Signal Routing Under Structural Distance: Transmission, Absorption, and Lossy Projection Under Pressure

Justin Parten*

January 29, 2026 v1.0

Abstract

Hierarchical institutions often convert rich, local signals into legible artifacts as they move upward. That conversion supports coordination and decision, but it can also discard information and collapse expressed uncertainty. This paper develops a minimal signal-routing model for that tradeoff. A nested organization is represented as a rooted tree. Each node tracks a scalar world state using a Gaussian belief summary and, at each time step, chooses one of three routing operators: *transmission* (send the current belief upward with channel noise), *absorption* (handle locally and send nothing), or *lossy projection* (quantize the estimate and shrink variance before sending).

Structural distance enters as a positional load that increases both observation noise and communication noise. Epistemic capacity enters as a stabilizing resource that mitigates degradation and makes absorption feasible. We also introduce an *epistemic headroom* rule in which time-varying pressure triggers projection when it exceeds a node's headroom, producing depth-stratified routing regimes. Comparative simulations show that widespread projection increases root-level error and produces recurrent overconfidence episodes relative to widespread transmission, while headroom routing yields a mixed regime in which upper layers project under pressure and lower layers absorb when slack permits.

Keywords: signal routing; structural distance; epistemic drift; uncertainty collapse; organizational epistemology; pressure and decision-making; hierarchical systems

Suggested citation: Parten, J. (2026). *Signal Routing Under Structural Distance: Transmission, Absorption, and Lossy Projection Under Pressure*. Preprint.

1 Introduction

Modern organizations rarely fail because nobody observes anything. They fail because what is observed does not remain intact as it is routed through layers that must summarize, standardize, and act. Upward communication tends to reward artifacts that are brief, comparable, and decision-ready. Those same artifacts can erase distinctions that mattered downstream and can convey a level of confidence that is not supported by the underlying signal.

This paper isolates that mechanism in a minimal model. The goal is not to reconstruct a particular institution, nor to offer a complete theory of organizational inference. The goal is to make a specific routing tension explicit: high-fidelity reporting preserves nuance but accumulates channel degradation, while lossy projection produces legible summaries that may collapse expressed uncertainty.

*Contact: `jparten1@proton.me`

1.1 Contributions and scope

The paper makes three contributions.

(1) A simple architecture for *signal routing under distance*: a rooted tree in which each node tracks a dynamic world state, combines local observation with upstream reports, and then chooses a routing action.

(2) A formal distinction between *transmission* (noisy but information-preserving in expectation), *absorption* (local handling without upward reporting), and *projection* (non-invertible compression that can reduce expressed uncertainty).

(3) A pressure-sensitive extension (*epistemic headroom*) that makes projection a threshold response to time-varying pressure, together with comparative root-level diagnostics.

The modeling choices are deliberately small. The world is one-dimensional, belief states are summarized by a mean and variance, and updates take a Gaussian form so that inference remains transparent.

2 Architecture and notation

2.1 Tree structure

Let the organization be a rooted k-ary tree of depth D. Nodes are indexed by i. The depth of node i is $d(i) \in \{0, \dots, D\}$, where $d = 0$ denotes the root and $d = D$ denotes the leaves. Each non-root node has a unique parent, written $\pi(i)$, and each non-leaf node has k children.

The tree is an information-routing scaffold. Leaves represent positions closest to ground conditions, where high-resolution observations originate. The root represents a meta-level decision layer, where many upstream signals are aggregated into a single operative representation. Edges are upward reporting channels: at each time step, a node may send an informational message to its parent, and messages are progressively aggregated as they move toward the root.

2.2 Belief representation

At each time t, the world has a scalar state $W(t)$. Each node i maintains a belief state summarized by a mean $\mu_i(t)$ and variance $v_i(t)$. It is often convenient to write $\sigma_i(t) = \sqrt{v_i(t)}$.

The Gaussian form is a representational convention: $\mu_i(t)$ is the node's best estimate of $W(t)$ given its current information, and $v_i(t)$ represents the node's acknowledged uncertainty about that estimate. All inference and routing operations in the model act on these two numbers.

2.3 Structural parameters

Each node is assigned three role-level parameters.

- **Structural distance** $\lambda_i \in [0, 1]$ is a reduced-form measure of insulation from ground conditions and corrective feedback. Higher λ_i increases both observation noise and channel noise.
- **Epistemic capacity** $\kappa_i \in [0, 1]$ represents local stabilizing resources (monitoring, redundancy, protected dissent, error-correction routines). Higher κ_i reduces effective noise and makes absorption feasible.
- **Baseline pressure** $\bar{p}_i \in [0, 1]$ represents steady decision pressure at the role. In the epistemic-headroom extension, instantaneous pressure fluctuates around $\bar{p}_i$.

