Virtues as Generators: A Zero-Parameter, Auditable Ethics from Recognition Science

Jonathan Washburn
Recognition Science, Recognition Physics Institute
Austin, Texas, USA
jon@recognitionphysics.org

March 2, 2026

Abstract

We present a companion to "Morality as a Conservation Law", extending the Recognition Science (RS) framework from feasibility to operation: fourteen *virtues* are formalized as the complete set of admissible generators of ethical dynamics on the reciprocity-conserving manifold $\sigma = 0$. Each generator is parameter-free and inherits the same invariants that calibrate physics: the unique convex cost $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ with $J''(1) = 1$, the eight-tick cadence, and the φ-fixed scaling identities. The operational calculus couples four auditable quantities: (i) reciprocity skew σ (feasibility), (ii) harm ΔS as the externalized action surcharge under least-action completion, (iii) a forced axiology $V = \kappa I(A; E) - \mathcal{C}_J^*$ (mutual-information coupling minus J-induced curvature penalty; κ fixed on a φ-tier), and (iv) consent as the directional derivative sign $D_j V_i \geq 0$ along feasible directions. A lexicographic decision rule selects among σ-feasible actions without weights: minimize $\max \Delta S$, then maximize $\sum f(V)$ (curvature-normalized), then maximize the σ-graph spectral gap λ_2, breaking residual ties on the fixed φ-tier. We outline per-virtue invariants and audits, state a completeness/minimality program (DREAM), and list crisp falsifiers. If RS is the true physics, these generators constitute an *inevitable* basis for moral action—auditable, falsifiable, and free of tunable knobs.

Keywords: Recognition Science; virtues as generators; reciprocity conservation ($\sigma = 0$); externalized action surcharge (ΔS); forced axiology (V); lexicographic selection; parameter-free ethics; auditability; DREAM theorem.

Contents

Notation (quick reference)

$J(x)$	Unique convex cost, $\frac{1}{2}(x+x^{-1})-1$, with $J(1)=0$, $J''(1)=1$
φ	Golden ratio, $(1+\sqrt{5})/2$; unique positive root of $\varphi^2=\varphi+1$
$\sigma_{ij}[C]$	Pairwise reciprocity skew (signed log-imbalance) over cycle C
$\sigma[C]$	Total skew, $\sum_{i<j} \lvert\sigma_{ij}[C]\rvert$ (feasibility requires $\sigma[C]=0$)
κ	Synonymous with σ (reciprocity skew, earlier "curvature")
Π_{LA}	Least-action completion operator (rebalances to $\sigma=0$)
$\Delta S(i \to j)$	Externalized action surcharge (harm) from i to j under Π_{LA}
V	Forced axiology: $\kappa I(A;E)-\mathcal{C}_J^*$ (MI–curvature), κ fixed by φ-tier
$I(A;E)$	Agent–environment mutual information
$\mathcal{C}_J^*$	J-induced curvature penalty under least-action completion
f	Curvature-normalized concave transform for welfare ($f(0)=0$, $f'(0)=1$, $f''(0)=-1$). Default: $f($
G_σ	σ-graph with weights w_{ij} (second variations at $\sigma=0$)
L_σ	Laplacian of G_σ; spectral gap $\lambda_2(L_\sigma)$
$H(a)$	$\max_{i,j} \Delta S(i \to j \mid a)$ (worst surcharge in a cycle)
$W(a)$	$\sum_i f(V_i \mid a)$ (cardinal welfare)
$R(a)$	$\lambda_2(L_\sigma \mid a)$ (robustness via spectral gap)
$\mathcal{F}$	Feasible manifold at $\sigma=0$ (balanced, least-action completable)
$\mathcal{V}$	Virtue (admissible generator $\mathcal{V}:\mathcal{F}\to\mathcal{F}$)

1 Positioning and contributions

Motivation. This paper moves from the conservation statement in the companion article to operation: from proving that admissible worldlines live on $\sigma=0$ to providing the *generators* that act on that manifold and the *audit* that selects among σ-feasible futures without introducing parameters.

Contributions.

- Define and formalize **fourteen RS-forced virtues** as generators on $\sigma=0$.
- Show per-virtue **invariants**, **ΔS bounds**, **V monotonicity**, and **φ-optimal ratios** where applicable.
- Present the **lexicographic selector** and **audit protocol** covering σ, ΔS, V, λ_2, and consent.
- Give a **completeness/minimality roadmap** (DREAM) and **falsifiers** that would defeat the framework.
- Release **reproducible Lean anchors** and **audit artifacts** for independent verification.

Relationship to the companion. Relative to *Morality as a Conservation Law*, this paper *operationalizes selection* on the $\sigma=0$ manifold: it specifies the generator set (virtues), the auditable quantities ($\sigma,\Delta S,V,\lambda_2$, consent), and the parameter-free rule that chooses among admissible actions.

2 RS scaffold (recap)

This section summarizes the Recognition Science (RS) substrate on which the virtues calculus operates. Full derivations and proofs are given in the companion paper *Morality as a Conservation Law.*

Unique convex cost and φ fixed point. RS fixes a single convex, symmetric cost on $\mathbb{R}_{>0}$,

$$J(x) \;=\; \tfrac{1}{2}\big(x + x^{-1}\big) \;-\; 1, \qquad x>0, \qquad J(1)=0, \quad J(x)=J(x^{-1}), \quad J''(1)=1. \tag{2.1}$$

Self-similarity at this curvature normalization forces the golden ratio φ as the unique positive fixed point of the induced scale recursion, $\varphi^2 = \varphi + 1$.

Eight-tick cadence and bridge identities. Time aggregates in minimal, cadence-invariant blocks of eight ticks (for $D = 3$), realized by a ledger-compatible Hamiltonian walk on the 3-cube. Calibration to physical units is knob-free through dimensionless identities, e.g.

$$c = \frac{\ell_0}{\tau_0}, \qquad \frac{c^3\,\lambda_{\text{rec}}^2}{\hbar G} = \frac{1}{\pi} \quad\Longleftrightarrow\quad \lambda_{\text{rec}} = \sqrt{\frac{\hbar G}{\pi c^3}}. \tag{2.2}$$

Reciprocity skew and cycle minimality at $\boldsymbol{\sigma = 0}$. Let σ denote cycle-wise reciprocity skew (signed log-imbalance of bestowed vs received adjustments between parties). Strict convexity of J at $x = 1$ yields, for any $\varepsilon \neq 0$,

$$J(1+\varepsilon) + J(1-\varepsilon) \;>\; 2J(1) \;=\; 0, \tag{2.3}$$

so replacing any bidirectional pair $(1{+}\varepsilon,\, 1{-}\varepsilon)$ by $(1, 1)$ *strictly lowers* the per-cycle action. Iterating this *pairwise smoothing* over all imbalanced pairs yields:

> *A balanced cycle ($\sigma = 0$) is action-minimizing; any persistent $\sigma \neq 0$ carries a strictly positive, avoidable action surplus.*

Operational objects. On the $\sigma = 0$ manifold, selection among feasible futures uses four audit quantities:

- **Harm $\boldsymbol{\Delta S}$** (externalized action surcharge): for an act by i, the marginal increase in j's required portion of the per-cycle action under the *least-action completion* that preserves balance and $\sigma = 0$.
- **Value $\boldsymbol{V}$** (forced axiology): uniquely fixed (up to a φ-tier scale κ) by gauge invariance, additivity, concavity, and curvature normalization,

  $$V \;=\; \kappa\, I(A;E) \;-\; \mathcal{C}_J^{*}, \tag{2.4}$$

  where $I(A;E)$ is agent–environment mutual information and $\mathcal{C}_J^{*}$ is the J-induced mechanical penalty under least-action completion.
- **Consent** as a *directional derivative* on the feasible manifold: for j's contemplated feasible direction ξ, agent i consents iff

  $$D_j V_i[\xi] \;=\; \tfrac{d}{dt}\, V_i\big(\text{proj}_{\sigma=0}(t\xi)\big)\big|_{t=0} \;\geq\; 0. \tag{2.5}$$

- **Robustness** via the σ-graph spectral gap λ_2: the algebraic connectivity of the reciprocity network's Laplacian at $\sigma = 0$, controlling worst-case surcharge amplification under bounded shocks.

