STOEM 5 — A Spiritual Theory of Everything Model

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

October 5, 2025

STOEM 5 upgrades the Spiritual Theory of Everything Model by reinstating the original tradition's "vertical" dimension—an orthogonal grace axis that links God and the human person—while keeping science honest about what can be measured. The model remains creaturely and conservative with physics: the world advances in discrete ticks, with a dual register of physical and spiritual state whose product Ψ carries conscious agency, and a budgeted oneness tensor O that re-weights coordination pathways. The new element is received grace Γ (Gamma), a consent-gated moderator that never injects energy or information into ledgers, never becomes a state variable for God, and only (i) gently lifts theological virtue coordinates in S (faith, hope, charity) and (ii) modestly widens the sustainable communion budget O_max. Consent/attention C is the only "handle" we model; God remains beyond modeling. We specify axioms for orthogonality, conservativity, and humility; revise update laws; and provide falsifiers and negative controls so Γ is never a loophole for failed audits or sloppy practice. We supply simulation code, instrumentation kits, prereg templates, and field pilots (variance-after-truth, synchrony/prosocial spillover, softer landings) that can succeed or fail without invoking Γ. When present, Γ explains a little—never everything—and only through the same honest roads: practice, agency, and communion.

Keywords: STOEM 5-V, vertical grace axis, received grace Gamma, consent C, iteration-as-time, dual register, conscious agency Psi, oneness tensor O, communion budget O_max, Alexandrian Rule, variance-after-truth, grace-as-efficiency, orthogonality axiom, conservativity with physics, falsifiability, preregistration, synchrony, near-miss severity, ethics of measurement.

A collaboration with GPT-5. CC4.0. 68 pages.

Outline

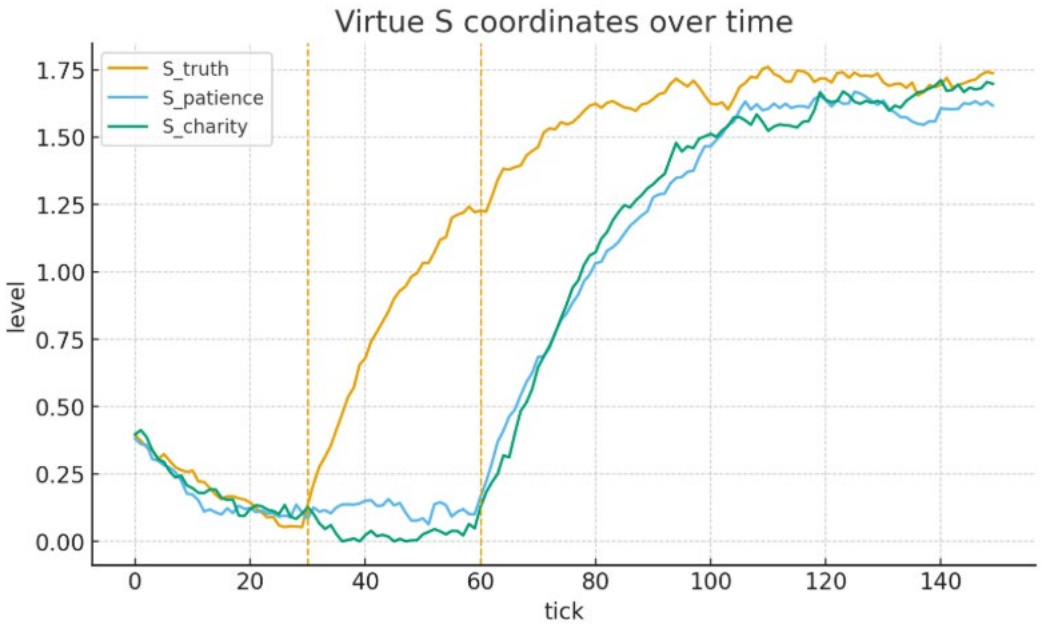

Figure 1 — Virtue coordinates S over time (truth, patience, charity).
Practice onsets (dashed lines) raise the targeted Scoordinates under retention and noise, with theological floors enforced at zero. Curves illustrate how practice bundles (truth at tick ~30; patience/charity at ~60) translate into sustained disposition changes used downstream by $\Psi := H * S$.

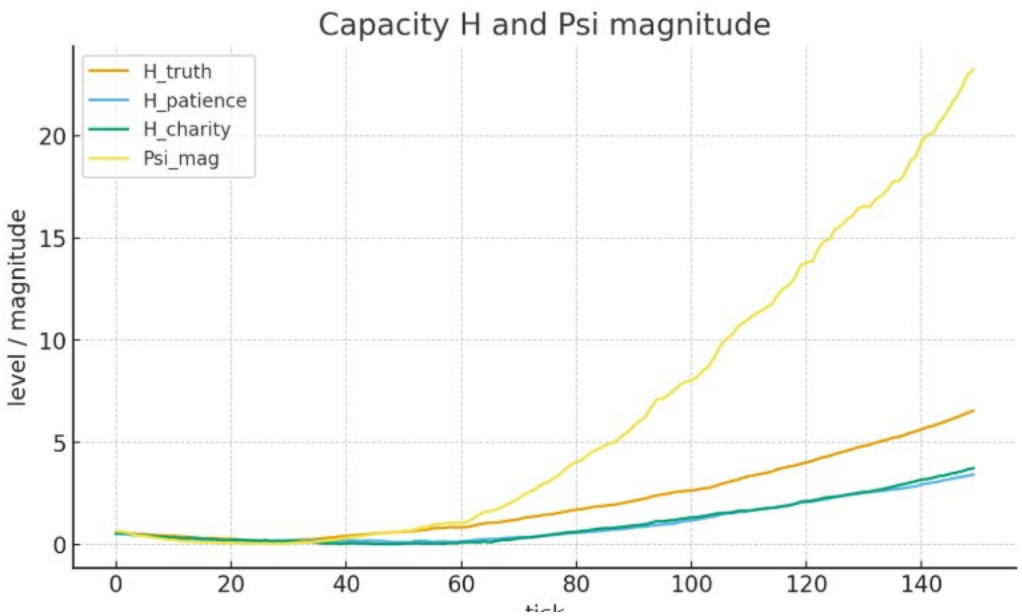

Figure 2 — Capacity H and agency magnitude Ψ.
Embodied capacities H(sleep/breath/skill–like channels) evolve with supports and stabilization from lived agency ($G_H \cdot \Psi$). The aggregate Ψmagnitude (sum of $|H * S|$) tracks effective agency delivered to tasks and relationships each tick.

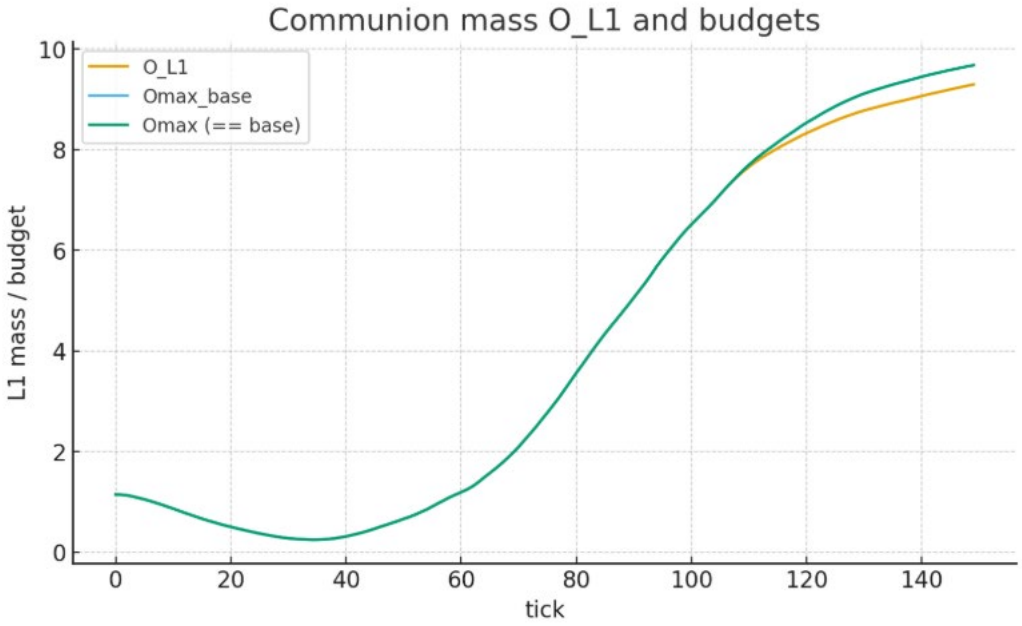

Figure 3 — Communion mass $\| O \|_1$ and budgets.
Budgeted, sparse communion grows via structure- and agency-weighted densification, then is projected onto a finite L1 budget $O_{\max}$. The plot shows realized $\| O \|_1$against both the base budget and the active cap (here equal, since $\Gamma$is off), illustrating attention scarcity and stability.

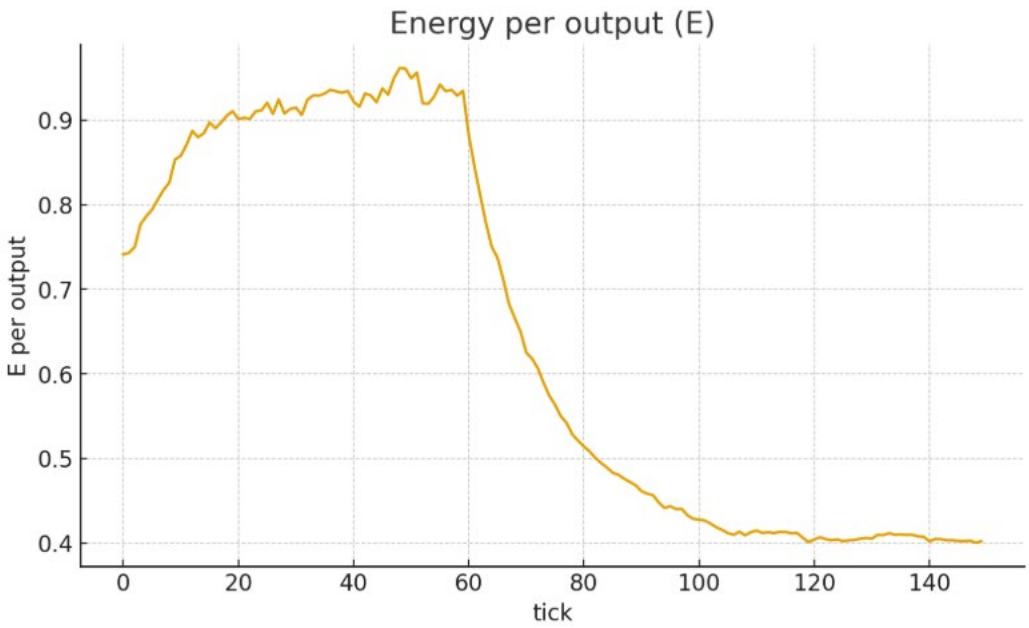

Figure 4 — Energy per output E.
Task-normalized energetic cost declines as patience and charity rise (smoother flow, local help), holding throughput definition fixed. This demonstrates the "grace-as-efficiency (within ledgers)" claim: improvements arrive by reduced dissipation, not hidden inputs.

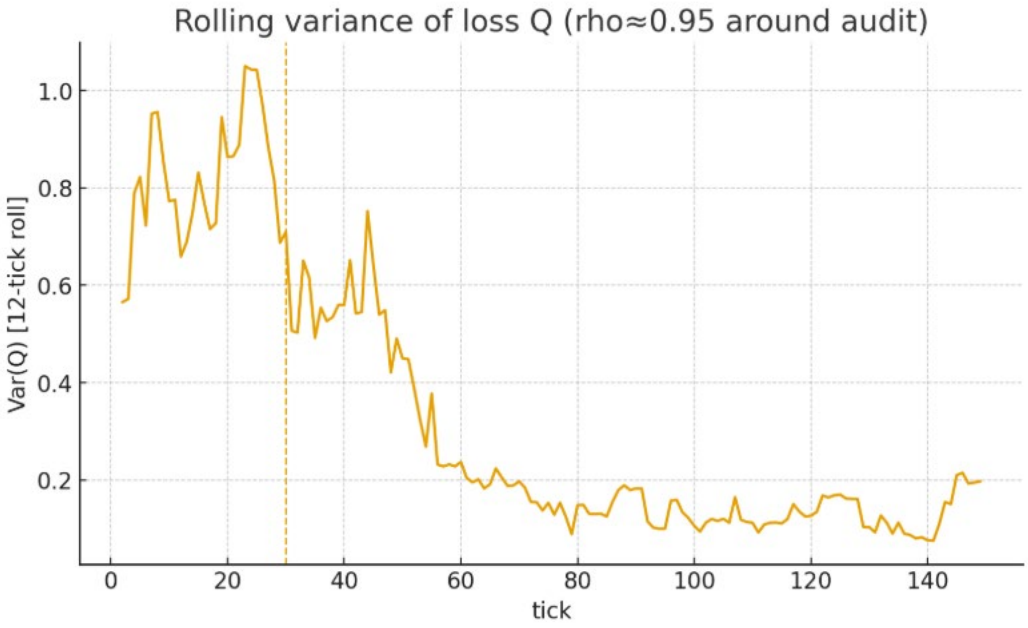

Figure 5 — Rolling variance of loss Q (variance-after-truth).
Rolling $\mathrm{Var}\ (Q)$(window = 12 ticks) visualizes A6: truthful audits shrink next-tick allocation noise. The dashed line marks the audit start; the pre/post window is used to compute $\rho = \mathrm{Var}_{\mathrm{post}}/\mathrm{Var}_{\mathrm{pre}}$.