2.4 Notation and interpretation

The model separates three kinds of objects.

World and time. $W(t)$ is the ground state at time t. Terra Firma (Section 3) specifies how $W(t)$ changes.

Beliefs and signals. Node i summarizes its belief about $W(t)$ by $(\mu_i(t), v_i(t))$. The pair $(\mu_i^-(t), v_i^-(t))$ denotes the time-t prior before incorporating time-t evidence. Each node also draws a local observation $y_i(t)$ and may receive messages $m_{j\to i}(t)$ from children.

Structure and routing. Structural distance λ_i and epistemic capacity κ_i modulate observation and channel noise (Section 6). A routing action $\mathcal{A}_i(t)$ determines whether the node transmits, absorbs, or projects its belief (Section 5). In the epistemic-headroom extension, the node compares instantaneous pressure $p_i(t)$ to headroom $H_i(t)$.

Table 1 provides a compact symbol reference.

Symbol	Meaning
$W(t)$	World state at time t
$\mu_i(t), v_i(t)$	Node i posterior mean and variance
$\mu_i^-(t), v_i^-(t)$	Node i prior mean and variance before time-t evidence
$y_i(t)$	Node i local observation
$m_{i\to\pi(i)}(t)$	Message from node i to its parent
$\sigma_{\text{obs},i}$	Observation noise standard deviation at node i
$\sigma_{\text{edge},i}$	Edge noise standard deviation on $i \to \pi(i)$
λ_i	Structural distance (insulation) at node i
κ_i	Epistemic capacity (stabilizing resource) at node i
$\bar{p}_i$	Baseline pressure at node i
$H_i(t)$	Epistemic headroom at node i (time-varying)
$p_i(t)$	Instantaneous pressure at node i (time-varying)
$\mathcal{A}_i(t)$	Routing action chosen by node i at time t
$\Delta, \rho, v_{\min}$	Projection parameters (bin width, shrink factor, variance floor)

Table 1: Notation quick reference.

With architecture and notation in place, we now specify the world dynamics and the node-level update.

3 Terra Firma: a dynamic world state

The model requires a ground state that changes over time. Terra Firma is a scalar process:

$$W(t) = W(t-1) + \varepsilon_W(t), \qquad \varepsilon_W(t) \sim \mathcal{N}(0, \sigma_W^2). \tag{1}$$

The process noise σ_W encodes the rate at which the environment moves. Even if a node were perfectly aligned at time $t-1$, uncertainty at time t increases because the world itself may have shifted.

4 Local inference at a node

Nodes are processed bottom-up each time step. When node i updates at time t, it forms a prior from its previous posterior, draws a local observation, and fuses all available signals into a posterior.

4.1 Time propagation and priors

The superscript $-$ denotes the prior at time t before incorporating time-t evidence. In the simplest propagation rule,

$$\mu_i^-(t) = \mu_i(t-1), \qquad v_i^-(t) = v_i(t-1) + \sigma_W^2. \tag{2}$$

The mean is carried forward, while the variance expands by the world process noise. The expansion reflects dynamical uncertainty: yesterday's certainty does not automatically survive a moving environment.

4.2 Local observation

Node i receives a noisy observation

$$y_i(t) = W(t) + \varepsilon_{\text{obs},i}(t), \qquad \varepsilon_{\text{obs},i}(t) \sim \mathcal{N}(0, \sigma_{\text{obs},i}^2), \tag{3}$$

where $\sigma_{\text{obs},i}$ depends on the node's structural distance and epistemic capacity (Section 6).

4.3 Messages from children

If child j transmits or projects, it sends a message consisting of a mean and variance pair. The parent treats that message as another noisy measurement of $W(t)$, with additional uncertainty from the channel. Concretely, a message from j to i is treated as having mean $\mu_j(t)$ and variance $v_j(t) + \sigma_{\text{edge},j}^2$.