These invariants and objects ground the virtues and the lexicographic selector in a parameter-free, auditable calculus shared with the companion derivations.

3 MoralState, reciprocity, and feasibility

MoralState. An agent-level ethical state is represented as

$$\text{MoralState } s \; = \; (L, \, \{\mathcal{E}_i\}_{i \in \mathcal{A}}, \, E; \, \text{valid}),$$

where L is a ledger snapshot (positive multipliers $x_e > 0$ on directed bonds with site-by-site balance), $\{\mathcal{E}_i\}$ partitions bonds by agent/group, $E > 0$ is an energy/audit budget, and `valid` asserts global admissibility (below). All updates are evaluated over one eight-tick cycle via the least-action completion operator Π_{LA} that rebalances non-acting bonds while keeping $\sigma = 0$.

Reciprocity and curvature ($\boldsymbol{\kappa \equiv \sigma}$). For any ordered pair (i, j) and cycle C, the pairwise skew is the signed log-imbalance

$$\sigma_{ij}[C] \; := \sum_{e \in i \to j} \ln x_e \;\; - \sum_{e \in j \to i} \ln x_e, \qquad \sigma_{ij} = -\sigma_{ji}.$$

We write κ synonymously for reciprocity skew (earlier "curvature"): $\kappa \equiv \sigma$. Global skew is $\sigma[C] := \sum_{i<j} |\sigma_{ij}[C]|$. Per-agent projections aggregate a domain's imbalances, e.g. $\kappa_i := \sum_{j \neq i} |\sigma_{ij}|$ or, when needed, its signed components $\{\sigma_{ij}\}_j$.

Global admissibility and per-agent projections. A MoralState s is *globally admissible* iff its pre- and post-cycle ledgers satisfy

$$\sigma[C_{\text{pre}}] = 0, \qquad \sigma[C_{\text{post}}] = 0, \qquad E > 0, \tag{3.1}$$

with the post state obtained by applying the contemplated act on finitely many bonds and completing the remainder via Π_{LA}. Per-agent projections $\Pi_i(L)$ restrict audit quantities (required-action portions, local skews) to $\mathcal{E}_i$ and feed $\Delta S(i \to j)$ and V_i.

Feasible transformations and virtues. Let $\mathcal{F}$ be the *feasible manifold* at a $\sigma = 0$ state: the set of nearby ledger states reachable in one eight-tick cycle that preserve site balance and $\sigma = 0$ under Π_{LA}. A *feasible transformation* is a map $T : \mathcal{F} \to \mathcal{F}$ realized by acting on finitely many bonds and rebalancing the rest by least action.

Definition 3.1 (Virtue, admissible generator on $\sigma = 0$)**.** A *virtue* is an admissible generator $\mathcal{V} : \mathcal{F} \to \mathcal{F}$ that maps σ-feasible states to σ-feasible states in one eight-tick cycle and exposes the audit tuple $(\sigma, \Delta S, V, \lambda_2, \text{consent})$ computed with Π_{LA}.

In particular: (i) *feasibility* holds cycle-wise, (ii) ΔS is the externalized surcharge under least-action completion, (iii) $V = \kappa I(A; E) - \mathcal{C}_J^*$ is computed per agent with the fixed φ-tier, (iv) *consent* uses the directional derivative on $\mathcal{F}$, $D_j V_i[\xi] \geq 0$, and (v) *robustness* is summarized by the post-action σ-graph spectral gap λ_2. The next sections instantiate fourteen such generators and their invariants.

4 Virtues calculus: definitions and invariants

Formal interface. A virtue exposes two maps: a *transform* on the feasible manifold and an *audit hook* for selection,

$$\mathcal{V}: \mathcal{F} \to \mathcal{F}, \qquad \mathsf{Audit}_{\mathcal{V}}: \mathcal{F} \to \mathsf{Report}, \tag{4.1}$$

with $\mathcal{F}$ the $\sigma = 0$ manifold (Sec. **??**). The report is a structured tuple of invariants and scores used by the lexicographic selector.

Invariants enforced. Every virtue must satisfy the following *hard* invariants per eight-tick cycle:

- **Feasibility (reciprocity conservation):**

$$\sigma[C_{\text{pre}}] = 0, \qquad \sigma[C_{\text{post}}] = 0, \qquad L_{\text{post}} = \Pi_{\text{LA}}(\text{act} \mid L_{\text{pre}}). \tag{4.2}$$

  Updates are applied on finitely many bonds; the least-action completion Π_{LA} rebalances the rest.

- **Energy constraints:** positivity of the audit budget and a per-cycle sustainability cap,

$$E_{\text{post}} > 0, \qquad \Delta E \;:=\; E_{\text{pre}} - E_{\text{post}} \;\le\; \frac{1}{\varphi}\, E_{\text{pre}}, \tag{4.3}$$

  unless explicitly tightened/neutralized by the specific virtue (Temperance saturates this bound).

- **φ-tier normalization:** all numerical scales inherited from the forced axiology and cadence: the information scale κ is fixed once (unit noiseless binary link $\mapsto$ one φ-tier), mechanical curvature is set by $J''(1) = 1$; no new free weights are permitted.

Audit hooks and metrics. For a pre/post cycle $(C_{\text{pre}}, C_{\text{post}})$ the audit hook returns

$$\mathsf{Report} \;=\; \big(\Sigma^{\text{pre}}, \Sigma^{\text{post}}, \Delta S, H, \Delta V, W, \lambda_2^{\text{pre}}, \lambda_2^{\text{post}}, \Delta\lambda_2, \mathsf{Consent}\big), \tag{4.4}$$

where:

- $\Sigma^{\text{pre}} = \{\sigma_{ij}[C_{\text{pre}}]\}_{i<j}$ and $\Sigma^{\text{post}} = \{\sigma_{ij}[C_{\text{post}}]\}_{i<j}$ are the σ *traces*.
- $\Delta S(i \to j)$ is the *harm matrix* (externalized action surcharge under Π_{LA}), with

$$H \;:=\; \max_{i,j}\, \Delta S(i \to j). \tag{4.5}$$

- $\Delta V_i := V_i^{\text{post}} - V_i^{\text{pre}}$ and the welfare aggregate $W = \sum_i f(V_i)$ (curvature-normalized concave transform).
- λ_2^{pre} and λ_2^{post} are the σ-graph spectral gaps; $\Delta\lambda_2 = \lambda_2^{\text{post}} - \lambda_2^{\text{pre}}$ summarizes robustness change.
- $\mathsf{Consent}$ is a table of derivative signs $\mathsf{C}(i \leftarrow j) = \big[\, D_j V_i[\xi] \ge 0 \,\big]$ for the contemplated feasible directions ξ.

These metrics are computed *after* applying the transform and projecting with Π_{LA}. The lexicographic selector first filters on feasibility (traces zero), then minimizes H, maximizes W on the tie set, then maximizes λ_2, and finally breaks any residual tie on the fixed φ-tier.

5 Operational pieces (brief)

This section summarizes the three core operators used by all virtues; full statements and proofs are in the companion *Morality as a Conservation Law.*

Harm ΔS (externalized action surcharge). Given an act by agent i on a finite bond set, the *required action* of j is the minimum portion of per-cycle cost borne by j's domain among all balanced, $\sigma = 0$ completions (least-action projection). Harm is the marginal increase

$$\Delta S(i \to j) \;:=\; S_j^*[\text{act}] \;-\; S_j^*[\mathbf{1}]. \tag{5.1}$$

Properties: (P1) *gauge-invariance on the ledger* (unchanged under node-potential transforms and unit re-anchoring preserving $c = \ell_0/\tau_0$); (P2) *additivity* on independent subsystems; (P3) *composition* over the eight-tick cadence, ΔS sums per cycle.