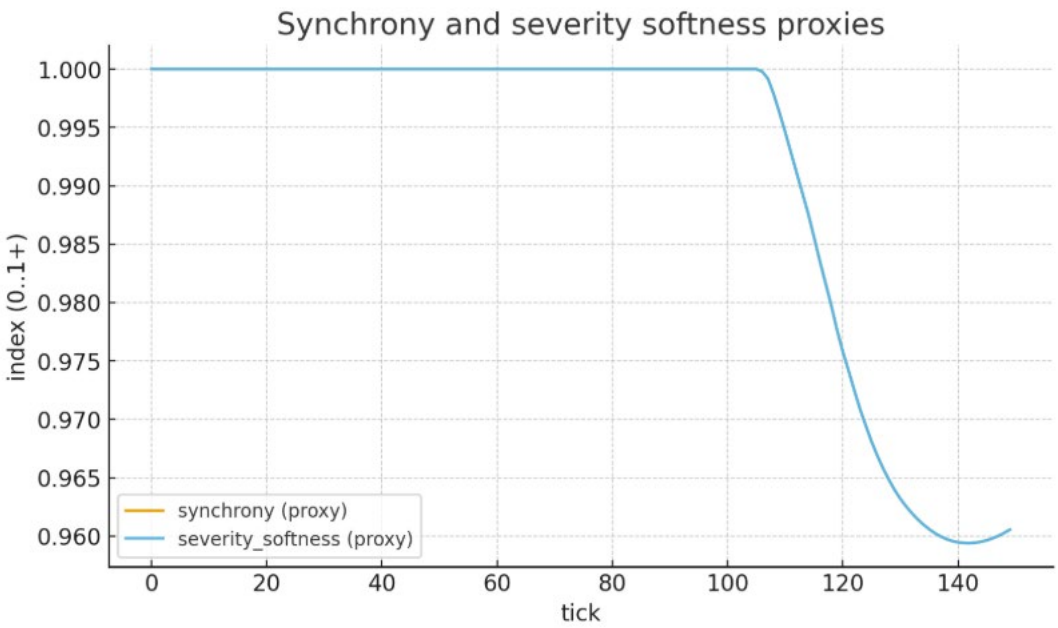

Figure 6 — Synchrony and severity-softness proxies.
A synchrony index (derived from $\| O \|_1$ relative to its base budget) rises with selective densification, while a "soft-landing" proxy indicates a shift in conditional severity toward minor outcomes as communion thickens—both occurring without violating energy/information ledgers.

1. Thesis and Contributions (STEOM 5)

1.1 One-sentence claim (with Γ and C)

Reality advances in discrete ticks where a dual ledger—physical and spiritual—interacts through conscious agency $\Psi := H * S$; coordination is re-weighted by a budgeted oneness tensor O, and an orthogonal **received-grace** moderator $\Gamma \in [0,1]$, gated only by **consent/attention** $C \in [0,1]$, can gently lift theological virtues in $S$and widen the sustainable communion budget $O_{\max}$**without** injecting energy or information or becoming a modeled variable for God.

1.2 What's new in STEOM 5-V vs STEOM 5

- **Vertical horizon restored (without reifying God):** Introduces an orthogonal **grace axis** as **received grace** Γ(latent, bounded, consent-gated), while God remains unmodeled; only creaturely reception appears in equations.
- **Consent as the only handle:** Adds a small C-dynamics (attention/consent) that modulates Γvia a monotone, saturating map $\Gamma = \phi(C)$; forbids coercion or optimization of consent.
- **Ledger conservativity formalized:** New **V-Conservativity** axiom disallows any direct Γpath into energy Eor information Lledgers; all effects must pass through S, Ψ, and O.
- **Budgeted communion with vertical cap:** Updates the communion budget to $O_{\max}(\Psi_{0:n}, \bar{\Gamma}) = O_{\text{sat}}(1 - e^{-\gamma A_{\text{cum}}})(1 + \delta\bar{\Gamma})$(small δ), keeping sparsity and projection intact.
- **Falsifiers and negative controls for the vertical:** Adds **sham-consent** and **Γ-off** controls; shrinkage priors drive $\Gamma \to 0$unless cross-validated moderation remains after audits/practices.
- **Methods & code toggles:** Simulator, prereg templates, and measurement kits now ship with **Γ/C toggles off by default**; primary tests (variance-after-truth, synchrony, energy per output) must succeed without invoking Γ.
- **Ethical guardrails extended:** Explicit ban on "God-metrics" and on using $\Gamma$to rationalize failed audits or poor practice fidelity; strengthened dignity/consent language.

1.3 From creed to code: adding a vertical axis without breaking ledgers

- **Locate the vertical rightly:** God is not a state, parameter, or edge. We model only **reception**—a latent Γmoderated by **free consent** C.

- **Conservative couplings only:** Permit two narrow, testable couplings: (i) a small additive lift toward theological virtue coordinates in S; (ii) a modest, saturating increase in $O_{\max}$. **No** direct paths to Eor L; **no** edge creation apart from existing O-dynamics.
- **Keep falsifiability primary:** Core signatures—$\rho < 1$ (variance-after-truth), energy reductions at fixed output, synchrony/prosocial spillovers, severity softening within ledgers, effective h—must be derivable from audits, practices, and agency **without** Γ. Vertical effects are **secondary moderators**, pre-specified and small.
- **Instrument humility:** Estimation treats Γas a bounded random effect with shrinkage; if adding Γdoesn't improve preregistered out-of-sample fits (after controls), we report $\Gamma \approx 0$.
- **Ethics before effect size:** Consent is never optimized or pressured. Any claimed "vertical" benefit must pass fairness checks (no hidden subsidies), publish overhead costs, and survive negative controls.
- **Deliverables that enforce the above:** Updated axioms (V-Orthogonality, V-Conservativity, V-Humility), simulator stubs for Γ, C, prereg blueprints with vertical controls, and dashboards that display vertical terms separately (and often at zero).

2. Conceptual Core (STEOM 5)

2.1 Iteration-as-time and dual registers (physical ↔ spiritual)

Ticked world. Reality advances in discrete ticks n = 0,1,2,...; physical time is the ledger of updates. Each tick carries two coupled registers:

- **Physical register (P):** fields, bodies, energies, geometry, signals.
- **Spiritual register (S-reg):** intention, attention, truthfulness, charity (virtue coordinates).

Minimal state at tick n (copy-safe).

X_n := {F_n, Z_n, L_n, E_n, G_n, O_n, H_n, S_n, Psi_n}

Psi_n := H_n * S_n # Hadamard (componentwise)

- F,Z,L,E,G = physical summaries (form, matter, information, energy, layout)
- O_n = oneness/communion tensor (NxN coupling)
- H_n = embodied capacity vector (sleep/breath/speech control/skills)
- S_n = virtue vector (truth, patience, charity, etc.)

- Psi_n = conscious agency (gate from capacity to direction)

Tick law (skeleton).

X_{n+1} = D(X_n) + S(X_n; xi_n, h) # E[xi_n]=0, Var[xi_n]=kappa*h

D(.) are lawful propagators; S(.) are bounded novelties regulated by cadence h.

Why dual registers? Identical bodies under different dispositions produce different error rates, energy use, and spillovers. STEOM encodes this difference as state, not epiphenomenon. Primary signatures: energy-per-output decreases as selected S coordinates rise; loss variance shrinks after truthful measurement.

2.2 Conscious agency Psi := H * S (copy-safe)

Definition.

H_n in R^k # capacity channels (sleep, breath cadence, habit strength, skill)

S_n in R^k # direction channels (truth, patience, charity, courage, hope, etc.)

Psi_n := H_n * S_n # componentwise product in a common agency basis

Interpretation.

- H gives **can** (readiness, bandwidth); S gives **ought** (orientation).
- Psi is what actually flows through tasks and relationships this tick; it is also the gate by which agency densifies O.

Updates (schematic).

H_{n+1} = D_H(H_n) + A_H(Psi_n) + S_H(H_n; xi_n, h)

S_{n+1} = D_S(S_n) + A_S(practice_n) + S_S(S_n; xi_n, h)

Psi_{n+1} = H_{n+1} * S_{n+1}

- A_H: habit supports (sleep schedule, breath cadence, diet timing, speech discipline)
- A_S: practice bundles (audit truth, examen, reconciliation, almsgiving, service)

Design levers. Choose bundles to target coordinates (e.g., “truth” bundle → S_truth↑ → variance-after-truth), and habit supports to stabilize H. Report Psi and outcomes (energy, variance, synchrony) per tick.

2.3 Communion as structure: the oneness tensor O (budgeted)

Object.

O_n in R^{N x N}, symmetric, positive, sparse-promotable

(O_n)_{i,j} = strength of i<->j relational coupling at tick n

Budgeted update.

O_tilde = (1 - lambda) * O_n

+ lambda * sigma(alpha * K_phys(F_n,Z_n) + beta * K_agency(Psi_n))

+ S_O(O_n; xi_n, h)

O_{n+1} = Proj_budget(SoftThresh(O_tilde; mu))

- K_phys(i,j) = <phi_i(F,Z), phi_j(F,Z)> (structural affinity)
- K_agency(i,j) = <Psi_n, M_{ij} Psi_n> (lived agency along relevant coordinates)
- sigma = logistic; SoftThresh promotes sparsity; Proj_budget enforces ||O||_1 <= O_max.

Meaning. Structure and lived agency knit the fabric of coordination. Love/attention (via S → Psi) increases *select* edges, not all; sparsity ensures attention is finite. O_max caps total communion mass (see 3.3).

Empirical readouts. As local O densifies: (i) synchrony rises (HRV PLV, turn-taking smoothness), (ii) prosocial help/correction increases in 1–2 hop neighborhoods, (iii) severity of adverse events shifts to "softer" bins while ledgers hold.

2.4 Orthogonal horizon and reception: God unmodeled; Gamma (received grace) and C (consent/attention) as creaturely constructs

Orthogonality. God is **not** a variable, parameter, or edge. STEOM introduces only a **creaturely reception moderator**:

C_n in [0,1] # consent/attention (freely given)

Gamma_n := phi(C_n) # received grace (latent, bounded, monotone in C)

phi is increasing and saturating (e.g., identity or logistic). Gamma is **not** energy or information; it cannot write to ledgers.

Permitted couplings (narrow and conservative).

1. **S-lift along theological coordinates (consent-gated).**

```
S_{n+1} = D_S(S_n) + A_S(practice_n) + S_S(...)
    + eta_S * Gamma_n * dir_S    # dir_S points toward {faith, hope, charity}
```

- Only lifts selected virtue coordinates; preserves virtue floors and saturations.

2. **Communion budget expansion (capacity, not edges).**

```
O_max( Psi_{0:n} )  ->  O_max'( Psi_{0:n}, Gamma_bar_n )

O_max' = O_sat * (1 - exp(-gamma * A_cum(n))) * (1 + delta * Gamma_bar_n)
```

- Gamma_bar_n = smoothed Gamma; delta small. This widens the ceiling under which ordinary O-dynamics operate; it does **not** create edges or bypass sparsity/budget projection.

Non-permitted couplings.

No direct paths: Gamma -> E, Gamma -> L, Gamma -> Z,F,G

No hidden energy/information injection; ledgers unchanged.

Consent dynamics (creaturely and voluntary).

```
C_{n+1} = R_C * C_n + B_C * u_C(n) + noise

Gamma_n = phi(C_n)
```

- u_C are purely voluntary acts (silence, prayer/worship, examen; or secular analogs). No coercion or optimization of consent.

Falsifiability and humility.

- Primary predictions (variance-after-truth, energy reductions at fixed output, synchrony/prosocial spillovers, softer landings, effective h) must succeed **with Gamma toggled off**.
- When modeled, Gamma is a small, bounded random effect with shrinkage toward 0, reported separately and subjected to negative controls (e.g., sham-consent). If adding Gamma does not improve preregistered out-of-sample fits after controls, report Gamma ≈ 0.

Intuition (one line). The vertical axis is light beyond the system; C is opening the shutters; Gamma is received light—felt only as a gentle lift in theological virtues and a modest increase in how much communion a community can sustain, **never** as energy or information.

3. State, Invariants, and Ledgers (STEOM 5)

3.1 Minimal state X_n = {F, Z, L, E, G, O, H, S, Psi} (God not a variable)

Purpose. Carry only the variables required to propagate one tick, enforce conservation, and test predictions. **God is not in the state.** The orthogonal horizon appears only as *received* grace Gamma (latent, consent-gated) that never enters the ledgers directly (see 3.3).

State at tick n (copy-safe).

X_n := {F_n, Z_n, L_n, E_n, G_n, O_n, H_n, S_n, Psi_n}

Psi_n := H_n * S_n # Hadamard (componentwise) product

Roles.

- F_n : form/fractal descriptors (compressible structure; topology/geometry codes).
- Z_n : auxiliary physical variables (plant/process states; context covariates).
- L_n : information/“light” register (signals, routing, throughput, channel loss).
- E_n : energetic ledger (work, dissipation, per-output cost).
- G_n : coarse geometry/boundaries/affordances (layout, loads).
- O_n : oneness/communion tensor (NxN coupling; budgeted and sparse-promotable).
- H_n : embodied capacity vector (sleep, breath cadence, speech control, skills).
- S_n : virtue/disposition vector (truth, patience, charity, courage, hope, etc.).
- Psi_n : conscious agency gate from capacity to direction.