4.4 Fusion

Under conditional independence and Gaussian noise, the posterior is the precision-weighted average:

$$v_i(t) = \left(\frac{1}{v_i^-(t)} + \frac{1}{\sigma_{\text{obs},i}^2} + \sum_{j \in \text{children}(i)} \mathbf{1}\{j \text{ sent}\} \frac{1}{v_j(t) + \sigma_{\text{edge},j}^2} \right)^{-1}, \tag{4}$$

$$\mu_i(t) = v_i(t) \left(\frac{\mu_i^-(t)}{v_i^-(t)} + \frac{y_i(t)}{\sigma_{\text{obs},i}^2} + \sum_{j \in \text{children}(i)} \mathbf{1}\{j \text{ sent}\} \frac{\mu_j(t)}{v_j(t) + \sigma_{\text{edge},j}^2} \right). \tag{5}$$

Each source receives weight proportional to its precision (inverse variance). Signals with high uncertainty contribute little to the posterior mean.
With local inference defined, the remaining modeling choice is how nodes route their posteriors upward.

5 Routing operators

After updating, node i chooses a routing action $\mathcal{A}_i(t) \in \{\text{TRANSMIT}, \text{ABSORB}, \text{PROJECT}\}$. The action determines what, if anything, is sent to the parent.

5.1 Transmission

If node i transmits, it sends its current belief state upward. The reporting channel adds noise, modeled as additional variance:

$$m_{i\to\pi(i)}(t) = \big(\mu_i(t),\ v_i(t) + \sigma^2_{\mathrm{edge},i}\big). \tag{6}$$

Transmission preserves the mean estimate but makes the report less reliable as structural distance increases.

5.2 Absorption

If node i absorbs, it does not send an upward report:

$$m_{i\to\pi(i)}(t) = \varnothing. \tag{7}$$

Absorption models local handling of uncertainty. The node may act, stabilize, or resolve the issue within its own scope. The parent receives no additional information from that branch at that time step.

5.3 Lossy projection

If node i projects, it converts its belief into a legible artifact that is easier to communicate and use for decision, at the cost of information loss. Projection has two components.

First, the mean is quantized into bins of width Δ:

$$\tilde{\mu}_i(t) = Q_\Delta(\mu_i(t)) := \Delta \cdot \mathrm{round}\left(\frac{\mu_i(t)}{\Delta}\right). \tag{8}$$

Second, the variance is shrunk toward a floor:

$$\tilde{v}_i(t) = \max\{\rho\, v_i(t),\ v_{\min}\}, \tag{9}$$

where $\rho \in (0,1)$ and $v_{\min} > 0$ are projection parameters.

The projected message is then sent over the channel:

$$m_{i\to\pi(i)}(t) = \big(\tilde{\mu}_i(t),\ \tilde{v}_i(t) + \sigma^2_{\mathrm{edge},i}\big). \tag{10}$$

Quantization makes the transformation many-to-one, so distinct upstream states can map to the same artifact. Variance shrinkage models the collapse of expressed uncertainty that often accompanies institutional projection into statuses, categories, or dashboards.

6 Distance and capacity as noise modifiers

Structural distance and epistemic capacity enter through observation noise and edge noise.

6.1 Observation noise

The observation noise at node i is

$$\sigma_{\mathrm{obs},i} = \sigma_{\mathrm{obs},0}\big(1 + s_{\mathrm{obs}}\lambda_i\big)\big(1 + s_\kappa(1-\kappa_i)\big), \tag{11}$$

where $\sigma_{\mathrm{obs},0}$ is a baseline noise scale, s_{obs} controls the impact of distance, and s_κ controls the impact of low capacity. Higher distance and lower capacity both degrade local sensing.

6.2 Edge noise

The reporting channel from i to $\pi(i)$ has noise

$$\sigma_{\text{edge},i} = \sigma_{\text{edge},0}\big(1 + s_{\text{edge}}\lambda_i\big)\big(1 + s_\kappa(1 - \kappa_i)\big). \tag{12}$$

This encodes a simple structural claim: insulation and weak scaffolding degrade not only what a node can see, but also what it can transmit.

7 Routing policies (scenarios)

To study how routing shapes root-level accuracy and calibration, we compare several policies.

7.1 Always transmit

Every node transmits at every time step.

7.2 Always project

Every node projects at every time step.

7.3 Random routing

Each node independently selects transmit, absorb, or project with equal probability at each time step.

7.4 Gain stage at depth d_g

Nodes with depth $d < d_g$ project, while nodes with depth $d \geq d_g$ transmit. This implements a stylized mid-level dashboard layer that converts upstream signals into artifacts before they reach the top.