Forced axiology V (MI–curvature). Axioms—gauge invariance, additivity on independent subsystems, concavity under coarse aggregation, and curvature normalization tied to $J''(1) = 1$—force a unique cardinal form (up to a φ-tier scale κ):

$$V \;=\; \kappa\, I(A;E) \;-\; \mathcal{C}_J^*, \tag{5.2}$$

with $I(A;E)$ the agent–environment mutual information and $\mathcal{C}_J^*$ the J-induced mechanical penalty under least-action completion. The normalization κ is fixed once by a canonical noiseless link; no free weights remain.

Consent (derivative sign on the σ-manifold). For agent j's contemplated feasible direction ξ at a $\sigma = 0$ state, agent i *consents* if the directional derivative of i's value is nonnegative under least-action completion,

$$D_j V_i[\xi] \;=\; \tfrac{d}{dt}\, V_i\big(\text{proj}_{\sigma=0}(t\xi)\big)\Big|_{t=0} \;\geq\; 0. \tag{5.3}$$

For non-competent agents, a conservative *proxy consent* uses a lower-envelope axiology consistent with the four axioms; selection may temporarily override individual consent only under the lexicographic rule when no consent-respecting feasible alternative exists and minimax harm compels repair.

Definition 5.1 (Consent on the feasible manifold)**.** Fix a $\sigma = 0$ state and let $\mathcal{T}_s\mathcal{F}$ denote the set of feasible directions at s (infinitesimal acts closed under Π_{LA}). For $\xi \in \mathcal{T}_s\mathcal{F}$ attributed to agent j, agent i *consents* iff

$$D_j V_i[\xi] \;\geq\; 0, \qquad \text{with} \;\; D_j V_i[\xi] := \tfrac{d}{dt}\, V_i\big(\text{proj}_{\sigma=0}(t\xi)\big)\Big|_{t=0}.$$

6 Lexicographic selector (decision principle)

We select among σ-feasible actions using a parameter-free lexicographic order. Let $\mathcal{A}_\sigma$ be the nonempty set of actions that map a $\sigma = 0$ state to a $\sigma = 0$ state over one eight-tick cycle under Π_{LA}. Define

$$H(a) := \max_{i,j} \Delta S(i \to j \mid a), \qquad W(a) := \sum_i f\big(V_i \mid a\big), \qquad R(a) := \lambda_2\big(L_\sigma \mid a\big), \tag{6.1}$$

and let $\tau(a)$ be the fixed φ-tier tie-break rank.

Order of operations.

$$
\begin{aligned}
\mathcal{A}_1 &:= \arg\min_{a\in\mathcal{A}_\sigma} H(a) && \text{(minimax max } \Delta S\text{)}\\
\mathcal{A}_2 &:= \arg\max_{a\in\mathcal{A}_1} W(a) && \text{(maximize welfare)}\\
\mathcal{A}_3 &:= \arg\max_{a\in\mathcal{A}_2} R(a) && \text{(maximize robustness)}\\
a^* &:= \arg\max_{a\in\mathcal{A}_3} \tau(a) && \text{(}\varphi\text{-tier tie).}
\end{aligned}
\tag{6.2}
$$

Selector algorithm (concise). Given audited candidates $\{a_k\}$ with σ-feasible reports, compute $H(a_k), W(a_k), R(a_k)$:

1. Filter: $\mathcal{A}_\sigma \leftarrow \{a_k : \ \sigma\text{-traces post} = 0\}$.
2. Harm stage: $\mathcal{A}_1 \leftarrow \arg\min_{a\in\mathcal{A}_\sigma} H(a)$.
3. Welfare stage: $\mathcal{A}_2 \leftarrow \arg\max_{a\in\mathcal{A}_1} W(a)$ (default $f(v) = v$).
4. Robustness stage: $\mathcal{A}_3 \leftarrow \arg\max_{a\in\mathcal{A}_2} R(a)$.
5. Tie-break: return $a^* \in \arg\max_{a\in\mathcal{A}_3} \tau(a)$.

Why no weights appear. Introducing a scalar weight to blend H and W (or W and R) would create a new, uncalibrated scale that is not fixed by the bridge or by $J''(1) = 1$, violating parameter-free constraints (gauge, cadence, curvature). The lexicographic order respects these invariances: feasibility first, then harm minimax, then fixed-cardinal welfare, then robustness, then the same φ-tier arithmetic used elsewhere.

Proposition 6.1 (Existence and selector optimality)**.** *If $\mathcal{A}_\sigma$ is finite or compact and H, W, R are continuous on $\mathcal{A}_\sigma$, then (**??**) yields nonempty $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ and a well-defined a^*. Moreover, for any $a \in \mathcal{A}_\sigma$ either a is infeasible (not in $\mathcal{A}_\sigma$) or*

$$
\big(H(a),\, -W(a)\big) \ \succeq_{\text{lex}} \ \big(H(a^*),\, -W(a^*)\big), \tag{6.3}
$$

with equality implying $R(a) \le R(a^*)$, and if R also ties then $\tau(a) \le \tau(a^*)$.

Robustness margins. Let $\mathcal{U}_{\Delta S}$ and $\mathcal{U}_V$ be uncertainty sets for ΔS and V. Define

$$
H^{\inf/\sup}(a) := \inf/\sup_{\theta\in\mathcal{U}_{\Delta S}} H_\theta(a), \qquad W^{\inf/\sup}(a) := \inf/\sup_{\phi\in\mathcal{U}_V} W_\phi(a). \tag{6.4}
$$

Declare a^* *robustly certified* if

$$
H^{\sup}(a^*) \ < \ \min_{b\neq a^*} H^{\inf}(b), \tag{6.5}
$$

or, when the first margins tie,

$$
H^{\sup}(a^*) = \min_b H^{\inf}(b) \quad \text{and} \quad W^{\inf}(a^*) \ > \max_{b:\ H^{\inf}(b)=H^{\inf}(a^*)} W^{\sup}(b), \tag{6.6}
$$

with the R and τ refinements applied only within any remaining ties. If these margins fail, the audit returns *indeterminate* and demands more measurement rather than weights.

7 The fourteen virtues

Each virtue is presented with a consistent template: informal role; formal definition (pre/post constraints and parameters); enforced invariants ($\sigma = 0$, energy, φ-tier); core theorems (e.g., σ-preservation, ΔS bounds, V monotonicity, φ-optimality when applicable); Lean anchors; and a worked micro-example with audit outputs.

7.1 Love — φ-ratio equilibration

Lean: `IndisputableMonolith/Ethics/Virtues/Love.lean`

Role. Bilateral equilibration of local skew with φ-optimal energy split.

Definition. Act on two domains (i, j); set pairwise skew to the average (unit in balanced normalization) and split energy E_{tot} as $E'_i = E_{\text{tot}}/\varphi$, $E'_j = E_{\text{tot}}/\varphi^2$; complete others by Π_{LA}.

Invariants. $\sigma_{ij}^{\text{post}} = 0$; $E'_i, E'_j > 0$; φ-tier normalization (no free weights).

Core theorems. `love_conserves_skew`; `love_reduces_variance`; `love_energy_split_is_golden`.

Micro-example. Two domains with $E = (1.0, 0.6)$; post: $E' = (0.618, 0.382)$, σ-traces zero; $H = 0.00$; $\Delta V = (+0.07, +0.05)$; $\Delta\lambda_2 \approx +0.02$.

7.2 Justice — accurate posting within eight-tick

Lean: `IndisputableMonolith/Ethics/Virtues/Justice.lean`

Role. Ensure accurate ledger posting within the cadence.

Definition. A protocol with `posting` such that `accurate`: balances update exactly; `timely`: recording within 8 ticks.

