}_J^*(p_{AE}, x),$$

with a single positive scale κ fixed once (placed on a φ -tier), and where \mathcal{C}_J^* is the J -induced mechanical penalty under least-action completion. The proof separates *informational coupling* from *mechanical over-strain*.

Objects and axioms. At a cycle C , an agent i with environment E is described by a coarse-grained joint distribution p_{AE} and a ledger microstate $x = \{x_e\}$, projected to the $\sigma = 0$ manifold by least-action completion. A cardinal value functional V maps (p_{AE}, x) to \mathbb{R} and satisfies:

(A1) *Gauge invariance under the bridge:* invariance under admissible re-anchoring $(\tau_0, \ell_0) \mapsto (s\tau_0, s\ell_0)$ (so $c = \ell_0/\tau_0$ fixed), under the calibration that fixes λ_{rec} , under ledger gauge transforms $x_e \mapsto g_{h(e)}x_e g_{t(e)}^{-1}$, and under invertible relabelings of A, E that preserve p_{AE} .

(A2) *Additivity on independent subsystems:* for independent pairs (A_1, E_1) and (A_2, E_2) with disjoint ledgers,

$$V((p_{A_1 E_1}, x^{(1)}) \oplus (p_{A_2 E_2}, x^{(2)})) = V(p_{A_1 E_1}, x^{(1)}) + V(p_{A_2 E_2}, x^{(2)}).$$

(A3) *Concavity (diminishing returns):* for $0 \leq \lambda \leq 1$,

$$V(\lambda p_{AE} + (1 - \lambda)q_{AE}, \lambda x + (1 - \lambda)y) \geq \lambda V(p_{AE}, x) + (1 - \lambda)V(q_{AE}, y).$$

(A4) *Curvature normalization tied to $J''(1) = 1$:* for purely mechanical, gauge-invariant over-strains with p_{AE} fixed,

$$x_e = 1 + \varepsilon_e, \sum_e \varepsilon_e = 0 \Rightarrow V(p_{AE}, x) = V(p_{AE}, \mathbf{1}) - \frac{1}{2} \sum_e \varepsilon_e^2 + o(\|\varepsilon\|^2).$$

Step 1: Splitting V into informational and mechanical parts. Define the *mechanical* deficit at (p_{AE}, x) as the J -Bregman distance to the unit configuration under least-action completion,

$$\mathcal{C}_J^*(p_{AE}, x) := \inf \left\{ \sum_e J(x'_e) : (p_{AE}, x') \text{ is gauge-equivalent to } (p_{AE}, x), \text{ balanced, } \sigma = 0 \right\}.$$

By construction, $\mathcal{C}_J^* \geq 0$, $\mathcal{C}_J^*(p_{AE}, \mathbf{1}) = 0$, it is invariant under the bridge and ledger gauge (A1), additive on independent ledgers (A2), and for small over-strains reduces to $\frac{1}{2} \sum_e \varepsilon_e^2$ (A4). Set

$$U(p_{AE}) := V(p_{AE}, x) + \mathcal{C}_J^*(p_{AE}, x),$$

which is well-defined by gauge invariance and independent of x (purely informational). Then U inherits from V :

(U1) invariance under relabelings of A, E (and under coarse refinements that preserve p_{AE});

(U2) additivity on independent pairs: $U(p_{A_1 E_1} \times p_{A_2 E_2}) = U(p_{A_1 E_1}) + U(p_{A_2 E_2})$;

(U3) concavity in the joint distribution: $U(\lambda p + (1 - \lambda)q) \geq \lambda U(p) + (1 - \lambda)U(q)$.

Step 2: Normalizations on U . Two normalizations fix baselines without adding knobs:

(N0) *Independence baseline:* If A and E are independent ($p_{AE} = p_{A}p_{E}$), then there is no coupling contribution; set $U(p_{A}p_{E}) = 0$ (this is a choice of origin; any constant would drop out of differences but we fix zero).