7.5 λ/κ-biased routing

Nodes sample an action from a softmax distribution whose logits favor projection for high λ and favor absorption for high κ:

$$\Pr(\mathcal{A}_i(t) = a) \propto \exp\{\eta\, u_i(a)\}, \tag{13}$$

where $\eta > 0$ is an inverse-temperature parameter and

$$u_i(\text{TRANSMIT}) = 0, \quad u_i(\text{ABSORB}) = \kappa_i, \quad u_i(\text{PROJECT}) = \lambda_i. \tag{14}$$

8 Epistemic headroom: pressure-triggered projection

The preceding policies treat routing as a stationary rule. Many institutional projections, however, occur intermittently, often during episodes of urgency. Epistemic headroom models that by making projection a threshold response to time-varying pressure.

8.1 Headroom and pressure

At each time t, node i has a headroom score

$$H_i(t) = h_0 + b_\kappa \kappa_i - a_\lambda \lambda_i + \varepsilon_H(t), \qquad \varepsilon_H(t) \sim \mathcal{N}(0, \sigma_H^2), \tag{15}$$

and an instantaneous pressure

$$p_i(t) = \bar{p}_i + \varepsilon_p(t), \qquad \varepsilon_p(t) \sim \mathcal{N}(0, \sigma_p^2). \tag{16}$$

Headroom is higher when capacity is higher and distance is lower. Pressure fluctuates around a role baseline.

8.2 Routing rule

Epistemic headroom (EH) selects actions by thresholding pressure against headroom:

$$\text{if } p_i(t) > H_i(t) \text{ then PROJECT.} \tag{17}$$

When pressure does not exceed headroom, projection is not forced. In that case, EH distinguishes absorption from transmission by requiring sufficient slack for absorption and by restricting absorption to nodes with adequate capacity and low distance:

$$\text{if } H_i(t) - p_i(t) \geq s_{\text{abs}} \text{ and } \kappa_i \geq \kappa_{\min} \text{ and } \lambda_i \leq \lambda_{\max} \text{ then ABSORB, else TRANSMIT.} \tag{18}$$

This encodes a conservative stabilizing idea: absorption is treated as feasible when the node is close to ground conditions and well-scaffolded, and when pressure is not consuming available headroom.

9 Simulation protocol

At each time step t:

1. Terra Firma advances according to the random-walk update.
2. Nodes are processed bottom-up from deepest level to root.
3. Each node forms its prior, samples its local observation, and fuses all available signals into a posterior $(\mu_i(t), v_i(t))$.
4. After updating, the node selects a routing action according to the scenario policy.
5. If the node transmits or projects, its message is degraded by edge noise and delivered upward.
6. If the node absorbs, no message is sent.

Default parameter values used in the illustrative run are listed in Appendix A.[1]

[1]The appendix reports the baseline tree geometry, noise scales, projection parameters, and the epistemic-headroom thresholds used for the example tables and figures.

10 Diagnostics

We report two root-level diagnostics.

Absolute error. How far the root's estimate is from the world:

$$e(t) = |\mu_{\mathrm{root}}(t) - W(t)|. \tag{19}$$

Calibration ratio. Error measured in standard-deviation units:

$$r(t) = \frac{e(t)}{\sqrt{v_{\mathrm{root}}(t)}}. \tag{20}$$

When beliefs are approximately Gaussian, values above 2 correspond to misses that should be rare under the claimed uncertainty. Persistent episodes with $r(t) > 2$ therefore indicate overconfidence at the top layer.

11 Results

Table 2 reports summary statistics for a default run with one replicate per scenario. The Terra Firma path and node parameters are held fixed across scenarios for comparability, while observation and channel noise draws are scenario-specific.

Scenario	$\mathbb{E}[e]$	median(e)	$P_{95}(e)$	$\mathbb{E}[r]$	$\Pr(r > 2)$	$\mathbb{E}[\sigma_{\mathrm{root}}]$
Best: transmit	0.0861	0.0688	0.2278	1.139	0.160	0.0757
Worst: project	0.1718	0.1626	0.3514	2.222	0.552	0.0775
Random routing	0.1232	0.1057	0.3073	1.410	0.232	0.0871
Gain stage at depth 2	0.1203	0.1142	0.2437	1.584	0.304	0.0758
λ/κ-biased routing	0.1578	0.1644	0.3154	1.941	0.500	0.0815
EH headroom routing	0.1350	0.1253	0.3179	1.747	0.352	0.0775

Table 2: Root-level error and calibration summaries for the illustrative run. Here $e(t) = |\mu_{\mathrm{root}}(t) - W(t)|$, $r(t) = e(t)/\sigma_{\mathrm{root}}(t)$, and $\sigma_{\mathrm{root}}(t) = \sqrt{v_{\mathrm{root}}(t)}$.

In this run, two patterns stand out.