Invariants. $\sigma^{\text{post}} = 0$; E unchanged; cadence bound ≤ 8.

Core theorems. `justice_preserves_sigma_zero`; `justice_compositional`.

Micro-example. Single entry posted; σ-traces unchanged; $H = 0.00$; $\Delta V = 0$; timing certificate ≤ 8.

7.3 Forgiveness — bounded skew transfer with energy cost

Lean: `IndisputableMonolith/Ethics/Virtues/Forgiveness.lean`

Role. Relieve debtor skew with creditor energy payment under caps.

Definition. Transfer bounded $\Delta\kappa \leq \kappa_{\max}$ from debtor to creditor; enforce $E_{\text{cred}}^{\text{post}} > 0$ with explicit cost schedule.

Invariants. Global $\sigma = 0$; $E_{\text{cred}}^{\text{post}} > 0$; parameter-free cap expressed in ledger units.

Core theorems. `forgiveness_conserves_total_skew`; `forgiveness_bounded_harm`; `forgiveness_improves`

Micro-example. Debtor $\kappa = +4$; transfer $\Delta\kappa = 2$; $H = 0.25$; $\Delta V_{\text{debtor}} = +0.12$; $\Delta V_{\text{cred}} = -0.03$.

7.4 Wisdom — φ-discount horizon optimization

Lean: `IndisputableMonolith/Ethics/Virtues/Wisdom.lean`

Role. Choose actions maximizing long-run value with φ-consistent discount (cadence-compliant).

Definition. Evaluate horizon policies with undiscounted eight-tick aggregation and fixed φ-tier normalization of V; select $\arg\max W$ on feasible set.

Invariants. $\sigma = 0$ per cycle; no ad-hoc discount; φ-tier.

Core theorems. `wisdom_respects_cadence`; `wisdom_monotone_V`.

Micro-example. Two-step policy: higher immediate ΔV but worse H is discarded; selected plan: $H = 0.30 \to 0.10$, $\sum f(V)$ increases.

7.5 Courage — high $|\nabla\sigma|$ moves, surcharge caps

Lean: `IndisputableMonolith/Ethics/Virtues/Courage.lean`

Role. Permit decisive moves when skew gradients are high, under capped surcharge.

Definition. Predicate $|\nabla\sigma| > \theta$ (with θ cadence-derived); allow act if projected $H \le H_{\max}$.

Invariants. $\sigma = 0$ post; $H \le H_{\max}$; energy positivity.

Core theorems. `courage_feasible_under_cap`; `courage_reduces_future_H`.

Micro-example. Shock zone: act lowers future max ΔS from 0.8 to 0.4 with cap $H \le 0.5$.

7.6 Temperance — energy spend $\le E/\varphi$

Lean: `IndisputableMonolith/Ethics/Virtues/Temperance.lean`

Role. Ensure sustainable per-cycle energy expenditure.

Definition. Enforce $\Delta E \le E/\varphi$ for any admissible act; reject otherwise.

Invariants. $\sigma = 0$; $E^{\text{post}} > 0$; ΔE bound.

Core theorems. `temperance_ensures_E_pos`; `temperance_is_sharp`.

Micro-example. Candidate act exceeding E/φ rejected; feasible variant: $\Delta E = 0.55/\varphi$; H unchanged; ΔV small positive.

7.7 Prudence — risk-adjusted expected skew

Lean: `IndisputableMonolith/Ethics/Virtues/Prudence.lean`

Role. Favor lower variance trajectories at comparable expectation.

Definition. Minimize $\mathbb{E}[|\kappa|] + \lambda\,\mathrm{Var}(|\kappa|)$ with λ fixed by φ-tier normalization (no free tuning) over σ-feasible policies.

Invariants. $\sigma = 0$; risk penalty scale fixed; energy positive.

Core theorems. `prudence_reduces_volatility`; `prudence_dominates_equal_mean`.

Micro-example. Two actions with equal $\mathbb{E}[|\kappa|]$; prudence selects the one with Var 30% lower; H equal; ΔV slightly higher.

7.8 Compassion — φ^2 energy, φ^4 skew relief

Lean: `IndisputableMonolith/Ethics/Virtues/Compassion.lean`

Role. Asymmetric relief to suffering states with φ-efficient conversion.

Definition. Transfer energy at rate $\le 1/\varphi^2$ and convert to debtor skew relief at factor φ^4 under Π_{LA}.

Invariants. Global $\sigma = 0$; energy bounds; φ-tier factors fixed.

Core theorems. `compassion_reduces_debtor_kappa`; `compassion_bounded_cost`.

Micro-example. Debtor $\kappa = 6$, helper $E = 2.0$; transfer 0.4; relief $\approx 0.4\,\varphi^4$; $H = 0.20$; net $\Delta W > 0$.

7.9 Gratitude — φ-rate reinforcing cooperation

Lean: `IndisputableMonolith/Ethics/Virtues/Gratitude.lean`

Role. Update cooperation propensity at a golden rate.

Definition. $p'_{AB} = p_{AB} + (1 - p_{AB})/\varphi$ after virtuous act; apply Π_{LA} to preserve feasibility.

Invariants. $\sigma = 0$; E neutral; geometric convergence bounds.
Core theorems. `gratitude_geometric_convergence`; `gratitude_monotone`.
Micro-example. Initial $p = 0.2$; sequence $0.2 \to 0.36 \to 0.49 \to \cdots$; $H = 0$; ΔV nonnegative.

7.10 Patience — eight-tick waiting

Lean: `IndisputableMonolith/Ethics/Virtues/Patience.lean`
Role. Avoid suboptimal, alias-prone moves; wait at least one cadence.
Definition. Enforce $t - t_{\text{last}} \geq 8$ ticks before non-repair actions; repairs follow Sec. **??** objective.
Invariants. $\sigma = 0$; cadence compliance; energy neutral.
Core theorems. `patience_avoids_alias`; `patience_improves_expected_W`.
Micro-example. Volatile environment: deferring act reduces H spike from 0.6 to 0.2; ΔW improved.

7.11 Humility — self-model correction to consensus

Lean: `IndisputableMonolith/Ethics/Virtues/Humility.lean`
Role. Adjust self-assessed skew toward external consensus.
Definition. Update local $\kappa_{\text{self}} \leftarrow \kappa_{\text{self}} + \alpha\, \text{sign}(\kappa_{\text{ext}} - \kappa_{\text{self}})$ with α fixed by φ-tier; apply Π_{LA}.
Invariants. $\sigma = 0$; energy bounded; normalization fixed.
Core theorems. `humility_shrinks_error`; `humility_stabilizes`.
Micro-example. Overconfident agent reduces model error by 40% in two cycles; $H \approx 0$; $\Delta V > 0$.

7.12 Hope — optimism prior support

Lean: `IndisputableMonolith/Ethics/Virtues/Hope.lean`
Role. Prevent paralysis by assigning non-zero mass to favorable outcomes.
Definition. Add optimism prior ϵ to high-utility outcomes s.t. $\sum \epsilon_i = 0$ and gauge invariance holds; evaluate V post-projection.
Invariants. $\sigma = 0$; E neutral; no free weights in prior shape.
Core theorems. `hope_enables_action`; `hope_no_harm` (under feasibility).
Micro-example. Low-probability good branch preserved; selected plan increases W with H unchanged.

7.13 Creativity — φ-chaotic exploration

Lean: `IndisputableMonolith/Ethics/Virtues/Creativity.lean`
Role. Explore admissible state space to escape local minima.
Definition. Generate candidates via φ-chaotic sequence over feasible directions; select by Sec. **??**.
Invariants. $\sigma = 0$ maintained; energy cap; φ-tier.
Core theorems. `creativity_improves_best_of_N`; `creativity_preserves_feasibility`.
Micro-example. Randomized admissible probes find $\Delta W = +0.05$ improvement with H unchanged; $\Delta\lambda_2$ small positive.