(N1) *Perfect coupling scale:* For a noiseless K -symbol channel with p_{AE} uniform on the diagonal $\{(a, e) : a = e\}$, declare $U(p_{AE}) = \kappa \log K$. Choosing one K fixes $\kappa > 0$ once (placed on a φ -tier in the main text).

Step 3: Characterization of U (informational term). Consider U on finite alphabets. By (U1) it depends only on the joint law. By (U2) and (N0), U is zero on independent products and additive across independent components. By (U3) U is concave. These three properties, together with the continuity implicit in concavity and with (N1), determine U uniquely as a positive multiple of the *mutual information*:

$$U(p_{AE}) = \kappa I(A; E),$$

where $I(A; E) = H(A) + H(E) - H(A, E)$ is the standard information measure built from the Shannon entropy $H(\cdot)$. In outline:

(i) *Grouping/chain property.* Partition E into (E_1, E_2) with E_2 a refinement of E_1 . Concavity and additivity on independent refinements force a grouping identity for U mirroring the chain rule, $U(A; E_1, E_2) = U(A; E_1) + U(A; E_2 | E_1)$, where the conditional term depends only on the conditional law of E_2 given (A, E_1) . Iterating along product refinements makes U linear in the logarithm of partition cardinalities in the noiseless case (by (N1)).

(ii) *Independence annihilates U .* If A and E become independent under a coarse-graining, U must drop to zero by (N0) and concavity; hence U measures departure from independence.

(iii) *Additivity on products.* For independent pairs, U sums. The only symmetric, continuous, concave, zero-on-independence, additive measure of dependence with the grouping property is a positive multiple of $I(A; E)$.

The proof is standard in spirit: construct U on simple finite cases (binary symmetric channels and noiseless copies), extend by additivity to product channels, and close by continuity and concavity. No alternative functional satisfies all of (U1)–(U3), (N0)–(N1).

Step 4: Characterization of \mathcal{C}_J^* (mechanical term). Let $\tilde{\mathcal{C}}(p_{AE}, x)$ be any nonnegative functional with the properties of \mathcal{C}_J^* : invariance under the bridge and ledger gauge (A1), additivity on independent ledgers (A2), and the small-strain expansion of (A4). Consider the map $x \mapsto \tilde{\mathcal{C}}(p_{AE}, x)$ at fixed p_{AE} . By (A1) it depends only on gauge-invariant local departures from $x = \mathbf{1}$; by (A2) it is separable across bonds; by (A4) its second variation at the unit point is the identity quadratic form. The unique convex even generator with these properties is $\Phi(\alpha) = \cosh \alpha - 1$ (equivalently J in multiplicative coordinates). Minimizing over least-action completions on the $\sigma = 0$ manifold therefore fixes $\tilde{\mathcal{C}} = \mathcal{C}_J^*$.

Step 5: Conclusion and φ -scale. Putting Steps 1–4 together,

$$V(p_{AE}, x) = U(p_{AE}) - \mathcal{C}_J^*(p_{AE}, x) = \kappa I(A; E) - \mathcal{C}_J^*(p_{AE}, x).$$

The sole multiplicative freedom $\kappa > 0$ is fixed once by declaring the value of a canonical noiseless coupling (e.g. one φ -tier for a unit binary link per cycle). No further knobs remain. Any attempt to introduce separate weights for I and \mathcal{C}_J^* would create an inadmissible new scale; any attempt to alter \mathcal{C}_J^* would violate the curvature normalization; any attempt to replace I by another functional would break at least one of (U1)–(U3) or the normalizations (N0)–(N1).

Remark (on concavity). While mutual information is not jointly concave in *all* arguments of an arbitrary parameterization, the concavity required here is with respect to convex mixtures of *coarse-grained* joint laws p_{AE} induced by admissible uncertainty/aggregation at fixed physical constraints; in that domain, the functional constructed above satisfies the axiom and matches the operational role of V in the lexicographic selection.