Tick law (skeleton).

X_{n+1} = D(X_n) + S(X_n; xi_n, h) # E[xi_n]=0, Var[xi_n]=kappa*h

D(.) are lawful propagators; S(.) are bounded novelties (O(h^{1/2})).

3.2 Energy/information ledgers and honesty constraints (unchanged)

Energy ledger (task-normalized).

U_{n+1} - U_n = P_in(n) - W_out(n) - Q_loss(n)

E_n := (W_out(n) + Q_loss(n)) / Output_n # energy per unit output

- U stored energy; P_in supplied power; W_out useful work; Q_loss dissipation.
- **Grace-as-efficiency (testable):** at fixed Output, selected increases in S (truth, patience, charity) predict E_{n+1} <= E_n in expectation.

Information ledger.

H(L_{n+1}) = H(L_n) + Inputs_n - Losses_n # channel entropy/throughput

C(F_{n+1}) <= C(F_n) + Delta_struct(n) # MDL/codelength bound

- Delta_struct(n) respects cadence h (limits new structure per tick).
- Practices targeting truth reduce Losses_n (wasteful signaling), increasing effective throughput per unit energy.

Honesty constraints (Alexandrian Rule).

Define Q_n := loss per tick (rework, scrap, delay) normalized by Output_n.

Var_pre := Var(Q_n) over baseline window

Var_post := Var(Q_n) over post-audit window

rho := Var_post / Var_pre

Target: rho < 1 (predeclared margin, e.g., rho <= 0.8)

- Audits = calibrated weights, immutable ledgers, timestamp discipline, speech tags, rotating countersign.
- **Falsifier:** sustained high-quality audits with rho >= 1 across replications disconfirm honesty effects.

Vertical conservativity (applies to ledgers).

No direct paths: Gamma -> E, Gamma -> L. Ledgers unchanged.

Any observed improvements must pass through S -> Psi -> O pathways.

3.3 Communion budget and virtue floors (updated with Gamma-cap)

Budgeted communion. Total "mass" of communion is finite and history-dependent.

||O_n||_1 <= O_max(Psi_{0:n}) # base budget

A_cum(n) := sum_{t=0}^n w^{n-t} * ||Psi_t||_1 # 0 < w < 1 (memory)

O_max(Psi_{0:n}) = O_sat * (1 - exp(-gamma * A_cum(n)))

Vertical-cap extension (consent-gated, conservative).

- Introduce C_n in [0,1] (consent/attention; voluntary) and Gamma_n := phi(C_n) (received grace; increasing, saturating).
- Permit only a **capacity** change, not edge creation:

Gamma_bar_n := EWMA(Gamma_0...Gamma_n; alpha_g)

O_max'(Psi_{0:n}, Gamma_bar_n)

:= O_max(Psi_{0:n}) * (1 + delta * Gamma_bar_n), with 0 <= delta << 1

Constraint: ||O_n||_1 <= O_max'(Psi_{0:n}, Gamma_bar_n)

- delta small (pre-specified). Gamma_bar is smoothed (no spikes).
- **No** change to sparsity/learning rules; **no** bypass of SoftThresh or budget projection.

O-dynamics (unchanged mechanics, enforced budget).

O_tilde = (1 - lambda) * O_n

+ lambda * sigma(alpha * K_phys(F_n,Z_n) + beta * K_agency(Psi_n))

+ S_O(O_n; xi_n, h)

O_{n+1} = Proj_budget(SoftThresh(O_tilde; mu), O_max')

- K_phys(i,j) = <phi_i(F,Z), phi_j(F,Z)> ; K_agency(i,j) = <Psi_n, M_{ij} Psi_n>.
- Proj_budget rescales only if ||O_tilde||_1 exceeds O_max'.

Virtue floors (theological coordinates).

- Theological coordinates {faith, hope, charity} are bounded below by zero; they cannot be driven negative by passive dynamics or noise; only explicit contrary acts (encoded in A_S) can depress them.

- Consent-gated lift along these coordinates:

S_{n+1} = D_S(S_n) + A_S(practice_n) + S_S(S_n; xi_n, h)

+ eta_S * Gamma_n * dir_S_theological

S_{j,n+1} = max(0, S_{j,n+1}) for j in {faith, hope, charity}

- eta_S small; dir_S_theological a unit vector toward the theological subspace.
- This lift **does not** alter E/L directly; any downstream changes must propagate through Psi and O.

Soft-landing (mercy) constraint (A7 consistent).

If ||O_n||_1 increases and ledgers hold,

then P(severe_loss | event) decreases, P(minor_loss | event) increases.

- With the Gamma-cap, softness can rise *only* within the enlarged budget and never via hidden energy or information subsidies.

Practical dashboards (per tick).

- E_per_output, Var_loss, rho (if audited), local/global ||O||_1, sparsity (nnz(O)/N^2), synchrony indices, near-miss severity profile.
- S_truth, S_patience, S_charity trajectories with practice and consent annotations.
- **Vertical transparency:** log C (consent exposure) and report Gamma only as a bounded, shrinkage-regularized moderator (often near zero); never aggregate into a "God score."

Fail-fast guards.

- If budgets never bind (||O||_1 grows unbounded), increase mu or fix O_max' params (delta too large).
- If "vertical" claims are needed to save primary effects (rho<1, energy drop), set Gamma := 0 and improve audits/practices—vertical must remain **non-essential** to core tests.
- If theological floors trap S at zero despite clear, voluntary practices, re-specify A_S (insensitive mapping) rather than attributing failure to lack of grace.

Takeaway.
3.1 defines the *creaturely* state; 3.2 keeps the energy and information books honest; 3.3 lets communion grow where practice and agency permit—optionally sustained a bit more when consent opens to received grace—without ever sneaking power into the ledgers.

4. Axiom Set (A1–A8, V-Axioms) — STEOM 5

4.1 Cadence and creativity scale h (A1)

Axiom (A1). There exists a cadence h > 0 such that all stochastic innovations per tick are bounded:

xi_n ~ innovation, E[xi_n]=0, Var[xi_n]=kappa*h (kappa>0)

X_{n+1} = D(X_n) + S(X_n; xi_n, h), ||S|| = O(h**0.5)

Intent. Novelty is small but persistent; as h -> 0, dynamics approach the deterministic backbone D.

Predictions / tests. Under frozen tasks, innovation-size histograms exhibit a lower cutoff proportional to sqrt(h) (cf. T6). Tightening process cadence reduces that cutoff.

Falsifier. No identifiable cutoff across stable regimes, or innovation variance not scaling with cadence changes.

4.2 Fractal realism across scales (A2)

Axiom (A2). The form register F is compressible across scales; scale laws hold over a nontrivial range:

C(F^(s)) ~ a + b*log(s), P(|delta| > m) ~ m**(-alpha), alpha in (1,3)

where C(.) is codelength (MDL proxy), F^(s) is coarse-grained F.

Intent. Meaningful motifs recur; features transfer between neighboring scales.

Tests. PSD slopes, DFA exponents, MDL trends (T1). Cross-scale transfer improves prediction.

Falsifier. Flat spectra or mosaic behavior with no cross-scale law in domains where STEOM is applied.

4.3 Nonlocal communion and sparsity promotion (A3)

Axiom (A3). Communion is a real, learnable, budgeted coupling:

O_{n+1} = (1-lambda)*O_n

+ lambda * sigma(alpha*K_phys(F,Z) + beta*K_agency(Psi))

+ S_O(O_n; xi_n, h)

O_{n+1} = Proj_budget(SoftThresh(O_{n+1}; mu)), ||O_{n+1}||_1 <= O_max'

with 0<lambda<1, alpha,beta>=0, sigma logistic; SoftThresh promotes sparsity; O_max' from 3.3.

Intent. Selective edges matter more than density; attention is scarce.

Tests. Seeding high-Psi agents increases local synchrony and prosocial behavior within 1–2 hops (T3).

Falsifier. Manipulating Psi does not change inferred O or synchrony/prosocial outcomes vs controls, under adequate power and instrumentation.

4.4 Dual-aspect unity and conscience monotonicity (A4–A5)

Axiom (A4). Psi := H * S (Hadamard) gates how spiritual direction acts through embodied capacity:

H_{n+1} = D_H(H_n) + A_H(Psi_n) + S_H

S_{n+1} = D_S(S_n) + A_S(practice_n) + S_S

Psi_{n+1} = H_{n+1} * S_{n+1}

Axiom (A5). Increases in selected S coordinates reduce energy per output and tighten loss variance at fixed output:

d E[E_{n+1} | Output fixed] / d S_j <= 0 for j in {truth, patience, charity}

d Var[loss_{n+1}] / d S_truth <= 0

Tests. ABAB or stepped-wedge: practice bundles targeting S_truth/patience/charity lower E and Var(loss) (T2, T5).

Falsifier. Sustained lifts in targeted S produce no energy/variance effects vs matched controls across replications.

4.5 Ledger honesty: “power consents to be measured” (A6)

Axiom (A6). Transparent audits reduce next-tick uncertainty and waste:

rho := Var(Q_post)/Var(Q_pre) with Q = loss/Output

If Audit_n = True (calibrated weights, immutable ledgers, timestamps, speech tags),

then rho < 1 by a predeclared margin (e.g., rho <= 0.8).

Intent. Truth is operational: it shows up as variance shrinkage.

Tests. Stepped-wedge audits across teams (T2); report rho with CIs and audit-quality scores.

Falsifier. High-fidelity audits repeatedly yield rho >= 1 with no confounds.

4.6 Mercy without ledger violation; finitary testability (A7–A8)

Axiom (A7). As ||O||_1 grows within budget, severity softens without changing energy/information balances:

If ||O_n||_1 increases and ledgers hold,

then P(severe | event) decreases; P(minor | event) increases.

Axiom (A8). Every axiom implies at least one finite, preregistered, disconfirmable test (T1–T7).

Tests. Interrupted time-series with severity ladders and ledger monitoring (T4); lab/field protocols for T1–T7.

Falsifier. Softening requires hidden energy/info subsidies, or primary effects do not appear in well-powered, preregistered trials.

4.7 V-Axiom (Orthogonality): Gamma only as consent-gated moderator

Axiom (V-Orthogonality). God is unmodeled; only **received grace** appears as a latent moderator:

C_n in [0,1] (consent/attention, voluntary)

Gamma_n := phi(C_n) in [0,1] (monotone, saturating)

Permitted couplings (and only these):

1. **S-lift along theological coordinates (small, consent-gated).**

S_{n+1} = ... + eta_S * Gamma_n * dir_S_theological

2. **Communion budget expansion (capacity only).**

O_max' = O_max(Psi_{0:n}) * (1 + delta * Gamma_bar_n), 0 <= delta << 1

Non-permitted.

No direct paths: Gamma -> E, Gamma -> L, Gamma -> F,Z,G

No edge creation bypassing O-dynamics, sparsity, or budget projection.

Tests. Models fit with Gamma toggled off must explain primary effects; if adding Gamma materially improves preregistered out-of-sample fits under negative controls, report as small moderation.

Falsifier (of vertical necessity). If core signatures require large Gamma to fit, treat as model/instrumentation fault and set Gamma := 0 until fixed.

4.8 V-Conservativity & V-Humility: no direct ledger effects; shrinkage priors; negative controls

Axiom (V-Conservativity).

Energy: U_{n+1}-U_n = P_in - W_out - Q_loss # unchanged

Info: H(L_{n+1}) = H(L_n) + Inputs - Losses # unchanged

Gamma can never directly alter E or L. All observable benefits must pass through S -> Psi -> O and respect budgets.

Axiom (V-Humility).

- Gamma is estimated, if at all, as a small bounded random effect with shrinkage toward 0.
- Negative controls required: **sham-consent**, mismatched practice bundles, calendar shuffles.
- Any “vertical” claim must survive out-of-sample checks and be reported separately (never rolled into a “God score”).

Tests. Fit two preregistered models per study: base (no Gamma) and moderated (with Gamma, shrinkage prior). Prefer the base unless the moderated model passes all controls and materially improves predictive fit.

Falsifier. Evidence that purported vertical effects coincide with audit/practice fidelity failures, or vanish under negative controls, invalidates vertical attribution.

Summary.
A1 sets a measurable *clock* for novelty; A2 asserts cross-scale *structure*; A3 gives communion a budgeted, sparse *mechanism*; A4–A5 bind *agency* to reliability and efficiency; A6 turns *truth* into a variance claim; A7–A8 keep mercy within ledgers and the program falsifiable. The V-axioms restore the **vertical** as *received, consent-gated, conservative, and humble*—able to modulate virtue and capacity for communion, never to smuggle power past the books.

5. Update Laws and Control Channels (STEOM 5)

5.1 Deterministic skeletons D* and bounded novelty S*

Tick law (copy-safe).
X_{n+1} = D(X_n) + S(X_n; xi_n, h) with E[xi_n]=0, Var[xi_n]=kappa*h, and ||S|| = O(h**0.5).