7.14 Sacrifice — φ-fraction debt absorption

Lean: `IndisputableMonolith/Ethics/Virtues/Sacrifice.lean`

Role. Absorb a φ-fraction of another's debt to reach lower global action.

Definition. Reduce beneficiary $\Delta\kappa$ while increasing sacrificer by $\Delta\kappa/\varphi$; enforce caps and $E^{\text{post}} > 0$.

Invariants. Global $\sigma = 0$; E positivity; φ-fraction fixed.

Core theorems. `sacrifice_lowers_total_action`; `sacrifice_bounded_harm`.

Micro-example. $\Delta\kappa = 3$ relieved; sacrificer takes ≈ 1.854; $H = 0.35$; net $\Delta W > 0$; $\Delta\lambda_2$ improves by 0.01.

8 Generators framework and the DREAM program

Statement (DREAM Theorem, informal). *Completeness.* Any σ-feasible, per-cycle admissible transform on the ledger decomposes (up to least-action projection) into a finite composition of the fourteen virtues in Sec. **??**.

Minimality. No proper subset of the fourteen virtues suffices to generate all σ-feasible transforms under the same invariants (reciprocity conservation, energy constraints, φ-tier normalization).

Proof architecture sketch. We work in a neighborhood of the $\sigma = 0$ manifold with the least-action projector Π_{LA}:

1. **Local closure on $\boldsymbol{\sigma = 0}$.** Linearize admissible deformations (in log-coordinates) subject to balance and $\sigma = 0$; show the space of feasible tangent directions is spanned by primitive moves corresponding to virtue types (equilibration, posting, bounded transfer, horizon-/risk-transforms, gradient-gated acts, energy-cap, asymmetric relief at φ-rates, cooperation update at φ, temporal wait, self-model correction, optimism prior, φ-chaotic probe, φ-fraction sacrifice).

2. **$\boldsymbol{\varphi}$-tier structure.** Uniqueness at the scale level (no weights) pins the coefficients of these primitives by the φ-tier normalization (information scale κ) and curvature normalization ($J''(1) = 1$); hence the basis is canonical.

3. **Energy and $\boldsymbol{\Delta S}$ constraints.** Feasible compositions must satisfy per-cycle energy positivity/caps and the minimax harm bound in the selector; show that the primitive set preserves feasibility under these constraints and closes under least-action completion.

4. **Normal form decomposition.** Any finite admissible act can be ordered (via projection and local commutator control) into a normal form: Justice/Love/Forgiveness blocks for pairwise rebalancing, followed by horizon/risk/meta policies (Wisdom/Prudence/Patience), capped by Temperance, adjusted by Humility/Hope, optionally enriched by Creativity, and with system-level transfers (Compassion/Sacrifice) and reinforcement (Gratitude). This yields completeness.

5. **Minimality by case separation.** For each virtue, construct a class of admissible transforms realizable only when that virtue (or an equivalent) is present (e.g., φ-ratio split for Love; timed posting for Justice; φ^4 relief for Compassion; φ-fraction absorption for Sacrifice; cadence gate for Patience; risk penalty for Prudence). Removing any virtue breaks coverage of at least one such class.

Work plan and intermediate lemmas. Proved. (i) J-convexity and pairwise smoothing; (ii) φ identities and bounds; (iii) reciprocity conservation $\sigma = 0$ minimality; (iv) ΔS invariants (gauge, additivity, composition); (v) $V = \kappa I - \mathcal{C}_J^*$ uniqueness; (vi) selector existence and robustness margins.

In progress. (A) *Local Generator Closure on* $\sigma = 0$ (span of primitive directions); (B) φ*-tier Uniqueness for Primitive Coefficients*; (C) *Normal Form Decomposition* (ordering lemma under Π_{LA}); (D) *Minimality Case Analysis* (virtue-by-virtue necessity).

Key lemmas (targets).

- L1 (closure): the tangent cone of $\mathcal{F}$ at $\sigma = 0$ equals the span of virtue primitives.
- L2 (φ-tier): any scale-compatible decomposition must use φ-fixed ratios; free weights are forbidden.
- L3 (harm composition): ΔS of composed blocks respects additivity/composition bounds.
- L4 (energy viability): sequences obeying Temperance remain within energy positivity region.
- L5 (consent composition): local consent signs compose under least-action projection.
- L6 (robustness): λ_2 is nondecreasing under specified reinforcement blocks.
- L7 (normal form): any finite act admits a block-ordered representation modulo Π_{LA} commutators.

Lean stubs and roadmap. Module. `IndisputableMonolith/Ethics/Virtues/Generators.lean`

Stubs (statements).

- `theorem virtues_complete:` any σ-feasible transform decomposes into virtues.
- `theorem virtues_minimal:` no proper subset generates all σ-feasible transforms.
- `lemma closure_sigma_zero:` local generator closure on $\sigma = 0$.
- `lemma phi_tier_uniqueness:` coefficient uniqueness under φ-tier.
- `lemma normal_form_decomposition:` block-ordering under Π_{LA}.
- `theorem minimality_case_analysis:` virtue-by-virtue necessity.

Roadmap.

1. Implement L1–L2 with explicit bases and φ normalization tests.
2. Prove L3–L4 using convex/separable structure of J and Temperance cap.
3. Establish L5 by differentiating V along feasible directions and composing projections.
4. Prove L6 via spectral perturbation of the σ-graph under reinforcement.
5. Conclude L7 and instantiate `virtues_complete`; finalize `virtues_minimal` via the case constructions.

9 Auditing and selection examples

We illustrate how the audit tuple and lexicographic selector operate in practice. Numbers are illustrative and chosen to make the logic visible; full derivations of the audit protocol are in the companion.

9.1 Case A: hidden skew vs repair-first (virtue framing)

Setup. Three domains (A, B, C) with $\sigma[C_0] = 0$. A naively beneficial proposal P_{good} increases $I(B; E)$, $I(C; E)$ but introduces hidden skew from A to both B and C.

Audit. σ-trace (post) nonzero $\Rightarrow$ fail Feasibility. By Sec. **??**, P_{good} is excluded at step one. A *repair-first* variant P_{rep} pairs each adjustment with a matched return (Love + Justice blocks), so $\sigma[C_1] = 0$.

Virtue framing. $P_{\text{rep}} = \texttt{Justice} \circ \texttt{Love}$ on agent pairs, then Π_{LA}. Outputs: $H(P_{\text{rep}}) = 0.40$; $\Delta W \approx +0.24$; $\Delta\lambda_2 > 0$. Selector picks P_{rep} immediately.

9.2 Case B: consent-sensitive reroute (courage vs prudence)

Setup. Two domains (D, R), feasible state $\sigma[C_0] = 0$. A reroute Q improves D's coupling but degrades R's, yielding $D_D V_R[\xi] < 0$.

Audit. Q: $H = 1.20$, $\Delta W \approx +0.05$, consent fails: $\mathsf{C}(R \leftarrow D) = \text{false}$. An alternative Q_{safe} staggers the change (Patience) and limits surcharge (Courage cap), preserving consent: $D_D V_R[\xi_{\text{safe}}] \geq 0$.

Virtue framing. $Q_{\text{safe}} = \texttt{Patience} \circ \texttt{Courage}$ under cap $H \leq 0.60$. Outputs: $H(Q_{\text{safe}}) = 0.60$, $\Delta W \approx +0.03$, consent passes. Selector discards Q at harm stage and picks Q_{safe}.

9.3 Virtue-specific audits

Love. Two agents (i, j) with $(E_i, E_j) = (1.0, 0.6)$. Apply Love: split E' as $(1/\varphi, 1/\varphi^2)$ of total; σ-trace zero; $H = 0$; $\Delta V = (+0.07, +0.05)$.

Forgiveness. Debtor $\kappa = +4$, creditor energy $E = 1.2$. Bound transfer $\Delta\kappa = 2$; least-action completion spreads surcharge: $H = 0.25$; $\Delta V_{\text{debtor}} = +0.12$, $\Delta V_{\text{cred}} = -0.03$.