Projection increases both error and overconfidence. Ubiquitous projection roughly doubles mean absolute error relative to ubiquitous transmission and substantially increases the rate of overconfidence episodes. Lossy compression discards information, and variance shrinkage reduces expressed uncertainty even when error remains.

Epistemic headroom produces a mixed regime. EH reduces the most extreme overconfidence burden relative to the λ/κ softmax policy, but it still produces repeated spikes. Relative to random routing, EH tends to deliver a less noisy root estimate (smaller average σ_{root}) while carrying higher mean error.

11.1 Action frequencies by depth

Figures 1 and 2 show action frequencies by depth for the λ/κ softmax policy and EH routing.

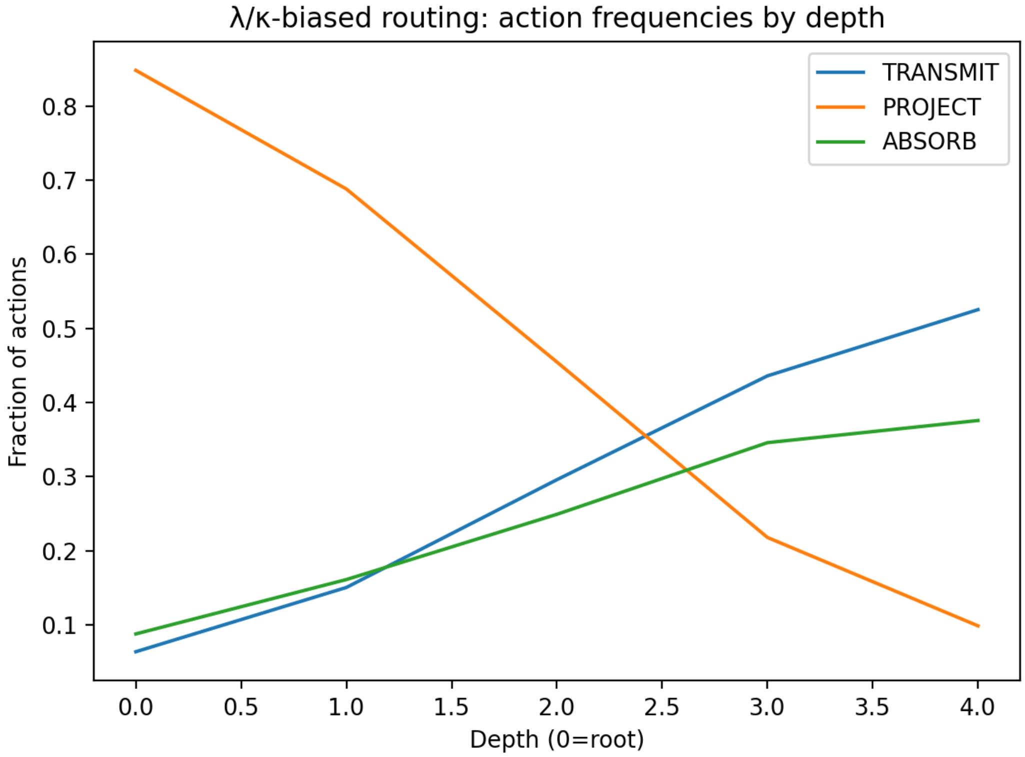

Figure 1: Action frequencies by depth under λ/κ-biased routing. The policy produces a smooth shift from projection at the top to transmission and absorption at deeper layers.