A Appendix C: σ -Completeness and selector existence

Set-up. Let \mathcal{A}_σ be the nonempty set of σ -*feasible* actions available over a single eight-tick cycle (balanced completions that preserve reciprocity). Endow \mathcal{A}_σ with a metric topology in which the least-action completion map is continuous. For $a \in \mathcal{A}_\sigma$ define

$$H(a) := \max_{i,j} \Delta S(i \rightarrow j \mid a), \quad W(a) := \sum_i f(V_i \mid a), \quad R(a) := \lambda_2(L_\sigma \mid a),$$

where ΔS is the externalized surcharge (Section 4), V the forced axiology (Section 5), f the fixed concave transform (Section 7), and $\lambda_2(L_\sigma)$ the σ -graph spectral gap (Section 9). Let $\tau(a)$ be the fixed RS φ -tier tie-break rank.

Assumptions.

(C1) **Compactness:** \mathcal{A}_σ is compact (or finite).

(C2) **Continuity:** H, W, R are continuous on \mathcal{A}_σ ; τ takes values in a finite, totally ordered set.

Lexicographic selector. Define the nested argmin/argmax sets

$$\mathcal{A}_1 := \arg \min_{\mathcal{A}_\sigma} H, \quad \mathcal{A}_2 := \arg \max_{\mathcal{A}_1} W, \quad \mathcal{A}_3 := \arg \max_{\mathcal{A}_2} R, \quad a^* := \arg \max_{\mathcal{A}_3} \tau.$$

Theorem A.1 (σ -Completeness and selector existence). *Under (C1)–(C2) the sets $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ are nonempty and the selector a^* is well-defined. Moreover, for any $a \in \mathcal{A}_\sigma$ either a is infeasible (not in \mathcal{A}_σ), or*

$$(H(a), -W(a)) \succeq_{\text{lex}} (H(a^*), -W(a^*)),$$

with equality only if $R(a) \leq R(a^)$ and, when $R(a) = R(a^*)$, $\tau(a) \leq \tau(a^*)$.*

Proof. By compactness and continuity (C1)–(C2), H attains a minimum on \mathcal{A}_σ (Weierstrass), so $\mathcal{A}_1 \neq \emptyset$. The set \mathcal{A}_1 is closed; W attains a maximum on it, giving $\mathcal{A}_2 \neq \emptyset$. The same argument yields $\mathcal{A}_3 \neq \emptyset$. Since τ takes finitely many ordered values, $\arg \max_{\mathcal{A}_3} \tau$ is nonempty; fix any a^* in it.

For the optimality claim, take $a \in \mathcal{A}_\sigma$. If $H(a) > H(a^*)$ then $(H(a), -W(a)) \succ_{\text{lex}} (H(a^*), -W(a^*))$. If $H(a) = H(a^*)$ but $W(a) < W(a^*)$ the same strict lexicographic inequality holds. If H and W tie, then $a \in \mathcal{A}_2$ and by construction $R(a) \leq R(a^*)$, with equality only if $\tau(a) \leq \tau(a^*)$. This establishes the stated dominance. If $a \notin \mathcal{A}_\sigma$, it violates feasibility and is excluded a priori. \square

Remarks on noncompact domains. If \mathcal{A}_σ is not compact, existence still holds under mild coercivity: it suffices that the sublevel sets $\{a : H(a) \leq \eta\}$ are compact (or precompact) for the attained minimum of H , and that W and R are upper semicontinuous on those sublevel sets; the above proof then applies verbatim on \mathcal{A}_1 and \mathcal{A}_2 .

Stability under uncertainty. If H, W are known within uncertainty sets $\mathcal{U}_{\Delta S}, \mathcal{U}_V$ (Section 10), continuity implies the existence of *robust* neighborhoods preserving the lexicographic order. In particular, strict margins

$$H^{\sup}(a^*) < \min_{b \neq a^*} H^{\inf}(b), \quad W^{\inf}(a^*) > \max_{b \in \mathcal{A}_1 \setminus \{a^*\}} W^{\sup}(b)$$

guarantee a^* remains selected for all realizations in $\mathcal{U}_{\Delta S}, \mathcal{U}_V$.