Sacrifice. Beneficiary relief $\Delta\kappa = 3$; sacrificer absorbs $\Delta\kappa/\varphi \approx 1.854$; global action drops; $H = 0.35$; $\Delta W > 0$.

Creativity. Generate $N = 16$ admissible probes via φ-chaotic sequence; select by Sec. **??**. Best candidate yields $\Delta W = +0.05$ with H unchanged; $\Delta\lambda_2 \approx +0.01$.

9.4 Robustness margins and uncertainty

Uncertainties on ΔS and V are summarized by sets $\mathcal{U}_{\Delta S}$ and $\mathcal{U}_V$ (intervals/bootstraps). A choice a^* is *robustly certified* when

$$H^{\sup}(a^*) < \min_{b \neq a^*} H^{\inf}(b), \quad \text{or on ties } H^{\sup} = \min H^{\inf} \text{ and } W^{\inf}(a^*) > \max W^{\sup}(b).$$

Otherwise the audit returns *indeterminate* and requests more measurement.

10 Robustness and σ-graphs

At $\sigma = 0$, linear response of the reciprocity network induces a weighted graph G_σ with Laplacian $L_\sigma = D - W$ ($D_{ii} = \sum_j w_{ij}$). For a bounded, mean-zero shock s with $\|s\|_2 \le S$, the least-action response solves $L_\sigma u = s$, $u \perp \mathbf{1}$, and has Dirichlet energy $\mathcal{E}(u) = \frac{1}{2} s^\top L_\sigma^+ s$.

Bound on worst surcharge. Node-wise surcharge obeys $\max_i \Delta S_i \le \mathcal{E}(u) \le \frac{S^2}{2\,\lambda_2(L_\sigma)}$. Thus larger λ_2 guarantees smaller worst-case ΔS under shocks.

Lemma 10.1 (Robust-Preference)**.** *Among σ-neutral W-maximizers tied on present-cycle H, prefer the arrangement with larger λ_2: it minimizes the worst future surcharge across all bounded shocks.*

This provides the third-stage tie-break in Sec. **??**.

Interpretation and policy. Higher algebraic connectivity means strain spreads quickly and thinly: fewer routes for skew to re-emerge and smaller worst bills per cycle under disturbances. When feasible designs tie on harm and welfare, policy should prefer those that increase λ_2 (e.g., add reciprocal ties that do not create skew).

11 Falsifiability and crisp defeats

The framework admits sharp defeats; three top-level falsifiers (from the companion) and virtue-specific defeats are enumerated.

Core falsifiers.

(F1) A durable $\sigma \neq 0$ process with lower action than any $\sigma = 0$ alternative.

(F2) An alternative axiology $\widehat{V}$ satisfying (A1)–(A4) that systematically outpredicts V in audits.

(F3) A temporal aggregator obeying gauge/cadence invariance and separability that differs from the undiscounted eight-tick sum.

Virtue-specific defeats.

- Love: reproducible violation of the φ-ratio optimality (energy split outperforms golden split under the same invariants).
- Forgiveness/Compassion/Sacrifice: systematic breaches of the proved ΔS bounds when audits enforce Π_{LA} and invariants.
- Justice/Patience: consistent improvements using late/early posting that break cadence or accuracy assumptions without increasing H.
- Prudence/Wisdom: alternative risk/horizon operators that satisfy the axioms but dominate W on preregistered audits.
- Creativity: admissible exploration procedure that preserves feasibility and dominates the best-of-N improvement rates.

12 Implementation details and reproducibility

Lean module map.

- Virtues (14 generators): `IndisputableMonolith/Ethics/Virtues/{Love, Justice, Forgiveness, Wisdom, Courage, Temperance, Prudence, Compassion, Gratitude, Patience, Humility, Hope, Creativity, Sacrifice}.lean`
- Harm ΔS: `IndisputableMonolith/Ethics/Harm.lean`
- Value V: `IndisputableMonolith/Ethics/ValueFunctional.lean`
- Consent: `IndisputableMonolith/Ethics/Consent.lean`
- Audit and selector: `IndisputableMonolith/Ethics/Audit.lean`
- RS scaffold and support: `IndisputableMonolith/Ethics/ConservationLaw.lean`, `IndisputableMonolith/`
- DREAM framework (statements): `IndisputableMonolith/Ethics/Virtues/Generators.lean`

How to run audits (data structures and flow).

1. Construct a pre-cycle `MoralState` s with ledger L, partition $\{\mathcal{E}_i\}$, and $E > 0$.
2. Choose an action a (virtue transform or composition) acting on finitely many bonds.
3. Apply least-action completion $L' \leftarrow \Pi_{\mathrm{LA}}(a \mid L)$; verify feasibility $\sigma[C_{\mathrm{post}}] = 0$.
4. Compute audit tuple: σ-traces, ΔS matrix and H, per-agent V and $W = \sum f(V)$, λ_2 pre/post, and consent signs $D_j V_i[\xi] \geq 0$.
5. Select via Sec. **??**; emit a `Report` record with inputs, outputs, and uncertainty sets.

Small examples.

- Two-agent Love audit: split energy $(E_i, E_j) \mapsto (E/\varphi, E/\varphi^2)$; $H = 0$; $\Delta V > 0$.
- Forgiveness audit: bounded $\Delta\kappa$ transfer; check H cap and debtor $\Delta V > 0$.
- Sacrifice audit: φ-fraction absorption; verify global action drop and λ_2 change.
- Creativity audit: generate N candidates; select by lexicographic rule.

Artifact checklist.

- **Proofs**: Lean files for virtues, ΔS, V, consent, selector; theorem names cited in Sec. **??**.
- **Code**: Virtue transforms, least-action projector, audit computation (module map above).
- **Notebooks**: example audits reproducing Sec. **??** cases; plots of σ-traces, ΔS distributions, and λ_2.
- **Audit JSON**: machine-readable `Report` bundles with pre/post states, metrics, uncertainty sets, and selector outcome.

Reproducibility and availability. All Lean sources and audit tooling are available in the repository containing this manuscript (module paths in Sec. ??). We will tag a release with an archival DOI (e.g., Zenodo; placeholder doi:10.5281/zenodo.XXXXXXX) and provide a `repro/` folder with scripts to regenerate all tables/figures and JSON reports from Sec. ??. A `LOCKFILE` pins mathlib and Lean toolchain versions to ensure determinism.

13 Limits, scope, and ethics policy

Conditional on RS. All claims are conditional on the RS scaffold: unique J, eight-tick cadence, bridge identities, and $\sigma = 0$ minimality. Empirical refutation of these premises narrows or overturns the ethics layer.

Boundaries and proxy consent. Agent/group partitions are preregistered; boundary rules for shared bonds are fixed and gauge-invariant. When agents are non-competent, audits use conservative proxy consent (lower-envelope axiology) and provide *override-with-repair* only under the lexicographic rule.

Bypass channels and governance. Only typed, audited channels may effect changes. Proposed actions outside instrumented channels are treated as unmeasured and rejected until instrumented. Governance should enforce transparency of audit artifacts and reproducibility checks.

Out of scope. This paper does not claim to derive qualia, private phenomenal experience, or metaphysical theses. It specifies a conservation-plus-selection law on a recognition ledger and an audit protocol to evaluate admissible actions under parameter-free constraints.

14 Related work

Classical ethics. Deontic theories emphasize rules; consequentialism emphasizes utilities; virtue ethics emphasizes dispositions. In this framework, those strands align without weights: $\sigma = 0$ plays the role of a hard *rule* (derived from J-convexity), $V = \kappa I - \mathcal{C}_J^*$ is a single *cardinal* consequential measure forced by invariances, and the fourteen *virtues* are the sparse generators that reliably improve V on the feasible set. No free parameters are introduced at any stage.