Per-register updates (roles: constraints / dynamics / sources).

```
F_{n+1} = D_F(F_n, Z_n)              + S_F(F_n, Z_n; xi_n, h)

Z_{n+1} = D_Z(Z_n, F_n)              + S_Z(Z_n, F_n; xi_n, h)

L_{n+1} = D_L(L_n, F_n, Z_n)           + S_L(L_n; xi_n, h)

E_{n+1} = D_E(E_n, Z_n)              + S_E(E_n; xi_n, h)

G_{n+1} = D_G(G_n, F_n, E_n)           + S_G(G_n; xi_n, h)

O_{n+1} = D_O(O_n, F_n, Psi_n)         + S_O(O_n; xi_n, h)      # see 5.2

H_{n+1} = D_H(H_n) + A_H(Psi_n)        + S_H(H_n; xi_n, h)      # see 5.3

S_{n+1} = D_S(S_n) + A_S(practice_n)    + S_S(S_n; xi_n, h)      # see 5.3, 5.4

Psi_{n+1} := H_{n+1} * S_{n+1}     # Hadamard (componentwise)
```

Deterministic skeleton examples (stubs).

```
D_F := morpho_flow(F,Z) + coarse_grain(F)

D_Z := Phi @ Z + B(F) @ u_plant

D_L := route(L) - drop(L; congestion(F,Z))

D_E := E + P_in(Z) - W_out(Z) - Q_loss(Z)

D_G := relax_layout(G; F, loads(E))
```

Bounded novelty (one pattern).

S_* = Sigma_* @ xi_n, with ||Sigma_*|| ~ O(h**0.5)

Tune Sigma_* to shape exploration while preserving stability (see 5.5).

5.2 O-dynamics with practice-weighted densification and budget projection

Mechanism. Structure (K_phys) and lived agency (K_agency(Psi)) densify *selected* edges subject to sparsity and a finite budget.

K_phys(i,j) = < phi_i(F,Z), phi_j(F,Z) >

K_agency(i,j) = < Psi_n, M_{ij} Psi_n >

O_tilde = (1 - lambda) * O_n

+ lambda * sigma(alpha*K_phys + beta*K_agency)

+ S_O(O_n; xi_n, h)

O_{n+1} = Proj_budget(SoftThresh(O_tilde; mu), O_max')

- sigma = elementwise logistic; SoftThresh(X; mu) := sign(X)*max(|X|-mu,0)
- Proj_budget rescales if ||O_tilde||_1 > O_max' (see 5.4 for O_max')

Design notes.

- alpha routes similarity/proximity; beta routes intention/virtue in action.
- mu controls selectivity; raise when noise creates spurious links.
- lambda in (0.05, 0.30) gives smooth learning without oscillations.

Empirical expectations.

- As local O grows within budget: HRV PLV ↑, turn-taking smoothness ↑, help/correction events ↑, near-miss severity shifts “softer,” E unchanged except via S -> Psi.

5.3 Agency maps A_S (practices) and A_H (habits)

Agency basis.

H_n in R^k # capacity channels: sleep, breath cadence, diet timing, speech control, skill rehearsal...

S_n in R^k # direction channels: truth, patience, charity, courage, hope, humility...

Psi_n := H_n * S_n

Linear-retention + control inputs (copy-safe).

H_{n+1} = R_H @ H_n + B_H @ u_H(n) + N_H @ xi_n + G_H @ Psi_n

S_{n+1} = R_S @ S_n + B_S @ u_S(n) + N_S @ xi_n

- R_* diagonal (0.90–0.98): forgetting/retention.
- u_H: habit supports (sleep schedule, breath protocol, diet timing, speech drills).
- u_S: practice bundles (audit truth, examen, reconciliation, almsgiving, service).
- G_H @ Psi_n: lived agency stabilizes capacity (truth reduces stress → sleep/breath regularity).
- Virtue floors for theological coords enforced after 5.4 vertical lift.

Practice bundles (examples, target → expected readouts).

- **Truth bundle:** S_truth ↑ → rho < 1, rework variance ↓, modest E gain.
- **Patience bundle:** S_patience ↑ → start/stop losses ↓, coherence in L ↑.
- **Charity bundle:** S_charity ↑ → local O densification, prosocial spillovers.
- Combine with u_H to avoid energy regressions (e.g., add sleep hygiene to intense truth protocols).

5.4 Vertical couplings: S-lift along theological coordinates; O_max expansion via Gammā; consent dynamics for C

Orthogonality. God is unmodeled. Only *received grace* appears as a consent-gated moderator:

C_n in [0,1] # consent/attention (voluntary)

Gamma_n := phi(C_n) in [0,1] # increasing, saturating (e.g., identity or logistic)

Gamma_bar_n := EWMA(Gamma_0...Gamma_n; alpha_g)

(a) Consent dynamics (creaturely).

C_{n+1} = R_C * C_n + B_C * u_C(n) + noise, 0 < R_C < 1

- u_C: voluntary acts (silence, prayer/worship, examen; or secular analogs).
- **Ethics:** Never coerce or "optimize" consent; participation must be free.

(b) S-lift (small, consent-gated) along theological coordinates.

S_{n+1} = ... + eta_S * Gamma_n * dir_S_theological

S_{j,n+1} = max(0, S_{j,n+1}) for j in {faith, hope, charity}

- eta_S small (e.g., 0.01–0.05 per tick); dir_S_theological unit vector.
- This lift does **not** alter E or L directly.

(c) Budget expansion (capacity, not edges).

O_max'(Psi_{0:n}, Gamma_bar_n)

= O_max(Psi_{0:n}) * (1 + delta * Gamma_bar_n), with 0 <= delta << 1

Constraint: ||O_n||_1 <= O_max'

- delta predeclared and small (e.g., 0.05–0.15).
- No bypass of sparsity or the standard O-dynamics; edges still shaped by K_phys and K_agency.

Non-permitted couplings (V-Conservativity).

No direct: Gamma -> E, Gamma -> L, Gamma -> F,Z,G

No edge injection: Gamma cannot create O edges or override SoftThresh/Proj_budget.

Estimation posture (V-Humility).

- Fit base models with Gamma toggled OFF; primary endpoints must not depend on Gamma.
- If a moderated model with shrinkage prior on Gamma materially improves preregistered out-of-sample fit **and** passes negative controls (sham-consent), report Gamma as a small moderator—separately, never as a "God score."

5.5 Parameterizations and stability regimes (with Gamma, C)

Nominal parameter set (suggested ranges).

h : 1.0 # tick cadence (day/session)

kappa : 1.0 # noise multiplier

lambda : 0.12 # O learning/forgetting

alpha : 0.6 # K_phys weight

beta : 0.4 # K_agency weight

mu : 0.05 # O sparsity threshold

R_H, R_S : 0.93–0.97 # retention (diag)

B_H, B_S : 0.02–0.10 # input gains

N_H, N_S : 0.01–0.03 # noise gains

G_H : 0.02–0.05 # Psi→H coupling

Vertical

R_C : 0.90–0.97 # consent retention

B_C : 0.02–0.06 # consent input gain

phi(x) : min(1, x) or logistic(k*x)

eta_S : 0.01–0.05 # S-lift gain (small)

delta : 0.05–0.15 # O_max expansion cap (small)

alpha_g : 0.2–0.4 # EWMA smoothing for Gamma_bar

Stability heuristics.

1. **Bounded agency:** Ensure ||R_H||, ||R_S|| < 1 (operator norm). Clip H,S to safe bands; enforce virtue floors for theological coords.
2. **Non-expansive communion:** With 0<lambda<1, SoftThresh, and Proj_budget, O is L1-nonexpansive; adjust mu upward if spurious edges appear.
3. **Energy compatibility:** Track E_per_output. If truth bundles raise stress (and E) transiently, add u_H supports (sleep/breath) or slow the audit cadence.

4. **Vertical moderation kept small:** Choose eta_S, delta s.t. removing Gamma does not flip primary results. If it does, set Gamma := 0 and improve practice/audit fidelity.
5. **Anneal exploration:** Start with higher kappa, reduce after baseline; increase mu when networks over-densify.
6. **Avoid clique capture:** Penalize high-degree nodes or add degree caps when beta is large; monitor local ||O||_1 / O_max' and nnz(O)/N^2.

Failure modes & mitigations.

- **No variance-after-truth (rho ~ 1):** Audit MVI score low; upgrade immutable ledgers/speech tags; extend baseline; add negative controls.
- **Energy regressions:** Overzealous S without H supports; add sleep/breath or reduce B_S.
- **Runaway O (cliques):** Raise mu, lower beta, strengthen budget (lower delta), add diversity penalty.
- **"Vertical dependency":** Effects vanish unless Gamma is large; disable vertical paths, re-specify A_S, improve fidelity—do not attribute to Gamma.
- **Consent pressure risk:** If uptake of u_C lowers fairness or autonomy scores, pause vertical elements; vertical is optional, never required.

Reference tick (pseudocode, with vertical toggles).

```
def tick(X, P, U, rng, use_gamma=False):
    F,Z,L,E,G,O,H,S,Psi = X
    xi = rng.normal(size=P.noise_dim) * (P.kappa * P.h) ** 0.5

    # Backbone
    F = D_F(F,Z) + S_F(F,Z,xi,P.h)
    Z = D_Z(Z,F) + S_Z(Z,F,xi,P.h)
    L = D_L(L,F,Z) + S_L(L,xi,P.h)
    E = D_E(E,Z) + S_E(E,xi,P.h)
    G = D_G(G,F,E) + S_G(G,xi,P.h)
```

```
# O-dynamics
O_tilde = (1-P.lambda_) * O + P.lambda_ * sig(P.alpha*K_phys(F,Z)
        + P.beta*K_agency(Psi)) + S_O(O,xi,P.h)

# Consent & vertical (optional)
C = None
Gamma = 0.0
if use_gamma:
   C = P.RC * (getattr(P,'C',0.0)) + P.BC * U.uC + noise_small()
   Gamma = phi(C)
   P.C = C  # store
   O_max_prime = O_max(P.Psi_hist) * (1 + P.delta * EWMA(Gamma, P.alpha_g))
else:
   O_max_prime = O_max(P.Psi_hist)

O = Proj_budget( SoftThresh(O_tilde, P.mu), O_max_prime )

# Agency
H = P.RH @ H + P.BH @ U.uH + P.NH @ xi + P.GH @ Psi
S = P.RS @ S + P.BS @ U.uS + P.NS @ xi
if use_gamma:
   S = S + P.etaS * Gamma * dir_S_theological
S = virtue_floor(S)  # theological coords

Psi = H * S
return (F,Z,L,E,G,O,H,S,Psi)
```

What to log every tick.

- Primary: E_per_output, Var_loss, rho (if audited), local/global ||O||_1, sparsity nnz(O)/N^2, synchrony indices, severity distribution.
- Agency: S_truth, S_patience, S_charity, H supports, Psi magnitude.
- Vertical (if used): C exposure, Gamma (bounded/shrunk), O_max' vs O_max.
- Integrity: audit score, data completeness, mix controls.

Bottom line.
5.1 gives the lawful backbone with bounded novelty; 5.2 makes communion grow where practice and agency warrant, under a hard budget; 5.3 turns practices and habits into actionable controls; 5.4 introduces a strictly conservative, consent-gated vertical moderation that can *sustain* virtue and communion without touching the ledgers; 5.5 pins down tunings, stability, and failure handling so experiments remain interpretable and falsifiable.

6. Operational Ethics (STEOM 5)

6.1 The Alexandrian Rule as measurement doctrine

Doctrine. *Power consents to be measured before it moves.* Any locus of authority binds itself to auditable weights, ledgers, calendrics, and speech.

Minimum Viable Implementation (MVI).

- **Weights:** Calibrate instruments; publish uncertainty; lock configurations; keep an immutable change log.
- **Ledgers:** Double-entry with hashed diffs; role-separated write access; rotating countersign.
- **Calendrics:** ISO-8601 timestamps; close dates are one-way doors; no retro-edits without public errata.
- **Speech:** Tag every claim at the point of utterance as *observation*, *inference*, or *intention*; corrections are first-class events.

Primary KPI (copy-safe).

Q_n := loss per tick (e.g., rework+scrap+delay) / Output_n

rho := Var(Q_post) / Var(Q_pre) with pre/post windows preregistered

Target: rho < 1 (e.g., rho <= 0.8 with CI not crossing 1)

Governance primitives.

- **No measure → no move:** high-impact actions are gated by passing audits.
- **One-way doors:** evidence thresholds rise with irreversibility.
- **Disclosure:** publish rho, audit checklist scores, and raw schemas; accept adversarial review.

Vertical note. Γ (received grace) has **no place** in audit logic. Audits succeed or fail on their own merits. Never invoke Γ to excuse weak ledgers or vague speech.

6.2 Truth shrinks variance; love re-weights edges

Truth ⇒ variance. Honest weights/ledgers/speech reduce hidden-state uncertainty and misallocation noise in the next tick.

Let V_n := Var(Q_n | history_n).

A6 predicts E[V_{n+1} - V_n | truthful_audit_n] < 0.

Mechanism. Audits tighten posterior beliefs, trimming exploratory thrash in D_* and lowering the novelty gains needed in S_*.

Love ⇒ edges (within budget). Practiced charity and patient attention (S → Ψ) selectively densify O via the agency kernel:

O_{n+1} = Proj_budget((1 - lambda) O_n

+ lambda * sigma(alpha*K_phys + beta*K_agency(Psi_n))

+ S_O)

Observable footprints.

- **Synchrony:** ↑ HRV phase-locking, smoother turn-taking, fewer overlaps.
- **Prosociality:** ↑ help/correction events, faster hand-offs in 1–2 hop neighborhoods.
- **Softer landings:** conditional severity shifts toward minor outcomes while energy/info ledgers hold.

Ethical biconditional.

- Truth without love concentrates power but leaves O thin.
- Love without truth over-couples the wrong places (spurious edges), improving neither variance nor mercy.
  Operational ethics therefore binds *both*: audits first, then charity targets attention where help is needed—always under budget.