B Appendix D: No-Arbitrary-Discount theorem (full derivation)

Objects. Let \mathcal{H} be the set of finite sequences $\mathbf{h} = (h_0, \dots, h_T)$ with $h_t \geq 0$ representing per-cycle worst surcharges $H(a_t)$ over an eight-tick cadence. A temporal aggregator is a map $\mathfrak{A}: \mathcal{H} \rightarrow \mathbb{R}$ used to rank repair paths.

Axioms.

- (D1) **Gauge invariance (bridge):** Joint re-anchoring $(\tau_0, \ell_0) \mapsto (s\tau_0, s\ell_0)$ that preserves $c = \ell_0/\tau_0$ does not change rankings induced by \mathfrak{A} .
- (D2) **Cadence invariance:** For any cyclic rotation R by a multiple of the eight-tick period, $\mathfrak{A}(R\mathbf{h}) = \mathfrak{A}(\mathbf{h})$. For concatenation $\mathbf{h} \parallel \mathbf{k}$, $\mathfrak{A}(\mathbf{h} \parallel \mathbf{k}) = \mathfrak{A}(\mathbf{h}) + \mathfrak{A}(\mathbf{k})$ when the underlying ledgers are independent.
- (D3) **Separability & monotonicity:** \mathfrak{A} is strictly increasing in each argument and continuous on \mathcal{H} .

Theorem B.1 (No-Arbitrary-Discount). *Under (D1)–(D3) there exists $\alpha > 0$ such that*

$$\mathfrak{A}(\mathbf{h}) = \alpha \sum_{t=0}^T h_t \quad \text{for all } \mathbf{h} \in \mathcal{H}.$$

In particular, any nonconstant weight sequence (w_t) (exponential, hyperbolic, phase-dependent) violates at least one axiom.

Proof. Step 1 (symmetry across positions). By cadence invariance (D2), rotation by any multiple of eight ticks leaves \mathfrak{A} unchanged. Hence for the unit vectors $e^{(t)}$ (length one at position t , zeros elsewhere) we must have $\mathfrak{A}(e^{(t)}) = \mathfrak{A}(e^{(s)})$ whenever $t \equiv s \pmod{8}$. But rotation by blocks of eight shows all positions are equivalent, so there is a common $\alpha := \mathfrak{A}(e^{(t)}) > 0$ by monotonicity.

Step 2 (additivity on sums). For $n \in \mathbb{N}$, let $\mathbf{1}_n = (1, \dots, 1)$ with n entries. By concatenation invariance and Step 1,

$$\mathfrak{A}(\mathbf{1}_n) = \sum_{t=1}^n \mathfrak{A}(e^{(t)}) = n \alpha.$$

For any $\lambda \geq 0$, continuity and monotonicity give $\mathfrak{A}(\lambda \mathbf{1}_n) = n \mathfrak{A}(\lambda \mathbf{1}_1)$; define $\phi(\lambda) := \mathfrak{A}(\lambda \mathbf{1}_1)$ so $\mathfrak{A}(\lambda \mathbf{1}_n) = n \phi(\lambda)$ with ϕ continuous, increasing, and $\phi(1) = \alpha$.

Step 3 (Cauchy on the positive cone). For scalars $\lambda, \mu \geq 0$ and independent horizons, (D2) yields

$$\phi(\lambda + \mu) = \mathfrak{A}((\lambda + \mu) \mathbf{1}_1) = \mathfrak{A}(\lambda \mathbf{1}_1 \parallel \mu \mathbf{1}_1) = \phi(\lambda) + \phi(\mu).$$

Continuity implies $\phi(\lambda) = \alpha \lambda$ on $\mathbb{R}_{\geq 0}$. Hence $\mathfrak{A}(\lambda \mathbf{1}_n) = \alpha n \lambda$.