Social choice. Arrow-type impossibilities concern aggregation of *ordinal* preferences over an *unrestricted* domain with independence axioms. Here, rankings occur on a *constrained* feasible set ($\sigma = 0$ under least action), with a *single cardinal* value functional V fixed by physical invariances, and the multi-criteria comparison is *lexicographic* (feasibility $\to$ harm $\to$ welfare $\to$ robustness $\to$ φ-tier). The conditions that drive Arrow's theorem do not obtain; the paradox dissolves rather than being contradicted.

Convex optimization. The use of $J(x) = \frac{1}{2}(x + 1/x) - 1$ frames feasibility and least-action completion as convex programs; ΔS is a marginal of a separable convex objective; $\mathcal{C}_J^*$ is a J-Bregman penalty. The selector is *not* a weighted sum (which would introduce a free scale) but a lexicographic order compatible with the invariances.

Information theory. The informational part of V follows the classical characterizations of mutual information (e.g., Faddeev/Csiszar grouping, additivity, concavity). The scale κ is fixed once by a φ-tier normalization; $J''(1) = 1$ pins the curvature of the mechanical penalty, closing any room for tuning.

Why paradoxes dissolve. Many familiar tensions stem from (i) mixing infeasible acts, (ii) aggregating incomparable ordinals, or (iii) introducing ad-hoc weights. Enforcing $\sigma = 0$, using a single forced cardinal V, and comparing actions lexicographically remove those failure modes: no skew cycles, no arbitrary tradeoffs, and no domain pathologies.

15 Conclusion

We presented a complete, parameter-free operational ethics on top of the RS scaffold. The fourteen virtues form a *sparse, φ-tied generator basis* for admissible moral dynamics on the $\sigma = 0$ manifold. Selection is *audited*, not argued: feasibility first; then minimize worst externalized surcharge; then maximize a single, curvature-normalized welfare; then prefer robustness; with residual ties broken on the same φ-tier arithmetic that calibrates physics. The calculus is falsifiable, reproducible, and free of tunable knobs. If RS is correct, this is the unique way to turn a conservation law into an operational moral code.

A φ-identities and numeric bounds

We collect elementary identities and bounds for the golden ratio

$$\varphi \;=\; \frac{1+\sqrt{5}}{2} \;\approx\; 1.618033989\,.$$

Algebraic identities.

$$\varphi^2 = \varphi + 1, \qquad \frac{1}{\varphi} = \varphi - 1, \qquad \frac{\varphi}{1+\varphi} = \frac{1}{\varphi}, \qquad \frac{1}{1+\varphi} = \frac{1}{\varphi^2}. \tag{A.1}$$

Numeric bounds (useful constants). For audit tolerances and rate estimates we use:

$$\frac{3}{5} \;<\; \frac{1}{\varphi} \;<\; \frac{13}{20}, \qquad \frac{7}{20} \;<\; \frac{1}{\varphi^2} \;<\; \frac{2}{5}, \tag{A.2}$$

$$0 \;\le\; 1 - \frac{1}{\varphi} \;<\; 1, \qquad \varphi \;>\; 1, \qquad \frac{1}{\varphi} \;<\; 1. \tag{A.3}$$

Convergence facts. The geometric sequence with ratio $(1 - 1/\varphi)$ converges absolutely; successive φ-rate updates (e.g., gratitude protocol) are monotonically increasing and bounded by 1, guaranteeing convergence.

B J-convexity and pairwise smoothing calculus

Let

$$J(x) = \tfrac{1}{2}(x + x^{-1}) - 1, \qquad x > 0,$$

and write $x = e^{\alpha}$ so that $J(x) = \Phi(\alpha) := \cosh\alpha - 1$ with $\Phi''(0) = 1$.

Strict convexity near unity. For any $\varepsilon \neq 0$,

$$J(1+\varepsilon) + J(1-\varepsilon) \;=\; 2\,\Phi\big(\ln(1+\varepsilon)\big) \;>\; 0. \tag{B.1}$$

Closed-cycle inequality. If $\{\alpha_k\}_{k=1}^m$ satisfies $\sum_k \alpha_k = 0$ (product-one constraint on multipliers), then by Jensen's inequality for cosh,

$$\sum_{k=1}^{m} \Phi(\alpha_k) \;\geq\; m\big(\cosh 0 - 1\big) \;=\; 0, \tag{B.2}$$

with equality iff $\alpha_k \equiv 0$. Thus the unit configuration minimizes the sum of per-bond costs on closed chains.

Pairwise smoothing step. For opposite-signed logs $\alpha > 0$ and $-\beta < 0$, replacing them by $(\alpha - t)$ and $-(\beta - t)$ for $t \in (0, \min\{\alpha, \beta\}]$ preserves the pairwise sum and strictly reduces $\Phi(\alpha) + \Phi(-\beta)$ by convexity. Iterating over all imbalanced pairs strictly decreases the total action until all pairs balance at unity.

Cycle minimality at $\boldsymbol{\sigma = 0}$. Decomposing a cycle into bidirectional exchanges and applying the smoothing step yields: a per-cycle action is minimized subject to balance constraints iff all pairwise net skews vanish ($\sigma = 0$). Any $\sigma \neq 0$ implies a strictly positive, avoidable action surplus.

C MP → Ledger formalization and ledger equivalence

We sketch the formal path from the Meta-Principle (MP) to a double-entry ledger with exactness and balance. Lean anchors live in `Verification/Necessity/LedgerNecessity.lean`.

Local finiteness. From atomicity (T2: no concurrency) and δ-units (T8) we obtain finiteness of in/out neighborhoods for each event e:

$$\textbf{AtomicFiniteOut:}\quad \{v \mid \texttt{evolves}\ e\ v\} \text{ is finite,} \tag{C.1}$$

$$\textbf{UnitsFiniteIn:}\quad \{v \mid \texttt{evolves}\ v\ e\} \text{ is finite.} \tag{C.2}$$

These yield a locally finite digraph representation (**LocalFinite_from_RS**).

Exactness ⇒ balance (node conservation). Define a totalized edge value `edgeVal` and finitary inflow/outflow sums based on the local neighborhoods:

$$\text{inflowSum } f\ e := \sum_{v \in \text{inNeigh}(e)} \text{edgeVal } f\ (v, e), \quad \text{outflowSum } f\ e := \sum_{v \in \text{outNeigh}(e)} \text{edgeVal } f\ (e, v). \tag{C.3}$$

Under exactness on real cycles (closed-chain sums vanish), node-wise balance holds strongly:

$$\textbf{ConservationLawStrong:}\qquad \forall e,\ \text{inflowSum } f\ e \;=\; \text{outflowSum } f\ e. \tag{C.4}$$

Exactness on real cycles. Represent a finite cycle as a vector $\langle e_0, \ldots, e_{n-1} \rangle$ with successor map; define the closed-chain sum `closedChainSum_real`. Well-foundedness of the reverse-evolution relation (MP forbids cycles on the empty type) implies no directed cycles; hence

$$\textbf{MP\_implies\_exactness\_real:} \qquad \forall \,\text{real cycles } C,\ \texttt{closedChainSum\_real}(f, C) = 0. \tag{C.5}$$

Combined with local finiteness this yields **exactness_implies_conservationStrong_real**.

Divergence decomposition. On locally finite subgraphs, node divergence (outflow−inflow) decomposes into finitely many real cycles through the node; exactness kills each closed-chain sum, giving balance. (Lean skeleton: `divergence_decomposes_into_cycles_real`.)

Ledger equivalence. A balanced flow graph is equivalent to a double-entry ledger: there exists a ledger L whose carrier is isomorphic to `Event` and whose debit/credit columns realize the same conservation law (**graph_with_balance_is_ledger**). The bridge theorem **mp_implies_conservation** packages the path

$$\text{MP} \ \Rightarrow\ \text{LocalFinite} \ \wedge\ \text{Exactness} \ \Rightarrow\ \text{Node Balance} \ \Rightarrow\ \text{Ledger}.$$

D MI–curvature uniqueness and κ normalization

We fill in details of the uniqueness proof for the forced axiology $V = \kappa I(A;E) - \mathcal{C}_J^*$ and the φ-tier normalization of κ.