Vertical note. If Γ is modeled at all, it may at most **expand the ceiling** (O_max') or gently lift theological coordinates in S. It does **not** re-weight edges directly and can never replace practiced charity.

6.3 Grace-as-efficiency within ledgers; Γ never a substitute for audits

Claim (A5, conservative form). As selected virtue coordinates in S rise—especially *truth*, *patience*, *charity*—the energy-per-output metric decreases in expectation at fixed output/mix; waste variance tightens.

Metrics.

- **Effort:** E_n := (W_out + Q_loss) / Output_n.
- **Waste:** components of Q_loss with units (rework hrs, scrap ppm, delay penalties).
- **Coherence:** C_n via routing efficiency or inverse spectral flatness in L.

Protocol (copy-safe).

1. **Baseline** 3–5 ticks: lock Output, mix, and instrumentation; compute E, Var(Q), C.
2. **Introduce one bundle** (e.g., Truth; or Patience + H-supports for sleep/breath).
3. **Hold mix constant**; keep audit fidelity ≥ threshold.
4. **Readouts** after 2–4 ticks:

DeltaE := E_baseline - E_post

rho := Var(Q_post) / Var(Q_pre)

DeltaC := C_post - C_baseline

5. **Stop rules:** if E worsens >10% or rho > 1.1 across 2 ticks, halt and revise mapping (B_S/B_H) rather than pressing the bundle.

Fairness & mercy checks. Gains must not be extracted by shifting burden to vulnerable subgroups. Require:

- Non-inferiority of severity distributions by subgroup,
- Parity (or justified difference) in E_per_output by subgroup,
- Transparent accounting of measurement overhead (Landauer-aligned).

Γ policy (V-Conservativity & V-Humility).

- **Never** posit Γ to "explain" missing audit effects; **never** count Γ toward success metrics.
- If modeled, Γ is a small, shrinkage-regularized moderator that can *sustain* virtue growth (S theological) or *modestly* raise O_max'—and only that.
- Publish vertical terms **separately** and be ready to report Γ ≈ 0 without embarrassment.

Failure interpretation matrix.

- **rho<1, ΔE≈0:** audits reduced variance; add H-supports to translate to energy gain.
- **ΔE>0, rho≈1:** energy gains from capacity alone; upgrade truth practices to bind power.
- **Synchrony↑, no efficiency:** O grew without targeting bottlenecks; retarget charity toward blocking edges.
- **No effects with good fidelity:** revisit parameterization (beta, mu, lambda), check mix confounds and instrument drift, and rerun with negative controls.

Bottom line. Operational ethics in STEOM 5 is a *discipline*: bind power to measurement, aim attention by love, and count all costs in open ledgers. Any talk of grace remains conservative—consent-gated, secondary, and never a bypass for the hard, honest work of audits and practice.

7. Testable Predictions (T1–T7) and Vertical Addenda — STEOM 5

7.1 Fractal spectra shifts under contemplative practice (T1)

Hypothesis. Regular contemplative practice (patience+truth bundles with H-supports) increases cross-scale coherence: PSD slope αin task-relevant signals steepens (Δα>0) vs controls.
Signals. HRV RR intervals, keystroke intervals, cursor micro-motions, pause durations, error inter-arrival times.
Design. 2-arm RCT, 4–6 weeks, n≈80. Treatment: daily 20 min breath cadence + 5 min examen. Control: silent reading.
Metrics. Fixed bandpass PSD slope (log–log), DFA exponent, multiscale entropy; prereg

estimators.
Primary test. $\Delta\alpha_{\text{treat}} - \Delta\alpha_{\text{ctrl}} > 0$.
Falsifier. Null/opposite shifts under good adherence and fixed pipelines.

7.2 Variance-after-truth in audited organizations (T2)

Hypothesis. High-fidelity audits shrink next-tick loss variance: $\rho = \frac{\text{Var}(Q_{\text{post}})}{\text{Var}(Q_{\text{pre}})} < 1$.
Contexts. Software (defects/merged change), clinics (no-show deltas), plants (scrap ppm).
Design. Stepped-wedge across ≥6 teams; audits = calibrated weights, immutable ledgers, timestamps, speech tags, rotating countersign.
Primary test. Team-level $\rho \leq 0.8$(CI<1) in ≥4/6 teams; pooled meta $\rho < 1$.
Falsifier. $\rho \geq 1$across replications with audit MVI scores high.

7.3 Communion contagion in social/physio synchrony (T3)

Hypothesis. Seeding high-Ψ agents increases 1–2 hop synchrony and prosocial behavior vs placebo seeds.
Design. 20 groups (6–10 ppl). Randomize 10 groups to 2 trained seeds (truth+patience bundles + sleep/breath). Observe 4–6 weeks.
Measures. HRV PLV, turn-taking smoothness, overlap rates, help/correction counts; inferred local $\| O \|_1$.
Primary test. Seeded > control on synchrony/prosocial deltas; strongest within 2 hops.
Falsifier. No changes despite verified seed compliance.

7.4 Retimed hazards and softer landings (T4)

Hypothesis. As $\| O \|_1$grows within budget, conditional severity softens: $P(\text{severe} \mid \text{event}) \downarrow$, $P(\text{minor}) \uparrow$, ledgers hold.
Contexts. Hospitals (falls/near-miss), factories (safety), devops (incidents), classrooms (discipline).
Design. Interrupted time-series with matched controls; fixed severity rubric.
Primary test. Ordinal model shows negative shift in severity odds for treated units; energy/info ledgers unchanged.
Falsifier. Softening requires hidden energy/data subsidies or does not occur with $\| O \|_1$growth.

7.5 Energetic cost reductions at fixed output (T5)

Hypothesis. Raising $S_{\text{truth}}, S_{\text{patience}}, S_{\text{charity}}$lowers $E = \frac{W+Q}{\text{Output}}$ holding mix constant.
Design. ABAB within 6–10 units over 8–12 weeks; one bundle at a time; fixed throughput band.
Primary test. Within-unit level change shows 3–8% Edrop in A (on) vs B (off); replication across

units.
Falsifier. No Eimprovement with strong audit fidelity and stable mix.

7.6 Effective $h$inference from innovation sizes (T6)

Hypothesis. Stable processes exhibit a lower innovation-size cutoff $\tau \propto \sqrt{h}$.
Design. 2-week process freeze; fit deterministic backbone, collect residual innovations; fit truncated mixture to estimate τ.
Primary test. Nonzero, stable τthat decreases when tasks are retimed (operationally smaller h).
Falsifier. No identifiable cutoff or instability under freezes.

7.7 Portability across sites and cultures (T7)

Hypothesis. T1–T6 replicate across diverse sites when doctrine fidelity and measurement are held constant.
Design. 4–12 sites (clinic/plant/software/school); shared kits and prereg.
Primary tests. $\rho < 1$(T2) in ≥3 sites; at least one of T1/T3/T5 positive in ≥3 sites; no ledger violations.
Falsifier. Effects collapse outside a single culture or require ad-hoc metrics.

7.V Vertical addenda — Γ as secondary, consent-gated moderation (pre-specified, small)

Principle. All primary tests (T1–T7) must succeed **with Γ toggled OFF**. Vertical effects, if any, are small, *pre-specified*, and depend on voluntary consent C. Report them separately; never count them toward success criteria.

Pre-specification (copy-safe).

- **Form.** Add one moderated contrast per T:
  Effect $= \beta_0 + \beta_1 \cdot$ Intervention $+ \theta \cdot \bar{\Gamma} + \epsilon$,
  with $0 \le \bar{\Gamma} \le 1$(EWMA of Γ), **shrinkage prior** on θ(centered near 0).
- **Bounds.** Cap | θ |ex ante (e.g., “vertical moderation cannot exceed 25% of the unmoderated effect size”).

How Γ may manifest (allowed patterns).

- **T1 (spectra):** Slightly larger Δα in participants with higher $\bar{\Gamma}$(θ>0 small).
- **T2 (variance):** Extra few-percent reduction in ρ beyond audit effect (θ<0 small).
- **T3 (contagion):** Mildly stronger synchrony uplift in 1–2 hops when seeds’ $\bar{\Gamma}$is higher (θ>0 small).

- **T4 (softening):** Additional shift toward minor severity within ledger bounds (θ>0 small).
- **T5 (energy):** Modest incremental ΔE beyond virtue-driven gains at fixed output (θ>0 small).
- **T6 (h):** No direct moderation pre-specified (vertical does not retime tasks).
- **T7 (portability):** Moderation may vary by site but remains bounded and non-essential.

Negative controls (must pass).

- **Sham-consent:** Ritualized "consent" without practices should not yield θ≠0.
- **Calendar shuffle:** Randomly permuting C timestamps removes moderation.
- **Audit-failure screen:** If audit MVI < threshold, forbid attributing any gains to Γ.

Reporting rules.

- Present base effects first (with Γ off).
- If moderated model improves prereg out-of-sample fit and passes all negative controls, report θ with CI, the cap used, and a plain-language note: "vertical moderation small, consent-gated, non-essential."
- Publish null vertical results (θ≈0) without reinterpretation.

Failure handling.

- If primary effects require large θ to fit, declare **model/instrumentation fault**, disable Γ, and repair audits/practices or O-inference.
- If vertical moderation appears only where fairness/consent metrics are weak, treat as **ethical failure** and halt vertical elements.

Bottom line. T1–T7 make STEOM 5 testable on creaturely grounds—audits, practices, agency, and budgeted communion. The vertical addendum allows a narrowly defined, consent-gated moderation that may *sustain* virtue and communion a bit further—never replace the ledgers, never rescue weak fidelity, never become the point.

8. Methods: From Simulation to Field Trials (STEOM 5)

8.1 Reference simulator (with Γ, C stubs off by default)

Purpose. Provide a minimal, reproducible engine that implements the tick law, enforces budgets, emits STEOM 5 readouts, and keeps all vertical couplings disabled unless explicitly toggled.

State (copy-safe).

X = {F, Z, L, E, G, O, H, S, Psi}; Psi := H * S # Hadamard

Core tick (pseudocode).

```
def tick(X, P, U, rng, use_gamma=False):
   F,Z,L,E,G,O,H,S,Psi = X
   xi = rng.normal(size=P.noise_dim) * (P.kappa * P.h) ** 0.5

   # 1) Backbone (deterministic + bounded novelty)
   F = D_F(F,Z) + S_F(F,Z,xi,P.h)
   Z = D_Z(Z,F) + S_Z(Z,F,xi,P.h)
   L = D_L(L,F,Z) + S_L(L,xi,P.h)
   E = D_E(E,Z) + S_E(E,xi,P.h)
   G = D_G(G,F,E) + S_G(G,xi,P.h)

   # 2) Communion dynamics (budgeted, sparse)
   O_tilde = (1-P.lambda_) * O \
         + P.lambda_ * sig(P.alpha*K_phys(F,Z) + P.beta*K_agency(Psi)) \
         + S_O(O,xi,P.h)

   Omax = O_max(P.Psi_hist)
   if use_gamma:
      C = P.RC * getattr(P, "C", 0.0) + P.BC * U.uC + noise_small()
```

```
    P.C = clip01(C)
    Gamma = phi(P.C)                # received grace (latent)
    Omax = Omax * (1 + P.delta * EWMA(Gamma, P.alpha_g))
  O = Proj_budget( SoftThresh(O_tilde, P.mu), Omax )

  # 3) Agency channels
  H = P.RH @ H + P.BH @ U.uH + P.NH @ xi + P.GH @ Psi
  S = P.RS @ S + P.BS @ U.uS + P.NS @ xi
  if use_gamma:
    S = S + P.etaS * Gamma * dir_S_theological
  S = virtue_floor(S)
  Psi = H * S

  return (F,Z,L,E,G,O,H,S,Psi)
```

Default config (vertical off).

use_gamma = False

h=1.0, kappa=1.0, lambda=0.12, alpha=0.6, beta=0.4, mu=0.05

RH=RS=diag(0.95), BH=0.05, BS=0.07, NH=NS=0.02, GH=0.03

O_sat, gamma (budget), w_decay tuned per domain

Logged metrics (per tick).

- E_per_output, Var_loss (sliding window), rho if an audit gate fired.
- ||O||_1 (global), local ||O||_1 around each node, nnz(O)/N^2.
- Synchrony indices (simulated PLV/turn-taking), severity histogram.
- S_truth/patience/charity, H supports, Psi magnitude.
- If use_gamma=True: C exposure, smoothed Gamma (bounded, shrinkage), Omax vs base.

Reproducibility. Fixed seed, config manifest (code hash, config hash, data hash), CSV+parquet outputs, and a one-click notebook for figures.

8.2 Synthetic worlds: ablations (β=0, μ→0, Γ on/off) and sensitivity

Worlds. N=40–120 agents with role labels, proximity graph for K_phys, and task queues affecting Z,L,E. Seeds and practice schedules controllable.

Ablations (isolate mechanisms).

1. **β=0** (no agency-to-O): physical affinity only. Expect little synchrony/prosocial change from practice bundles.
2. **μ→0** (no sparsity): dense O. Expect transient synchrony with instability; efficiency often flat or worse.
3. **No audits** (A6 off): truth bundles should not reduce Var_loss (ρ≈1).
4. **A_H only / A_S only / both**: identify whether energy gains need both capacity and virtue.
5. **Γ toggles**: use_gamma=False vs True. With True, allow only small etaS, delta. Primary effects must persist with False.

Sensitivity sweeps (one-factor-at-a-time).