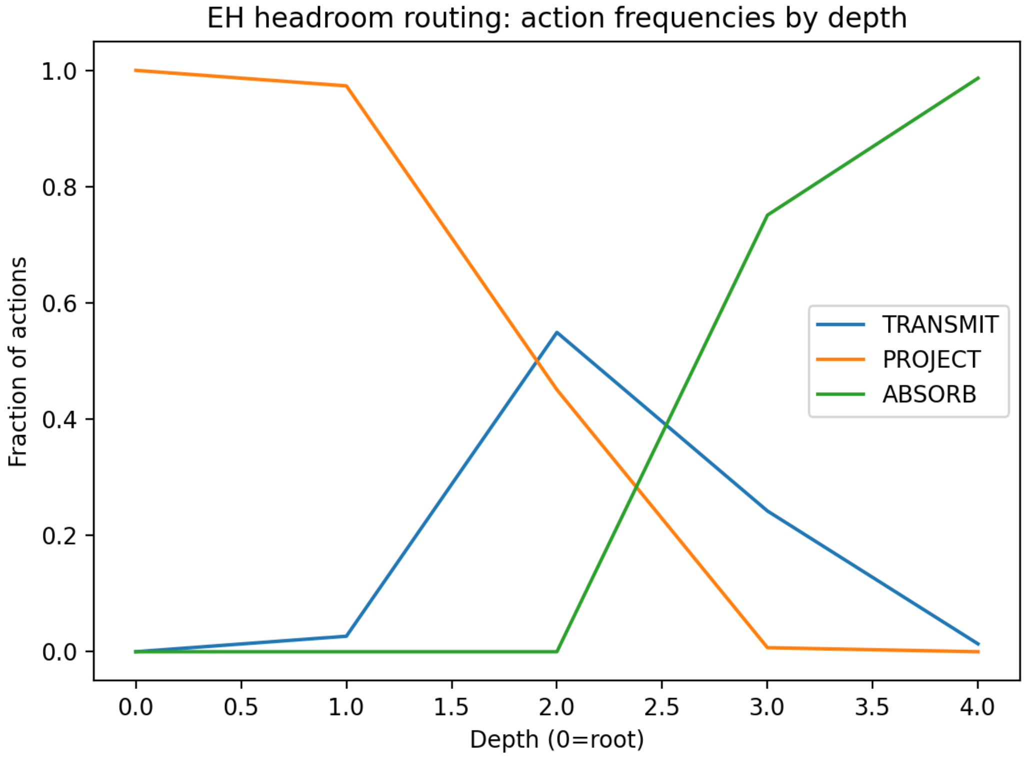

Figure 2: Action frequencies by depth under epistemic headroom routing. The policy induces a depth-stratified regime: projection dominates at shallow depth, transmission dominates near the middle, and absorption becomes frequent near the leaves.

Under the default depth gradients, EH tends to generate a banded routing structure. Near the root, headroom is low and pressure spikes are common, so projection dominates. Mid-level nodes frequently transmit, acting as relays. Near the leaves, capacity is higher and distance is lower, so absorption becomes eligible and frequent.

11.2 Root diagnostics over time

Figures 3 and 4 plot $e(t)$ and $r(t)$ over time for all scenarios.

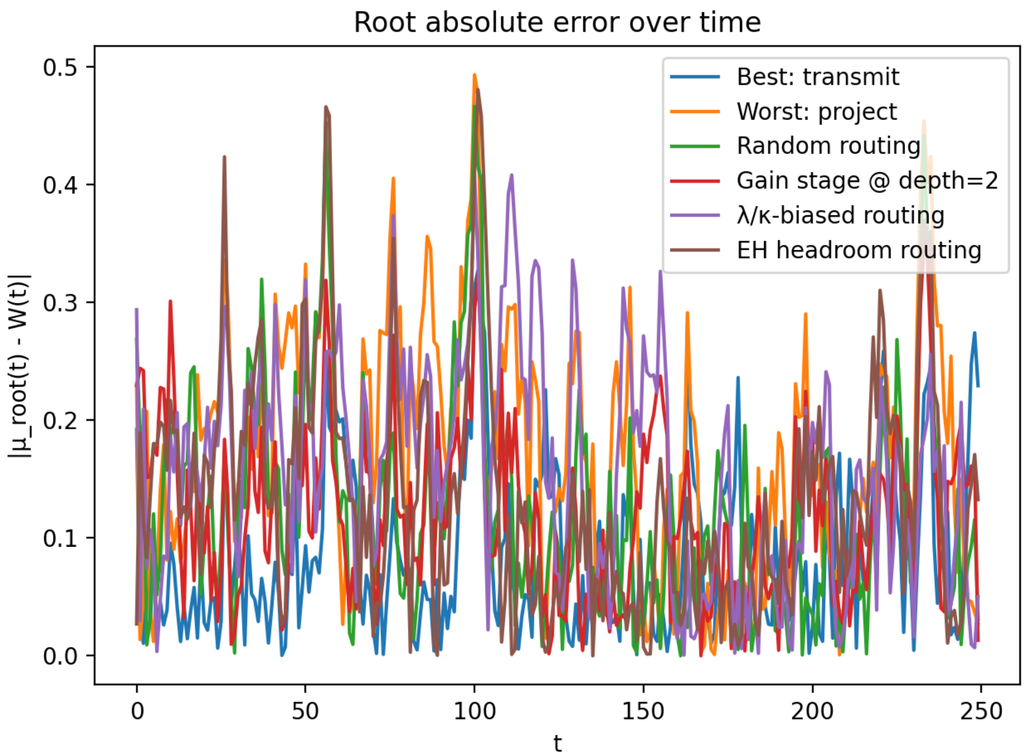

Figure 3: Root absolute error $e(t) = |\mu_{\text{root}}(t) - W(t)|$ over time for all scenarios.