Step 4 (general sequences). Any $\mathbf{h} \in \mathcal{H}$ decomposes as $\sum_{t=0}^T h_t e^{(t)}$. Using rotation invariance (Step 1) and independence additivity (D2),

$$\mathfrak{A}(\mathbf{h}) = \sum_{t=0}^T \mathfrak{A}(h_t e^{(t)}) = \sum_{t=0}^T \alpha h_t = \alpha \sum_t h_t.$$

Exclusion of nonconstant weights. A weighted sum $\sum w_t h_t$ violates (D2) unless w_t is constant across the eight-tick phase and across concatenations; any exponential or hyperbolic form introduces a rate parameter with units of inverse time, which rescales under $(\tau_0, \ell_0) \mapsto (s\tau_0, s\ell_0)$, violating (D1). \square

Consequence. Repair objectives are uniquely determined (up to a positive factor) by the undiscounted eight-tick sum $\sum_t H(a_t)$; there is no lawful discount knob to justify delayed clearance of skew.

C Appendix E: σ -graph robustness lemma (full proof)

Set-up. Let $G_\sigma = (V, E, W)$ be the σ -graph on N domains with weights $w_{ij} = w_{ji} \geq 0$ summarizing reciprocity conductances at $\sigma = 0$. Assume G_σ is connected. The Laplacian is $L_\sigma = D - W$ with $D_{ii} = \sum_{j \neq i} w_{ij}$. Let $0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_N$ be the eigenvalues of L_σ and let L_σ^+ be its Moore–Penrose pseudoinverse on the mean-zero subspace.

Shock model and response. A bounded, mean-zero disturbance is $s \in \mathbb{R}^N$ with $\sum_i s_i = 0$ and $\|s\|_2 \leq S$. The least-action, balanced response solves

$$L_\sigma u = s, \quad u \perp \mathbf{1},$$

and induces edge differences $u_i - u_j$. The Dirichlet energy of the response is

$$\mathcal{E}(u) = \frac{1}{2} \sum_{i < j} w_{ij} (u_i - u_j)^2 = \frac{1}{2} s^\top L_\sigma^+ s.$$

Node-wise surcharge bound. Let ΔS_i denote the externalized surcharge borne by node i in the cycle (sum of convex per-bond penalties on incident edges in the second-order/quadratic regime). By nonnegativity and separability of bond costs,

$$0 \leq \Delta S_i \leq \sum_{j: (i,j) \in E} \frac{1}{2} w_{ij} (u_i - u_j)^2 \leq \mathcal{E}(u),$$

hence

$$\max_i \Delta S_i \leq \mathcal{E}(u) = \frac{1}{2} s^\top L_\sigma^+ s. \quad (25)$$

Spectral gap control. On the mean-zero subspace, $L_\sigma^+ \preceq \lambda_2^{-1} I$, so $s^\top L_\sigma^+ s \leq \lambda_2^{-1} \|s\|_2^2$. Combining with (25) gives

$$\max_i \Delta S_i \leq \frac{1}{2\lambda_2(L_\sigma)} \|s\|_2^2 \leq \frac{S^2}{2\lambda_2(L_\sigma)}. \quad (26)$$

Lemma C.1 (Robust-Preference, full form). *Let a, b be two σ -neutral arrangements tied on present-cycle harm and welfare ($H(a) = H(b)$ and $W(a) = W(b)$). If $\lambda_2(L_\sigma \mid a) > \lambda_2(L_\sigma \mid b)$, then for every mean-zero shock with $\|s\|_2 \leq S$ the least-action responses satisfy*

$$\max_i \Delta S_i^{(a)} \leq \frac{S^2}{2\lambda_2(L_\sigma \mid a)} < \frac{S^2}{2\lambda_2(L_\sigma \mid b)} \geq \max_i \Delta S_i^{(b)}.$$

In particular, arrangement a minimizes the worst-case future surcharge under bounded shocks relative to b .

Proof. Apply (26) to each arrangement with their respective spectral gaps; the strict inequality follows from the strict ordering of gaps. \square

Tightness and assumptions. If s aligns with the eigenvector of λ_2 , the bound (26) is attained in the quadratic approximation: $s^\top L_\sigma^+ s = \lambda_2^{-1} \|s\|_2^2$. Connectivity is required to ensure $\lambda_2 > 0$. If G_σ has multiple components, the analysis applies componentwise; the relevant gap is the minimum λ_2 across components after contracting trivial ones.