- h: {0.25, 0.5, 1.0, 2.0} → verify innovation cutoff scales ∝ sqrt(h).
- lambda: {0.05, 0.12, 0.30} → learning smoothness vs responsiveness.
- alpha/beta grid → domains where structure vs agency dominates O.
- mu: [0.01, 0.10] → selectivity and clique suppression.
- Retentions RH,RS, gains BH,BS → habit vs practice memory.
- Vertical: etaS in [0, 0.05], delta in [0, 0.15] → confirm non-essential moderation.

Stability criteria (declare "good region").

- rho < 1 following audits; E_per_output non-increasing under targeted bundles; ||O||_1 / Omax < 0.9 on average; no node degree explosions; severity softening with budget intact.

Failure triage.

- If effects depend on large etaS/delta, set use_gamma=False and retune B_S, mu, beta.
- If nnz(O) blows up, raise mu, reduce beta, or tighten budget (lower O_sat or delta).
- If E rises, add u_H supports or reduce bundle intensity.

8.3 Prereg blueprints (org, lab, community) with vertical negative controls

Shared prereg spine.

- Units, inclusion/exclusion, tick definition, windows, endpoints, estimator, power, missingness, outlier rules.
- Audit MVI checklist with threshold for counting an "audited" tick.
- **Two-model plan:** Base (no Γ) and Moderated (Γ with shrinkage). Base determines success/failure; Moderated is exploratory unless explicitly powered.

Organization (T2, T5).

- **Design.** Stepped-wedge audits across 6–10 teams; ABAB for practice bundles with throughput band fixed.
- **Primary endpoints.** rho < 1 (CI<1); DeltaE > 0 at fixed output/mix.
- **Vertical controls.** If consent activities are offered, prereg a **sham-consent** control (time-matched, content-neutral). Any Γ effect must vanish under sham and must not rescue failed audits.

Lab (T1, T6).

- **Design.** 2-arm RCT (practice vs reading) and a process-freeze block for h-inference.
- **Preprocessing lockfile.** HRV cleaning, bandpass, PSD slope estimator.
- **Vertical controls.** Randomly permute consent timestamps in a placebo analysis; moderated model must lose its lift under this permutation.

Community (T3, T4).

- **Design.** Seed 2 high-Ψ participants in half of groups; 4–6 weeks.
- **Endpoints.** PLV change, turn-taking smoothness, help/correction counts, severity ladder shift.

- **Vertical controls.** Consent is voluntary; no grading or privileges tied to it. If a site includes consent practices, add a matched group with **consent exposure only** but no practice bundles; large effects here falsify vertical humility.

Reporting.

- Post prereg, code, raw/tidy data, audit scores, rho with CI, E_per_output, local/global ||O||_1, synchrony deltas, severity shift, h-hat.
- Publish nulls. Vertical terms (if any) in a separate table with caps and negative-control checks.

8.4 Measurement kits: HRV, variance ledgers, synchrony indices; Γ reporting rules

A. HRV & synchrony kit.

- **Hardware.** Chest straps or medical-grade RR at ≥250 Hz; synchronized start.
- **Preprocessing.** Artifact removal, ectopic filtering, interpolation; document method.
- **Single-stream metrics.** RMSSD, SDNN, LF/HF band powers, DFA exponent, multiscale entropy.
- **Synchrony metrics.** PLV on bandpassed RR, magnitude-squared coherence in LF/HF, cross-correlation near lag 0, speech turn-taking smoothness, overlap rate.
- **Outputs.** Per-session means, deltas, mixed-effects models; cluster-robust SEs for group designs.

B. Variance ledger kit (truth → variance).

- **Define** Q_n (loss/Output) with domain units; freeze Output and mix controls.
- **Compute** sliding-window Var_loss, pre/post windows Var_pre, Var_post, and rho = Var_post / Var_pre.
- **Integrity.** Immutable ledgers, data completeness %, audit MVI score; analyst blinding until code freeze.

C. O-inference & prosocial kit.

- **Behavioral.** Help/correction counts, on-time handoffs, two-hop spillover tallies.
- **Graphical inference.** Predeclared state-space/graph model from synchrony/prosocial features → local ||O||_1 deltas; sparsity and budget respected.

- **Ethics.** Consent, opt-out, minimal collection, no covert recording.

D. h-estimation kit.

- **Freeze** process 10–14 ticks; fit deterministic backbone; collect residual innovations.
- **Fit** truncated-normal or mixture with lower cutoff tau; set h_hat := (tau/c)**2 using simulation-calibrated c.
- **Manipulation check.** Retiming tasks (smaller h) must reduce tau.

E. Γ reporting rules (V-Humility in practice).

- **Default:** Do not fit Γ. Report base results.
- **If preregistered moderation is ON:**
    - Use shrinkage priors centered at 0; cap |θ| ex ante (e.g., ≤25% of base effect).
    - Report consent exposure C descriptives; never compute a "God score."
    - Run negative controls: sham-consent, timestamp permutation, audit-failure screen.
    - Only report θ if (i) out-of-sample fit improves, (ii) negatives pass, (iii) fairness and consent integrity are intact.
    - Present vertical table separately with clear language: "small, consent-gated, non-essential moderation."
- **Never** invoke Γ to explain away weak audits/practices. If primary effects depend on Γ, declare model or fidelity failure and rerun.

F. Fairness & harm monitors (shared).

- Severity non-inferiority by subgroup; E_per_output parity checks; measurement overhead accounting (include in E).
- Stop if harm thresholds trip (e.g., E worsens >10% over 2 ticks, severity odds worsen, audit score falls).

Bottom line. 8.1 gives a runnable scaffold (vertical off by default). 8.2 maps which levers actually drive the signatures and which merely decorate them. 8.3 turns everything into preregistered, falsifiable studies with explicit vertical negative controls. 8.4 ships the

instruments and the rules that keep results honest—especially about Γ, which, if present at all, remains small, consent-gated, and never the reason the program works.

9. Case Sketches and Early Use-Cases (STEOM 5)

9.1 Audit-then-operate (variance reduction in teams)

Setting. Two software teams (A,B), 12 devs each, same product area, 6 sprints (1 sprint = 2 weeks).

Baseline (Sprints 1–2). Freeze tooling; define Output := merged_changes. Loss Q_n := rework_hours + defect_hours + delay_hours. Compute Var_loss over a 2-sprint window. No practice bundles yet.

Intervention (Sprints 3–4) — Truth bundle to Team A only.

- Ledgers: immutable diffs, timestamp locks, role-separated writes.
- Speech tags at stand-up: each claim labeled observation / inference / intention.
- Weekly rotating peer audit; **no retro-edits** without public errata.
  Team B = control (stepped-wedge later).

Operate (Sprints 5–6). Keep A on; B stays off until study end.

Primary KPI (copy-safe).

rho_A := Var(Q_post(A))/Var(Q_pre(A)) Target: rho_A <= 0.8 (CI < 1)

rho_B ≈ 1 (no systematic change)

Secondary. E_per_output := (W_out + Q_loss)/Output ↓ 3–6% for A; post-mortem “clear root cause” flag rate ↑.

Guardrails. Mix normalization; audit MVI score ≥ 7/10 to count a sprint as “audited.” Publish schemas and outlier policy.

Failure handling.

- rho_A ~ 1: raise audit quality (weights, timestamps), extend windows, or add speech-tag coaching.
- E flat with rho_A < 1: add H-supports (sleep/breath) to convert less thrash into lower energy (9.2).

- Improvements in B too: attribute to external driver; switch to stepped-wedge or add Team C.

Variants.

- Helpdesk: Output = resolved_tickets, Q = reopen + escalation.
- Clinic: Output = arrived_appointments, Q = no-show_delta + rescheduling_cost.
- Plant: Output = conforming_units, Q = scrap_ppm + rework_min.

9.2 Breath-prayer micro-protocols and task energy

Setting. Documentation pod (10 writers), daily 2-hour release-note block.

Baseline (Week 1).

- Output := approved_words/hour (AWPH).
- E := staff_minutes / 1k AW.
- HRV (RMSSD), errors per 1k AW, keystroke-interval PSD slope alpha.

Intervention (Weeks 2–3) — Patience bundle + H-supports.

- Micro-protocol per hour: 2 min 4–7–8 cadence + one sentence breath-prayer ("receive / return").
- "One-edit rule": finish a paragraph before editing.
- Fixed break at 55 min (5-min walk), hydration cue.

Readouts.

DeltaE := E_baseline - E_post (target 5–10%)

Error_rate ↓ 10–20%

RMSSD ↑ ; PSD slope alpha ↑ (small–moderate)

Compliance. App timer logs; HRV coverage ≥ 70% sessions.

Stops & tweaks. If AWPH drops > 5%, shorten cadence to 1 min; maintain "one-edit rule." If RMSSD ↑ but DeltaE ~ 0, add Truth micro-audit (daily checklist, tag claims) to unlock variance shrinkage.

Why STEOM. H stabilizes, S_patience rises → Psi smoother; fewer start/stop losses; L routing thrash ↓; E falls at fixed output.

Variants.

- Call center: 30–45 sec cadence between calls; Output = resolved_calls; E = avg_handle_time/resolution.
- Code review: 60–90 sec cadence pre-review; "read once before comment" rule.

9.3 Prosocial seeding and O-densification in classrooms

Setting. Four seminar sections, 9 students each, 6 weeks. Randomly assign two sections to receive **two seeds** (coached weekly on truth, patience, concrete help + sleep/breath supports).

Seed coaching (30 min/week).

- Truth: "steel-man once before disagreeing."
- Patience: 2× 30-sec breath cadence during tense moments.
- Charity: 1 concrete help per class (note-sharing, prompt a quiet peer).
- Speech: no interruptions; ask clarifying questions first.

Instrumentation.

- HRV straps (opt-in); PLV in LF band.
- Turn-taking smoothness; overlap/interrupt count.
- Prosocial tally (help/correction rubric).
- Inferred local ||O||_1 from a predeclared state-space/graph model.

Primary endpoints.

delta_synchrony := PLV_post - PLV_pre (section-level)

Prosocial_delta := helps+corrections_post - baseline

Local O_L1 delta for non-seed students within 2 hops

Targets. Seeded > control (d ~ 0.5); strongest for 1–2 hop neighbors; instructor intervention time not ↑.

Guardrails & ethics. Consent; opt-out; no grading tied to signals; cap same-pair interactions to avoid clique capture; degree caps in O inference.

Failure handling.

- Synchrony ↑, prosocial flat: add "help quota" clarity.
- Prosocial ↑, synchrony flat: increase sampling or add speech-only synchrony (turn-taking).
- No changes: choose seeds with higher social centrality; increase coaching cadence.

Extensions. Wards/labs: pair senior seeds with juniors for handoffs; track error-catching and on-time handoffs.

9.4 Vertical note: consent practices (C) as optional, non-coercive adjuncts

Principle (V-Orthogonality & V-Humility). Consent/attention C is **voluntary** and **non-optimized**. Received grace Gamma := phi(C) is **never** counted toward success, never used to rescue weak audits or sloppy practice fidelity.

How to offer C ethically (optional tracks).

- **A-track (secular):** quiet reflection, gratitude journaling, mindful breathing in silence corners; no data collected unless volunteered.
- **B-track (faith-forward):** prayer, examen, worship times provided; participation free and private.

Protocol additions (only if offered).

Log C exposure descriptively (yes/no or minutes), never as a score.

Pre-specify moderated models with shrinkage priors on theta (centered at 0).

Run negative controls: sham-consent, timestamp permutation.

Cap |theta| ex ante (e.g., <= 25% of base effect).

Reporting. Present base results first (Γ off). If a moderated fit passes all negative controls and improves prereg out-of-sample performance, report θ separately as "small, consent-gated, non-essential."

Red lines.

- No incentives, penalties, or nudges tied to C.
- No "God metric."
- If vertical exposure correlates with fairness or burden shifts, **pause** vertical elements and remediate.

Practical examples.

- 9.1: Provide a quiet room before stand-up; no logging of who attends; audits proceed regardless.
- 9.2: Breath-prayer wording optional; secular breathing script available; either choice is acceptable.
- 9.3: Seeds may choose to engage in private consent practices; coaching and evaluation do not depend on it.

Bottom line. Use-cases 9.1–9.3 stand on audits, practices, and agency within budgeted communion. The vertical remains an **adjunct**: freely chosen, small if present, and never a substitute for the measured work.

10. Related Work (Bridging Traditions)

10.1 Theology of virtues and practices; grace and consent

Virtue as habituated power. Classical accounts (e.g., Aristotle as received by Aquinas) treat *habitus* as stable readiness for right action. STEOM's S (virtue coordinates) and H (capacity) echo this: practiced dispositions lower internal "friction," aligning with our grace-as-efficiency claim (lower E at fixed output).

Rule and examen as proto-instrumentation. Monastic and communal rules (Basil, Benedict) operationalize transformation via schedules, countersignature, chapter correction—early forms of our audit doctrine (weights, ledgers, timestamps, speech). Ignatian examen adds daily review and explicit intention tagging—precursors to speech tags (*observation/inference/intention*).

Grace and consent. Patristic and scholastic strands hold synergism without compulsion: grace invites, consent cooperates. STEOM's vertical layer adopts this stance mathematically: God unmodeled; only *received grace* $\Gamma = \phi(C)$ enters, bounded by freely given consent C, never as a force on ledgers.

Works of mercy as system repair. Practical theologies (e.g., almsgiving, reconciliation, hospitality) predict thicker local bonds. Our O-dynamics formalize this as selective edge densification under a finite budget, with mercy tested by *severity softening without hidden subsidies*.