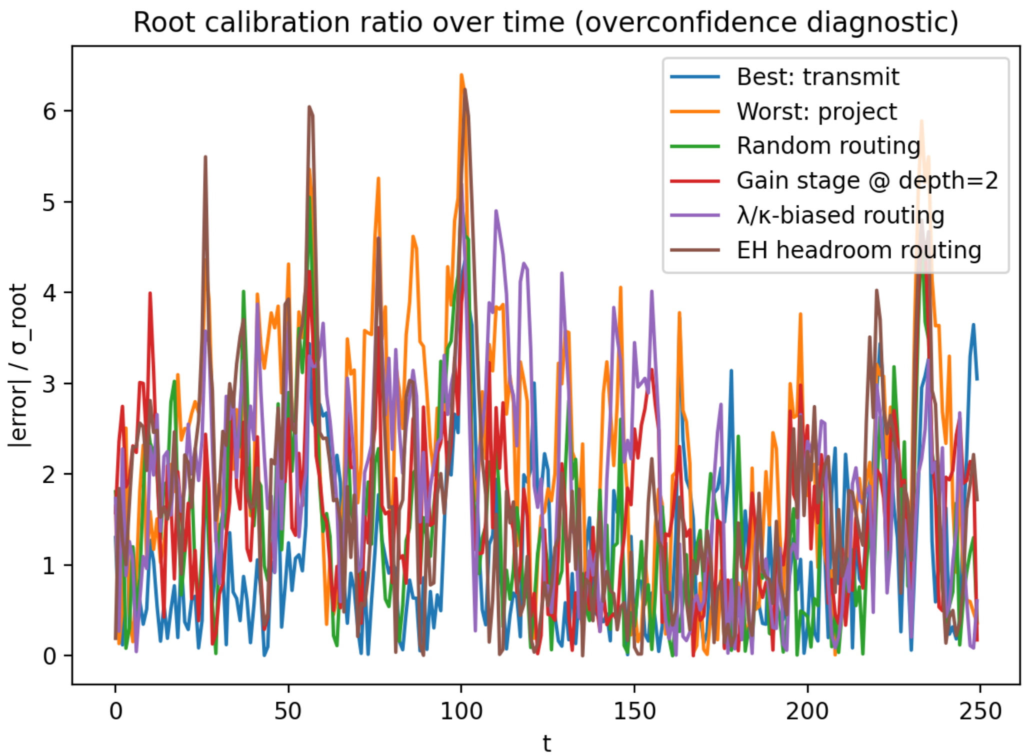

Figure 4: Root calibration ratio $r(t) = e(t)/\sigma_{\text{root}}(t)$ over time. Values above 2 indicate overconfidence episodes.

A consistent visual signature is that when projection is nontrivial, $r(t)$ exhibits repeated spikes above 2. Those spikes reflect the gap between two dynamics: projection can reduce expressed uncertainty quickly, while the underlying error may persist or even increase under information loss.

12 Interpretation as a control problem

The simulation can be read as a control problem with partial observability. The world evolves, nodes maintain belief states, and routing actions determine which information reaches the root. Structural distance loads the system by degrading observations and channels. Projection is a control action that produces a crisp artifact, but it does so by discarding distinctions and by compressing stated uncertainty.

On this reading, **pressure** is a decision-time switch. When pressure is low relative to headroom, the system can afford to keep uncertainty live, transmit richer signals, or absorb locally. When pressure exceeds headroom, the system converts uncertainty into a decisive artifact and routes it upward. Because projection is non-invertible in the model, repeated high-pressure episodes can gradually produce a top layer that is confident without being well calibrated. Since the world state evolves over time, routing choices affect not only instantaneous accuracy, but also the rate at which the system can track change; projection therefore alters learning dynamics across time, not merely one-step estimates.

13 Relation to SDT, TST, RST, and LBNT

The routing model can be situated within the broader structural framework.

13.1 Structural distance loading (SDT)

In this model, λ_i is a scalar that increases noise in both local observation and communication. That is a reduced-form choice. In SDT, structural distance is developed as a joint degradation of reciprocal visibility, signal fidelity, and consequence coupling (Parten, 2025a). The present model compresses those components into one loading parameter and then focuses on the routing question: given distance-dependent degradation, which routing actions amplify or mitigate the drift burden at the top?