Interpretation. The spectral gap is the unique, parameter-free measure of how a σ -neutral polity redistributes bounded shocks under least action. Larger gap means lower amplification of disturbances and smaller worst externalized bills, which is why the lexicographic rule prefers it when feasibility, harm, and welfare tie.

D Appendix F: Measurement protocol details (estimators, intervals, robustness)

This appendix gives practical estimators for the two core quantities used in audits—harm ΔS and value V —together with interval arithmetic and robustness margins. Everything here is gauge-invariant and eight-tick compliant.

F.1 Estimating ΔS (externalized action surcharge)

Definition recap. For an act by agent i in a cycle C , with prescribed multipliers $\alpha = \{\alpha_e\}_{e \in A_i}$ on a finite set of bonds A_i in i 's domain (neutral act $\alpha \equiv 1$), the required action of j is the minimum of $\sum_{e \in \mathcal{E}_j} J(x_e)$ over all balanced, $\sigma = 0$ completions consistent with α . The harm is

$$\Delta S(i \rightarrow j \mid C) := S_j^*[\alpha; C] - S_j^*[\mathbf{1}; C], \quad J(x) = \frac{1}{2}(x + x^{-1}) - 1, \quad x > 0.$$

Convex program (exact regime). Work in log-variables $x_e = \exp(\eta_e)$ so $J(x_e) = \cosh(\eta_e) - 1$, convex and even. Let B be the node-edge incidence matrix (each row a node, orientation $+1$ into the node, -1 out of the node). Balance is $B\eta = 0$. Enforce $\sigma = 0$ by restricting to *least-circulation* completions: represent η as a gradient $\eta = \nabla u$ of node potentials u on each connected component (this kills all cycle sums), and impose the act as linear equality constraints on edges in A_i :

$$(\nabla u)_e = \eta_e^{\text{act}} \quad \text{for } e \in A_i.$$

Then $S_j^*[\alpha; C]$ is obtained by minimizing $\sum_{e \in \mathcal{E}_j} (\cosh((\nabla u)_e) - 1)$ over the node potentials u subject to those equalities. Repeat with $\eta^{\text{act}} \equiv 0$ to get $S_j^*[\mathbf{1}; C]$ and take the difference. This program is convex in u (sum of convex terms of affine forms) and returns a certificate of feasibility; infeasibility indicates that the prescribed act cannot be completed on $\sigma = 0$ in one cycle (repair needed).

Quadratic program (small-strain regime). When all bond strains are small ($|\eta_e| \ll 1$), approximate $J(x_e) = \frac{1}{2}\eta_e^2 + O(\eta_e^4)$. The exact program becomes

$$\min_u \frac{1}{2} \sum_{e \in \mathcal{E}_j} ((\nabla u)_e)^2 \quad \text{s.t.} \quad (\nabla u)_e = \eta_e^{\text{act}} \quad (e \in A_i).$$

The solution is the least-squares gradient field matching the act on A_i . Efficient solvers reduce this to a sparse linear system for u ; $\Delta S(i \rightarrow j)$ is the difference of two quadratic forms.

Pairwise σ constraint (aggregate form). If one prefers not to parameterize by u , enforce $\sigma = 0$ per agent pair by linear constraints

$$\sum_{e \in i \rightarrow j} \eta_e + \sum_{e \in j \rightarrow i} \eta_e = 0 \quad \text{for all ordered pairs } (i, j),$$

alongside balance $B\eta = 0$ and the act constraints on A_i . This yields a convex program directly in η .

Practical notes. (i) *Gauge check*: add any node potential g to u ; the objective and constraints are invariant. (ii) *Boundary assignment*: bonds that sit on agent boundaries are assigned by a fixed, local, gauge-invariant tie-breaker; the *difference* $S_j^*[\alpha] - S_j^*[1]$ does not depend on the tie-breaker. (iii) *Certificates*: record KKT residuals and primal feasibility; these enter uncertainty bounds below.