10.2 Complexity/fractals and nonlocal coupling

Fractal realism. Physiology and behavior exhibit 1/f-like spectra and multiscale regularities (heartbeat, gait, keystrokes). STEOM packages this in F with MDL/codelength proxies and DFA/PSD tests (T1). Cross-scale transfer justifies why practices tuned at one timescale can aid neighbors.

Sparse, budgeted networks. Network science and graphical models show that coordination depends more on *where* edges are than on density. Our oneness tensor O uses sparsity (soft-thresholding) plus an L1 budget O_max' to emulate attention scarcity and prevent clique blow-ups.

Synchronization theory. From coupled oscillators to speech turn-taking, small edge weight changes reshape group timing. STEOM uses agency-weighted kernels K_agency(Psi) to channel practiced truth/patience/charity into measurable synchrony gains and prosocial spillovers (T3).

Criticality and soft tails. Self-organized criticality and buffer effects explain severity retiming: added *communion mass* in the right places shifts events toward “soft landings” (T4) without changing the input energy—mirroring our “mercy within ledgers.”

Information bounds. MDL and data-processing ideas motivate A6: truthful audits increase state–measurement information, which should reduce allocation variance next tick ($\rho<1$). STEOM keeps these gains conservative by charging measurement overhead inside E.

10.3 Ethics-as-engineering: measurement, accountability, mercy

Quality and safety lineages. Shewhart/Deming control, ISO-style process discipline, and modern safety science (near-miss analysis, Just Culture) ground A6: bind action to transparent, auditable measures and expect variance shrinkage. STEOM borrows their bones—immutable ledgers, timestamp discipline, peer countersign—then adds speech tagging to curb causal confusions.

Against Goodhart. Social systems game targets. STEOM therefore elevates *second-order* readouts (variance-after-truth ρ, severity distributions) and guards with negative controls (sham audits, calendar shuffles). Measurement is a covenant, not theater.

Mercy within constraints. Disaster ethics and triage remind us: distribution of harm matters. A7 encodes this technically—softening severity must occur *inside* energy/information ledgers, never by pushing cost onto the unseen. Landauer-style accounting requires we debit measurement overhead before declaring "grace-as-efficiency."

Human dignity and consent. Ethics isn't only correctness; it is care. STEOM's vertical layer is deliberately non-coercive: C (consent) is optional, unoptimized, and logged only at coarse granularity. No "God scores," no rewards for piety, no penalties for silence.

Synthesis. Complexity science contributes spectra, sparsity, and coupling math; quality/safety contributes audit craft and accountability; virtue traditions contribute disciplined practices and a grammar for consent and mercy. STEOM 5 braids them into a conservative program: **measure honestly, practice patiently, love selectively under budget—then test the claims.**

11. Falsifiability, Limits, and Guardrails (STEOM 5 — formerly "STOEM 5-V")

11.1 Disconfirmers and negative controls (including sham-consent)

Primary disconfirmers (map to core claims).

- **D1 — Truth→variance (A6).** High-fidelity audits (MVI≥7/10) produce rho := Var(Q_post)/Var(Q_pre) >= 1 (CI≥1) across replications.
- **D2 — Virtue→efficiency (A5).** Sustained increases in S_truth, S_patience, S_charity do not reduce E = (W_out + Q_loss)/Output at fixed mix, nor tighten Var(Q).
- **D3 — Agency→communion (A3).** Raising Psi via targeted bundles yields no change in synchrony/prosocial measures and no increase in local ||O||_1 vs controls.
- **D4 — Mercy within ledgers (A7).** As ||O||_1 rises (budget respected), severity distributions do **not** shift softer, or "improvements" require hidden energy/info subsidies.
- **D5 — Cadence (A1) & fractals (A2).** No identifiable innovation cutoff tau ~ sqrt(h) under freezes; spectra and MDL show no cross-scale regularities.

Vertical disconfirmers (V-axioms).

- **VD1 — Vertical necessity.** Primary effects (T1–T7) require large Gamma (received grace) to fit; with Gamma := 0, effects vanish.

- **VD2 — Consent mimicry. Sham-consent** (time-matched, content-neutral activity) produces the same "vertical moderation" estimate theta as genuine consent; or theta survives timestamp permutation tests.
- **VD3 — Ledger bypass.** Any model fit implies direct Gamma -> E or Gamma -> L paths (violates V-Conservativity).

Negative controls (must show nulls).

- Cosmetic "audits" (no immutable ledgers/speech tags) ⇒ rho ≈ 1.
- Mismatched practice (e.g., courage bundle on an accounting bottleneck) ⇒ no variance drop.
- Forced densification (μ→0; manual O-links) ⇒ synchrony transient, efficiency flat/worse.
- **Sham-consent** and **timestamp shuffles** ⇒ vertical moderation theta ≈ 0.

Stop rules (fail-fast).

- After 2 clean cycles missing a primary endpoint (pre-powered), pause, publish nulls, and revise mappings/parameters before any scale-up.

11.2 Confounds and mitigations (placebo, selection, Hawthorne)

Placebo / expectancy.

- *Risk:* "Something is happening" drives effort spikes.
- *Mitigation:* Stepped-wedge rollouts; equal contact time for controls; analyst blinding; prereg power & endpoints.

Selection / survivorship.

- *Risk:* High-Psi volunteers or attrition bias.
- *Mitigation:* Randomize seeds from a stratified pool; intent-to-treat + per-protocol; track and model attrition; predefine replacement rules.

Hawthorne / observation effects.

- *Risk:* Measurement alters behavior temporarily.
- *Mitigation:* Mirror instrumentation in controls; extend windows; include post-novelty decay checks.

Seasonality / mix shifts.

- *Risk:* Throughput/mix changes masquerade as gains.
- *Mitigation:* Fix Output definitions; include mix covariates; use matched controls and calendar terms.

Goodhart / gaming.

- *Risk:* Actors optimize the metric, not the work.
- *Mitigation:* Prefer second-order signals (variance-after-truth rho, severity ladders); lock schemas; rotate auditors; publish errata.

Instrumentation drift.

- *Risk:* HRV sensors, defect labeling drift.
- *Mitigation:* Weekly calibration logs; inter-rater reliability (κ≥0.7); device health checks; frozen preprocessing pipelines.

Contamination / spillover.

- *Risk:* Controls adopt practices informally.
- *Mitigation:* Cluster randomization; monitor fidelity; sensitivity analyses with spillover covariates.

Vertical-specific confounds.

- *Risk:* Social desirability around consent; coercion by perks.
- *Mitigation:* No incentives/penalties tied to C; coarse, opt-in logging; separate vertical tables; **never** aggregate a "God score."

11.3 Scope limits: what STEOM 5 does *not* claim (no God-metrics, no energy injections)

No God-metrics.

- God is **not** a variable, parameter, or edge. Only *received grace* Gamma := phi(C) may appear as a **small, consent-gated moderator**, and it is reported separately with shrinkage toward 0. No composite "piety" or "God score" is ever computed.

No energy or information injections.

- Ledgers remain physical:
  $U_{n+1} - U_n = P_{in} - W_{out} - Q_{loss}$ and $H(L_{n+1}) = H(L_n) + \text{Inputs} - \text{Losses}$.
  Gamma cannot modify E or L directly; all observable effects must pass through S -> Psi -> O under budget and sparsity.

No guarantees.

- Practices **bias** distributions; they do not guarantee success each tick. Shock events still occur; mercy is a statistical softening, not elimination.

No universal portability without fidelity.

- Effects depend on doctrine fidelity (audit quality, honest ledgers/speech), parameter ranges (μ, λ, β), and correct targeting. Low-fidelity deployments should fail—and be reported.

No carte blanche for surveillance.

- Measurement is bounded by consent, minimality, fairness, and dignity. Coercive or covert instrumentation violates the model's ethics even if "effective."

No moralizing of nulls.

- Null/negative results indicate model/fit/instrument issues unless independently established otherwise. Revise maps (B_S, B_H, β, μ), not people.

No unbounded communion.

- Attention is scarce: $||O||_1 <= O_{max}'$. Attempting to over-densify O yields cliques, burnout, and instability; budgets are part of the claim.

No single-metric governance.

- rho matters, but success requires coherence across **variance**, **energy per output**, **severity softening**, **consent/fairness integrity**, and **audit fidelity**. Any one can veto.

Accounting for overhead.

- Measurement costs are debited against E. We do not claim "free" gains—net improvements must exceed instrument overhead.

Guardrail checklist (operational).

1. **Preregister** endpoints, analysis code, negative controls (incl. sham-consent).
2. **Lock ledgers**: immutable diffs, timestamp discipline, speech tags; publish audit MVI scores.
3. **Enforce budgets** on O; monitor nnz(O)/N^2 and degree caps.
4. **Fairness/mercy**: severity non-inferiority and E parity by subgroup; stop if violated.
5. **Vertical humility**: keep Gamma OFF by default; if ON, cap moderation, require negative-control passes, and report separately.
6. **Fail fast**: two clean misses → pause, publish, revise, then retry.

12. Roadmap and Open Problems (STEOM 5)

12.1 Multi-site replication network (nulls welcomed)

Purpose. Build an open, lightweight consortium to test T1–T7 across clinics, plants, software teams, and classrooms—using identical kits, prereg, and dashboards—while explicitly **welcoming nulls**.

Phases.

- **P0 — Kernel (6–8 weeks).** Ship v1 "Core Kit": audit checklist + rho calculator; HRV/synchrony pipeline; severity ladder rubric; simulator with Γ/C toggles *off* by default; template dashboards; ethics templates. Form a 3-person steering triad (methods/ethics/data) and a rotating red-team reviewer.
- **P1 — Pathfinders (3–4 months).** 4 sites, one test each: T2 (org), T3 (classroom), T4 (clinic/ward), T5 (plant). Power for *minimal detectable effects*; publish all results (including null/negative) with fidelity scores.
- **P2 — Replication (6–12 months).** 8–12 sites; at least 2 per domain. Add T1 (fractal spectra) and T6 (h-hat). Random-effects meta-analyses; report heterogeneity (τ^2) and moderators (site type, fidelity, mix).
- **P3 — Standards (ongoing).** Freeze v1.1 specs: audit MVI ≥7/10; fixed PSD/DFA estimators; consent logging grammar; O-inference protocol; "vertical moderation" table format with caps and negative-controls.

Data model & governance.

- **Tables:** events (Q, Output, audits), sensors (RR intervals, speech marks), derived features (PLV, turn-taking), outcomes (rho, E, severity, O_L1), consent exposure (coarse, opt-in).
- **Privacy:** opt-in, minimal, retention caps; no covert collection.
- **Transparency:** publish config/code hashes; pre/post windows; CIs; ethics notes.
- **Incentives:** quarterly "Best Null" and "Best Red-Team" awards; small grants for replication of negative findings.

Open problems (network).

- Incentive-compatible data sharing that protects participants yet enables meta-analysis.
- Harmonizing "Output" across domains without overfitting to one.
- Building convincing **negative-control libraries** (sham audits, sham-consent) that are ethical *and* realistic.

12.2 Learning O from multimodal streams under causal constraints

Goal. Infer a time-varying oneness tensor $O_t$from physiology, behavior, and speech, such that edges represent *causal* coordination capacity, satisfy sparsity/budget constraints, and transport across sites.

Inputs.

- Physiological: RR intervals, respiration, actigraphy.
- Behavioral: turn-taking, interrupts, help/corrections, handoffs.
- Linguistic: claim tags (obs/inf/intent), lexical alignment.
- Interventions: audits, practice bundles (u_S), habit supports (u_H), seed placement.
- Ledgers: E per output, Q, severity ladder.

Modeling spine.

- **State-space graph.**
  $x_{t+1} = A_t x_t + B_t u_t + w_t$, $y_t = C_t x_t + v_t$; encode $O_t = g(A_t)$with sparsity + L1 budget.
- **Priors.** Two kernels: $K_{\text{phys}}(F, Z)$and $K_{\text{agency}}(\Psi)$; elastic weighting.
- **Causal regularizers.**

- **Invariance (ICP):** prefer edges stable across context shifts and interventions.
- **Do-anchors:** when seeds/audits fire, up-weight edges whose deltas match predeclared directionality.
- **Temporal guards:** constrained lags; forbid instantaneous cross-channel backdoors.

Outputs.

- O_t; local/global ‖ O_t ‖$_1$; edge stability; decomposition (phys vs agency contribution).
- Predictive checks: out-of-sample synchrony and severity forecasting; counterfactuals (β→0).

Validation plan.

- **Sim-to-real.** Train on synthetic worlds with known O; port to field data unchanged; report recovery and calibration plots.
- **Interventional tests.** Edge deltas around audits/seeds must follow priors; placebo (calendar) should not.
- **Transport.** Train on site A, test on B; measure drop in synchrony prediction and edge-stability curves.

Open problems (learning).

- Partial identifiability proofs for $O_t$under bounded interventions and missing-not-at-random sensors.
- Robust handling of charismatic outliers (one agent dominating synchrony).
- Online forgetting vs regime-change detection without erasing true communion.
- Disentangling simultaneous rises in truth and charity coordinates, given collinearity.

12.3 Formal conservativity: thermodynamic and information-theoretic bounds with Γ present but non-energetic

Aim. Prove STEOM 5 is a **conservative extension**: vertical moderation (Γ) may shape dispositions and budget ceilings but cannot break energy/information ledgers.

Targets & conjectures.