Axioms and splitting. On per-cycle, $\sigma = 0$ states, let V satisfy: (A1) gauge invariance under the RS bridge; (A2) additivity on independent subsystems; (A3) concavity under coarse aggregation; (A4) curvature normalization tied to $J''(1) = 1$. Define the mechanical deficit as the reduced J-penalty under least-action completion,

$$\mathcal{C}_J^*(p_{AE}, x) \ := \ \inf \Big\{ \sum_e J(x'_e) : \ (p_{AE}, x') \sim (p_{AE}, x), \ \text{balanced}, \ \sigma = 0 \Big\}. \tag{D.1}$$

Set $U := V + \mathcal{C}_J^*$. Then U depends only on p_{AE} and inherits: invariance to relabelings, additivity on independent pairs, and concavity on mixtures.

Informational term is mutual information. By standard characterization (Faddeev/Csiszar), the only symmetric, continuous, concave, zero-on-independence, additive dependence functional with a grouping/chain property is a positive multiple of mutual information:

$$U(p_{AE}) \ = \ \kappa \, I(A;E). \tag{D.2}$$

The multiplicative freedom $\kappa > 0$ is the *only* degree of freedom left by (A1)–(A3).

Mechanical term is J-Bregman. At fixed p_{AE}, $\mathcal{C}_J^*$ must be (i) invariant under ledger gauge, (ii) additive across bonds, (iii) strictly convex with unit curvature at the unit-effort point. The unique convex even generator with these properties is $\Phi(\alpha) = \cosh \alpha - 1$ in log-coordinates (equivalently J in multiplicative coordinates); minimizing over Π_{LA} gives $\mathcal{C}_J^*$.

Small-strain regime. For $x_e = 1 + \varepsilon_e$ with $\sum_e \varepsilon_e = 0$ (balance),

$$\mathcal{C}_J^*(p_{AE}, x) \;=\; \sum_e J(x_e) \;=\; \tfrac{1}{2}\sum_e \varepsilon_e^2 \;+\; o(\|\varepsilon\|^2), \tag{D.3}$$

so V reduces to $\kappa I(A;E)$ minus a quadratic mechanical penalty with unit curvature, as required by (A4).

κ normalization via φ-tier. Fix κ once by the RS convention: a canonical noiseless binary link per eight-tick cycle registers one φ-tier of value, $U = \kappa I = \kappa \cdot 1$. This pins κ; thereafter no continuous parameter remains. Any attempt to rescale κ independently of $J''(1) = 1$ or cadence invariance would introduce a forbidden free knob.

E No-Arbitrary-Discount: full proof and variants

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

Axioms.

(D1) **Gauge invariance (bridge):** joint re-anchoring $(\tau_0, \ell_0) \mapsto (s\tau_0, s\ell_0)$ preserving $c = \ell_0/\tau_0$ does not change rankings.

(D2) **Cadence invariance:** rotation by multiples of eight ticks leaves $\mathfrak{A}$ unchanged; concatenation on independent horizons adds: $\mathfrak{A}(\mathbf{h}\|\mathbf{k}) = \mathfrak{A}(\mathbf{h}) + \mathfrak{A}(\mathbf{k})$.

(D3) **Separability and monotonicity:** $\mathfrak{A}$ is strictly increasing in each argument and continuous.

Theorem E.1 (No-Arbitrary-Discount)**.** *Under (D1)–(D3), there is $\alpha > 0$ such that $\mathfrak{A}(\mathbf{h}) = \alpha \sum_t h_t$ for all $\mathbf{h} \in \mathcal{H}$.*

Proof. By (D2), all positions are equivalent under eight-tick rotations; define $\alpha := \mathfrak{A}(e^{(t)}) > 0$ for unit vectors $e^{(t)}$. By concatenation and continuity, $\mathfrak{A}(\lambda \mathbf{1}_n) = n\,\phi(\lambda)$ with ϕ continuous, increasing, and Cauchy-additive on $\mathbb{R}_{\geq 0}$; hence $\phi(\lambda) = \alpha\lambda$. For a general $\mathbf{h} = \sum_t h_t e^{(t)}$ use independence additivity across blocks and rotation invariance to obtain linearity: $\mathfrak{A}(\mathbf{h}) = \sum_t \mathfrak{A}(h_t e^{(t)}) = \alpha \sum_t h_t$. □

Variants.

- **Weighted rotations forbidden:** any phase-dependent weights (w_t) violate cadence invariance unless w_t is constant across the eight positions and stable under concatenation (reducing to the sum case).
- **Exponential/hyperbolic discounts excluded:** any rate parameter introduces a time scale that rescales under (D1), breaking gauge invariance.
- **Segmented horizons:** if the horizon factorizes into independent segments, (D2) implies segment-wise additivity and the same linearity argument applies per segment.
- **Upper/lower envelopes:** uncertainty sets for H produce interval-valued aggregators $\mathfrak{A}^{\inf/\sup}$; the same structure yields $\mathfrak{A}^{\inf/\sup}(\{h_t\}) = \alpha \sum h_t$ with the same α.

F DREAM scaffolding: closures, normal forms, minimality cases

We detail the scaffolding that underpins Sec. ??: local closures on $\sigma = 0$, normal forms for compositions under Π_{LA}, and witness classes for minimality.

Local closures (tangent space at $\sigma = 0$). Let $\mathcal{F}$ be the feasible manifold. In log-coordinates, feasible directions satisfy a homogeneous linear system (balance and pairwise skew constraints). The primitive directions induced by the virtues span this tangent space:

$$T_s\mathcal{F} \;=\; \mathrm{span}\{\, \mathbf{v}_{\texttt{Love}},\, \mathbf{v}_{\texttt{Justice}},\, \ldots,\, \mathbf{v}_{\texttt{Sacrifice}} \,\}. \tag{F.1}$$

Commutator-like combinations $[\mathbf{v}_i, \mathbf{v}_j]$ appear due to alternating act/project steps; closure is preserved since Π_{LA} is the minimizer of a convex functional.

Normal forms (block ordering). For any finite composition, there exists an equivalent block-ordered form modulo Π_{LA} commutators:

$$\mathcal{V} \;=\; \underbrace{(\texttt{Justice} \circ \texttt{Love} \circ \texttt{Forgiveness})^*}_{\text{pairwise rebalancing}} \circ \underbrace{(\texttt{Wisdom} \circ \texttt{Prudence} \circ \texttt{Patience})^*}_{\text{horizon/risk/meta}} \circ \underbrace{\texttt{Temperance}}_{\text{energy cap}} \circ \underbrace{(\texttt{Humility} \circ \texttt{Hope} \circ \texttt{C}}_{\text{model/prior/s}} \tag{F.2}$$

Here * denotes that repetitions are allowed until no further improvement occurs under the lexicographic rule.

Minimality witness classes. For each virtue, construct transforms only realizable with that generator (or an equivalent):

- `Love`: φ-ratio energy splits that minimize pairwise action at fixed total.
- `Justice`: posting with eight-tick timeliness and exact accuracy.
- `Forgiveness/Compassion/Sacrifice`: asymmetric $\Delta\kappa$ transfers with φ-fixed factors and bounded H.
- `Wisdom/Prudence/Patience`: horizon/risk/meta policies respecting cadence with no discount knobs.
- `Temperance`: per-cycle energy-capped feasibility that cannot be replicated by other blocks without violating $E > 0$.
- `Humility/Hope`: self-model correction toward consensus and optimism prior support without altering scales.
- `Creativity`: φ-chaotic feasible exploration producing improvements unreachable by deterministic local steps.
- `Gratitude`: φ-rate cooperation update with geometric convergence guarantees.

Removing any block eliminates coverage of at least one witness class, establishing minimality.