F.2 Estimating V (forced axiology)

Definition recap. Per agent i , at a cycle C , the axiology is

$$V_i = \kappa I_i(A; E) - \mathcal{C}_{J,i}^*,$$

where $I_i(A; E)$ is the agent–environment mutual information under a fixed coarse-graining, $\kappa > 0$ is fixed once (placed on a φ -tier), and $\mathcal{C}_{J,i}^*$ is the J -induced mechanical penalty attributed to i ’s domain under least-action completion.

Estimating the MI term. Fix finite alphabets for A and E by a preregistered coarse-graining. Collect counts $n(a, e)$ over the cycle (or a short, stationary window of cycles). Form smoothed frequencies

$$\hat{p}(a, e) = \frac{n(a, e) + \beta}{N + \beta K}, \quad K = |A| \cdot |E|,$$

with a fixed β (e.g. Jeffreys’ $\beta = \frac{1}{2}$); compute

$$\hat{I}_i = \sum_{a,e} \hat{p}(a, e) \log \frac{\hat{p}(a, e)}{\hat{p}(a) \hat{p}(e)}.$$

Report $\kappa \hat{I}_i$ as the coupling estimate. For small alphabets, exact confidence bands can be obtained by enumerating probability tables consistent with a total-variation ball around \hat{p} ; for larger alphabets, bootstrap the counts and take percentile bands. Both approaches define the uncertainty set \mathcal{U}_V used below.

Estimating the mechanical penalty. Solve the least-action completion once per cycle (the convex program in u or η). Attribute the per-bond cost $J(x_e)$ to agent domains by the same fixed boundary rule used for ΔS . Sum to get

$$\hat{\mathcal{C}}_{J,i}^* = \sum_{e \in \mathcal{E}_i} (\cosh(\hat{\eta}_e) - 1),$$

or, in the small-strain regime, $\hat{\mathcal{C}}_{J,i}^* = \frac{1}{2} \sum_{e \in \mathcal{E}_i} \hat{\eta}_e^2$.

Value estimate and welfare. Combine the two pieces

$$\widehat{V}_i = \kappa \widehat{I}_i - \widehat{\mathcal{C}}_{J,i}^*, \quad \widehat{W} = \sum_i f(\widehat{V}_i),$$

where f is the fixed concave transform (normalized by $f(0) = 0$, $f'(0) = 1$, $f''(0) = -1$).

F.3 Interval arithmetic (uncertainty sets)

Uncertainty on ΔS . Instrument noise and solver tolerances give intervals for η_e and residuals for the equality constraints. Define a box uncertainty set $\mathcal{U}_\eta = \prod_e [\underline{\eta}_e, \bar{\eta}_e]$ consistent with diagnostics (including KKT residual bounds). Compute outer bounds for harm by solving the two robust programs

$$\Delta S^{\text{inf}}(i \rightarrow j) = \inf_{\eta \in \mathcal{U}_\eta} \left(S_j^*[\alpha] - S_j^*[\mathbf{1}] \right), \quad \Delta S^{\text{sup}}(i \rightarrow j) = \sup_{\eta \in \mathcal{U}_\eta} \left(S_j^*[\alpha] - S_j^*[\mathbf{1}] \right).$$

In practice: (i) *Quadratic regime*—closed forms give

$$\Delta S^{\text{sup}/\text{inf}} \approx \frac{1}{2} \|P_j A \bar{\eta}^{\text{act}}\|_2^2 \pm L_{\Delta S} \|\delta\eta\|_2,$$

where A maps edge logs to the least-squares gradient field, P_j projects to \mathcal{E}_j , $\bar{\eta}^{\text{act}}$ is the midpoint act, $\delta\eta$ is the uncertainty radius, and $L_{\Delta S}$ is a computable local Lipschitz bound. (ii) *Exact regime*—solve the min/max by alternating convex optimization and box-projection (guarantees outer bounds).