1. **Thermodynamic compatibility.**
   Claim. For any admissible policy sequence $\pi_t$induced by audits/practices and any consent path C_t(with $\Gamma_t = \phi(C_t)$), the plant obeys
   $U_{t+1} - U_t = P_{\text{in}} - W_{\text{out}} - Q_{\text{loss}}$.
   Program. Map "grace-as-efficiency" to waste-reducing schedulers within resource-theory constraints; show expected Q_{loss}decreases via policy changes, not extra work extracted.

2. **No-signaling / causal order.**
   Claim. $O_t$alters conditional dependencies but cannot transmit superluminal signals or inject information.
   Program. Prove a data-processing inequality for *coordination information* $I_{\text{coord}}$with an upper bound set by channel capacities of L; show $\Delta I_{\text{coord}} \leq$ function of $\| O_t \|_1$and $H(L)$.

3. **Variance-after-truth bound.**
   Claim. Honest audits increase mutual information I(state; measurement)and lower Bayes risk for allocation, implying $E[V_{t+1}] < E[V_t]$.
   Program. Derive a bound linking $\rho$to an information gain term under misspecification and heavy tails (van Trees–style inequalities).

4. **Budgeted mercy bound.**
   Claim. With $\| O_t \|_1 \leq O_{\max}$, the achievable reduction in severe-tail probability is upper-bounded by a function of $O_{\max}$and network curvature.
   Program. Use flow/cut inequalities on the attention graph to upper-bound tail-softening.

5. **Vertical conservativity.**
   Claim. If Γonly affects S(theological coordinates) and $O_{\max}$multiplicatively, then any trajectory achievable with $\Gamma > 0$is approximable by a purely creaturely controller with slightly modified practice schedules, up to overhead.
   Program. Construct an ε-simulation that replaces Γ-lift with additional (bounded) $u_S$and budget tuning; quantify overhead and require it to be included in E.

6. **Landauer-aligned accounting.**
   Claim. Measurement overhead (audits, sensing) lower-bounds energy spent on information; net Eimprovements must exceed this cost.
   Program. Embed metering costs into E; forbid success declarations if overhead > savings.

Open problems (formal).

- Tight ρ–information inequalities under non-Gaussian losses.
- Conditions for **bundle commutativity** (order of truth/patience/charity interventions) vs provable non-commutativity.
- Renormalization view of effective h: invariants when task granularity and audit cadence scale together.
- Proving that bounded δ(Γ-cap on $O_{\max}$) guarantees stability (no clique attractors) under broad noise classes.

Operational checkpoints.

- In every field study: (i) report measurement overhead, (ii) net Eafter overhead, (iii) severity shifts with CIs, (iv) $H(L)$changes, (v) vertical tables (often θ≈0).
- Treat any claimed "miracle" as a bookkeeping error until proven otherwise.

Bottom line. The roadmap builds people/process infrastructure for honest replication, the learning program turns messy streams into causally disciplined O, and the formal agenda keeps STEOM 5 conservative with physics and information—so any measured "grace" is precisely what we say it is: *better coordination within the ledgers we already owe.*

Appendices (STEOM 5)

A. Copy-safe math & symbol glossary (adding Gamma, C)

Conventions (ASCII-safe).

- Scalars: lower-case (h, kappa, rho, tau). Vectors: bold in prose; in code use arrays (H, S, Psi). Matrices/tensors: capitals (O, K_phys, K_agency).
- Products: Hadamard (componentwise) *; matrix multiply @; dot <a,b>.
- Norms: ||X||_1 L1; budgets like ||O||_1 <= O_max'.
- Updates: X_{n+1}; differences DeltaE := E_post - E_pre.
- Expectations/variance: E[.], Var[.].

Core state (God not a variable).

X_n := {F_n, Z_n, L_n, E_n, G_n, O_n, H_n, S_n, Psi_n}

Psi_n := H_n * S_n

New vertical symbols.

C_n in [0,1] # consent/attention (voluntary)

Gamma_n := phi(C_n) # received grace (latent, bounded, monotone in C)

Gamma_bar_n # EWMA of Gamma (smoothing alpha_g)

Key maps.

- K_phys(F,Z) structural kernel, K_agency(Psi) agency kernel.
- O_max'(Psi_{0:n}, Gamma_bar_n) = O_max(Psi_{0:n}) * (1 + delta * Gamma_bar_n).
- S theological subspace: {faith, hope, charity} with **virtue floors**: S_j >= 0.

Primary metrics.

Q_n := loss per tick / Output_n

rho := Var(Q_post) / Var(Q_pre)

E_per_output := (W_out + Q_loss) / Output

h_hat := (tau/c)^2 # from innovation cutoff

B. Full pseudocode and parameter tables (Γ, C toggles and priors)

B.1 Reference tick (vertical OFF by default)

```
def tick(X, P, U, rng, use_gamma=False):
   F,Z,L,E,G,O,H,S,Psi = X
   xi = rng.normal(size=P.noise_dim) * (P.kappa * P.h) ** 0.5

   # Backbone
   F = D_F(F,Z) + S_F(F,Z,xi,P.h)
   Z = D_Z(Z,F) + S_Z(Z,F,xi,P.h)
   L = D_L(L,F,Z) + S_L(L,xi,P.h)
   E = D_E(E,Z) + S_E(E,xi,P.h)
   G = D_G(G,F,E) + S_G(G,xi,P.h)

   # O-dynamics (budgeted + sparse)
   O_tilde = (1-P.lambda_) * O \
         + P.lambda_ * sig(P.alpha*K_phys(F,Z) + P.beta*K_agency(Psi)) \
         + S_O(O,xi,P.h)

   Omax = O_max(P.Psi_hist)
   if use_gamma:
      # Consent dynamics
      C_prev = getattr(P, "C", 0.0)
      C_now  = clip01(P.RC * C_prev + P.BC * U.uC + noise_small())
      P.C = C_now
      Gamma = phi(C_now)  # received grace (latent)
      Omax  = Omax * (1 + P.delta * EWMA(Gamma, P.alpha_g))
```

```
  O = Proj_budget( SoftThresh(O_tilde, P.mu), Omax )

  # Agency channels
  H = P.RH @ H + P.BH @ U.uH + P.NH @ xi + P.GH @ Psi
  S = P.RS @ S + P.BS @ U.uS + P.NS @ xi
  if use_gamma:
     S = S + P.etaS * Gamma * dir_S_theological
  S = virtue_floor(S)  # enforce nonnegativity on theological coords

  Psi = H * S
  return (F,Z,L,E,G,O,H,S,Psi)
```

B.2 Parameter table (suggested ranges; copy-safe)

Param	Meaning	Typical
h	tick cadence	0.5–2.0
kappa	noise scale	0.5–2.0
lambda	O learning/forget	0.05–0.30
alpha	weight on K_phys	0.4–0.7
beta	weight on K_agency	0.3–0.6
mu	O sparsity thresh	0.01–0.10
RH, RS	retention (diag)	0.90–0.98
BH, BS	input gains	0.02–0.10
NH, NS	noise gains	0.01–0.03
GH	Psi->H gain	0.02–0.05
O_sat, gamma, w_decay	budget curve	domain-tuned

Param	Meaning	Typical
RC	consent retention	0.90–0.97
BC	consent input gain	0.02–0.06
phi	C→Gamma map	identity or logistic(k*x)
etaS	S-lift gain	0.01–0.05
delta	O_max expansion	0.05–0.15
alpha_g	EWMA for Gamma	0.2–0.4

B.3 Priors for moderated fits (Γ, shrinkage)

When (and only when) preregistered:

theta ~ Normal(0, sigma_theta^2), with sigma_theta small (e.g., 0.05*effect_scale)

0 <= Gamma_bar <= 1   (derived from C; coarse, opt-in logging)

Cap: |theta| <= theta_cap   # e.g., 0.25 * base effect size

Always fit and report the **base model (theta=0)** first; moderated model must pass negative controls.

C. Instrumentation: audit checklist, variance-after-truth worksheet, virtue index (with theological floor), Γ reporting rubric

C.1 Audit MVI (0–10 score; need ≥7 to count as "audited")

- (1) Calibrated weights & uncertainty posted (0/1)
- (2) Immutable ledger with hashed diffs (0/1)
- (3) ISO-8601 timestamps; no retro-edits without errata (0/1)
- (4) Role-separated writes; rotating countersign (0/1)
- (5) Data completeness monitor (0/1)
- (6) Speech tags on claims (obs/inf/intent) (0/1)
- (7) Outlier policy predeclared (0/1)
- (8) Mix controls frozen (0/1)
- (9) Analyst blinding until code freeze (0/1)
- (10) Ethics note posted (0/1)

C.2 Variance-after-truth worksheet (copy-safe)

Define Q_n # loss/Output with domain units

Choose windows: [n0:n1]=pre, [n1+1:n2]=post

Qn_norm := Q_n

Var_pre := Var(Qn_norm[n0:n1])

Var_post := Var(Qn_norm[n1+1:n2])

rho := Var_post / Var_pre

Report: rho, 95% CI (bootstrap or delta), audit MVI score, completeness %

C.3 Virtue index (short-form, 0–1 scaled; includes theological floor)

- **Truth**: (a) "no exaggeration" compliance; (b) peer-checked factual accuracy; (c) timely errata.
- **Patience**: (a) breath cadence adherence; (b) delay-before-respond rule; (c) impulsive edit rate ↓.

- **Charity**: (a) service acts logged; (b) help-per-tick; (c) reconciliation attempts.
  Scoring. Average normalized sub-items; for {faith, hope, charity}, enforce max(0, .) after noise—never negative.

C.4 Γ reporting rubric (V-Humility in practice)

- **Default**: Γ OFF; report base results.
- **If prereg moderation is ON**:
  - Use small, centered priors; predeclare theta_cap.
  - Log **consent exposure (C)** coarsely (yes/no or minutes), opt-in only.
  - Run negative controls: sham-consent, timestamp permutation, audit-failure screen.
  - Present a **separate vertical table** with theta, CI, cap, controls passed.
  - Never aggregate into a composite "God score."
  - If moderation needed to meet primaries → declare **model/fidelity fault**, not success.

D. Prereg templates (vertical negative controls; shrinkage priors for Γ)

D.1 Header

- Title, site, PI, version hash/date.
- Primary hypothesis (exact metric), secondary endpoints.
- Units of analysis; tick definition; inclusion/exclusion.
- Baseline/post windows; missing/outlier policy.
- Power calc; stopping rules; ethics note URL.

D.2 Analysis plan (two-model)

Base model: Effect ~ beta0 + beta1 * Intervention + controls

Moderated model: Effect ~ beta0 + beta1 * Intervention + theta * Gamma_bar + controls

theta ~ N(0, sigma_theta^2), |theta| <= theta_cap

Out-of-sample evaluation preregistered (CV or holdout)

- Base determines success; moderated reported only if it improves OoS fit and passes negatives.

D.3 Negatives & shams (mandatory)

- **Sham-audit** (cosmetic) → expect rho ≈ 1.
- **Sham-consent** (time-matched, content-neutral) → expect theta ≈ 0.
- **Timestamp permutation** for consent exposure → expect theta ≈ 0.
- **Calendar shuffle** when seasonality suspected.

D.4 Domain inserts

- **Org (T2/T5)**: stepped-wedge audits; ABAB energy tests; fixed throughput band; mix covariates.
- **Lab (T1/T6)**: fixed PSD/DFA pipeline; process freeze for h-inference.
- **Community (T3/T4)**: seed selection/assignment; synchrony/prosocial rubrics; severity ladder.

D.5 Publication policy

- Post prereg, code, raw/tidy data, audit MVI, rho with CI, E per output net of overhead, severity shifts, synchrony, h-hat.
- Publish nulls/negatives with fidelity scores; award “Best Null/Red-Team.”

E. Ethics note: transparency, consent, non-coercion, pastoral safeguards (explicit ban on "optimizing" consent)

E.1 Transparency. Plain-language purpose: "We test whether truthful measurement reduces waste and whether patient charity improves coordination, within fixed energy/information ledgers." Post metrics, stop rules, and how overhead is charged.

E.2 Consent & dignity.

- Participation voluntary; opt-out without penalty; no covert recording.
- Individual signals private by default; group-level reporting.
- Minimal data; retention caps; secure storage; role-separated access.

E.3 Non-coercion for vertical elements.

- **Explicit ban on 'optimizing' consent (C).** No incentives, penalties, quotas, or performance linkage. No "consent dashboards."
- Consent exposure (if any) logged coarsely and only with opt-in.

E.4 Fairness & mercy constraints.

- Success requires **all**: rho<1 **and** E net of overhead non-increasing **and** severity non-inferior by subgroup.
- If benefits correlate with burden shifts (e.g., vulnerable groups take the hit), declare failure and remediate.

E.5 Accountability without humiliation.

- Audits evaluate processes, not persons; errata welcomed.
- Prohibit public league tables of individual virtue; use team-level dashboards.

E.6 Oversight & red-teaming.

- Independent reviewer signs off prereg, audit quality, consent materials, harm checks.
- Red-team metrics for Goodhart surfaces; maintain a library of negative controls.

E.7 Stop rules (harm).

- Stop if E worsens >10% for 2 ticks, or severity odds worsen, or audit MVI < threshold.
- Publish negatives, debrief, revise maps/params, then retry.

E.8 Theology boundary.

- No metaphysical claims from metrics; God remains unmodeled. Γ, if reported, is a small, consent-gated moderator—often 0—and never counted toward success.

Use & packaging. Drop this appendix as /docs/appendix_A_to_E.md in your repo. Keep versioned. Require each site to check C.1, complete D., and post E. before collecting data.