Uncertainty on V . For the MI term, take $\widehat{I}_i^{\text{inf}/\text{sup}}$ as the lower/upper ends of a bootstrap band or the extrema over a total-variation ball $\{q : \|q - \hat{p}\|_1 \leq r\}$ (grid search suffices for small alphabets). For the mechanical term, propagate \mathcal{U}_η through J :

$$\mathcal{C}_{J,i}^{\text{inf}} = \sum_{e \in \mathcal{E}_i} \min_{\eta_e \in [\underline{\eta}_e, \bar{\eta}_e]} (\cosh \eta_e - 1), \quad \mathcal{C}_{J,i}^{\text{sup}} = \sum_{e \in \mathcal{E}_i} \max_{\eta_e \in [\underline{\eta}_e, \bar{\eta}_e]} (\cosh \eta_e - 1),$$

which are attained at interval endpoints because \cosh is convex. Then

$$V_i^{\text{inf}} = \kappa \widehat{I}_i^{\text{inf}} - \mathcal{C}_{J,i}^{\text{sup}}, \quad V_i^{\text{sup}} = \kappa \widehat{I}_i^{\text{sup}} - \mathcal{C}_{J,i}^{\text{inf}},$$

and

$$W^{\text{inf}} = \sum_i f(V_i^{\text{inf}}), \quad W^{\text{sup}} = \sum_i f(V_i^{\text{sup}}).$$

F.4 Robustness margins (decision guarantees)

Per-cycle selection. Given a finite candidate set \mathcal{A}_σ of σ -feasible actions, compute

$$H^{\text{inf}/\text{sup}}(a) = \inf / \sup_{\theta \in \mathcal{U}_{\Delta S}} \max_{i,j} \Delta S_\theta(i \rightarrow j \mid a), \quad W^{\text{inf}/\text{sup}}(a) = \inf / \sup_{\phi \in \mathcal{U}_V} \sum_i f(V_{i,\phi} \mid a).$$

Declare a^* *robustly selected* if

$$H^{\text{sup}}(a^*) < \min_{b \neq a^*} H^{\text{inf}}(b),$$

or, when the first margins tie,

$$H^{\text{sup}}(a^*) = \min_b H^{\text{inf}}(b) \quad \text{and} \quad W^{\text{inf}}(a^*) > \max_{b: H^{\text{inf}}(b) = H^{\text{inf}}(a^*)} W^{\text{sup}}(b).$$

Apply the spectral-gap and φ -tier refinements only within any remaining ties.

Repair horizons. For $\{a_t\}_{t=0}^T$, compute

$$\mathcal{J}_{\text{rep}}^{\text{inf/sup}} = \sum_{t=0}^T H^{\text{inf/sup}}(a_t),$$

and select the plan with the smallest $\mathcal{J}_{\text{rep}}^{\text{sup}}$; if multiple plans tie, use the per-cycle welfare and robustness refinements.

Slack and determinacy. Introduce a small, fixed technical slack $\epsilon_H, \epsilon_W > 0$ (on the order of solver tolerances). If margins fall below slack, return *indeterminate* and request more measurement (larger samples for MI, tighter instrument bounds for logs).

F.5 Sanity checks and failure modes

Sanity checks. (i) *Gauge test*: random re-anchoring $(\tau_0, \ell_0) \mapsto (s\tau_0, s\ell_0)$ and node-potential shifts leave ΔS and V unchanged. (ii) *Zero-act test*: $\alpha \equiv 0$ yields $\Delta S \equiv 0$ and leaves V unchanged within noise. (iii) *Symmetry test*: swapping agent labels that preserve p_{AE} leaves V invariant.

Failure modes. (i) *Infeasible act*: convex program infeasible \Rightarrow repair required; do not impute ΔS . (ii) *Boundary sensitivity*: different boundary rules change absolute S_j^* , but differences defining ΔS and V must agree within slack; if not, report *indeterminate*. (iii) *Non-stationary MI*: if coarse-grained statistics drift within the cycle window, widen the window to the eight-tick cadence or report separate estimates per tick and aggregate without discount.