13.2 Non-invertible compression

Projection is many-to-one. Once distinctions are erased, downstream aggregation cannot recover them without fresh independent signal. This is the informational core of why legibility artifacts can be stabilizing for action while degrading truth-tracking.

13.3 Uncertainty collapse and corrigibility (TST)

Projection also alters a system's responsiveness. By shrinking expressed variance, the model makes the upper layer behave as if it knows more than it does. Under distance, that can reduce sensitivity to disconfirming evidence. In TST terms, reduced corrigibility is one mechanism through which misalignment becomes persistent under load (Parten, 2025b).

13.4 Pressure as a loading term (TST)

TST formalizes drift at a position using a local multiplier of the form $M_i = 1 + \alpha\lambda_i - \beta\kappa_i$ (Parten, 2025b). The routing model does not reuse that drift recurrence, but it is compatible with the same control intuition: distance loads a position and capacity stabilizes it. The headroom extension suggests a refinement in which pressure acts as an additional time-varying loading term,

$$M_i(t) = 1 + \alpha\lambda_i - \beta\kappa_i + \gamma p_i(t), \tag{21}$$

with $\gamma > 0$. The routing model operationalizes the same idea by directly comparing $p_i(t)$ to $H_i(t)$ to decide when uncertainty is collapsed into a projected artifact.

13.5 From nested routing to coupled ecologies (RST)

This paper models one nested system with internal routing. RST generalizes to interacting ecologies of systems and asks when misalignment propagates across coupling pathways (Parten, 2026a). One extension is to treat subtrees as systems in the RST sense and study how routing policies shape effective coupling between modules.

13.6 Structural decisiveness and load-bearing nodes (LBNT)

Not every node's local behavior affects the root equally. Some positions are structurally decisive because they sit near bottlenecks or because their projections gate what the upper layers can see. LBNT develops sensitivity-based tools for identifying such load-bearing positions in coupled drift models (Parten, 2026b). The routing model provides a mechanism-level way to generate bottlenecks and to study where projection-induced information loss accumulates.

14 Limitations and extensions

Several limitations point to clear extensions.

- **Scalar world state.** Real representational spaces are high-dimensional; a single scalar can hide topic dependence and geometry.
- **Gaussian belief summaries.** The mean-variance representation keeps inference transparent but omits multimodality and structured uncertainty.
- **Simple fusion and independence assumptions.** Signals are treated as conditionally independent with clean precision-weighted fusion. This isolates routing effects but is not a general inference theory.
- **Stylized projection.** Quantization and variance shrinkage capture non-invertibility and uncertainty collapse, but real artifacts differ in bias, incentive structure, and cost. A natural extension is to introduce multiple projection operators with explicit distortion profiles.
- **Exogenous pressure in EH.** Here pressure fluctuates around a baseline. Many systems generate pressure endogenously, for example when accumulated error triggers accountability cycles. Endogenizing pressure would allow feedback loops between drift, pressure, and projection.

15 Conclusion

This paper isolates a structural failure mode in hierarchical signal routing. When information must traverse distance and decision pressure intermittently rewards legible, decisive artifacts, systems can collapse expressed uncertainty faster than they reduce underlying error. What emerges is not merely noisy inference, but a distinctive pattern of confident error: beliefs that are both incorrect and insufficiently corrigible.

At the core of this dynamic lies an asymmetry. High-fidelity transmission preserves distinctions while accumulating channel degradation, whereas lossy projection trades those distinctions for decisiveness through irreversible compression. Once projection occurs under distance, downstream aggregation cannot reconstruct what was erased without renewed contact with ground conditions. Uncertainty collapse therefore introduces a one-way transformation in the informational dynamics of the system.

Pressure determines when this transformation is most likely. Rather than treating projection as a constant preference, the model represents it as a threshold response that activates when epistemic headroom is exceeded. Under such conditions, systems tend to summarize and escalate information precisely when their capacity to validate those summaries is weakest. Repeated pressure-triggered projection can thus generate persistent overconfidence even in the absence of deception, strategic behavior, or incentive misalignment.

Epistemic headroom routing illustrates a mixed regime that partially mitigates these effects by permitting transmission and absorption when slack is available. That mitigation, however, does not eliminate the underlying tradeoff. As long as distance degrades feedback and pressure intermittently demands legibility, some degree of information loss and uncertainty collapse remains unavoidable.

A broader implication follows. Projection is neither intrinsically pathological nor costless; it is a control action whose structural costs depend on timing and location within a hierarchy. In large, layered systems, where and when uncertainty is collapsed can matter as much as how often it occurs. Making those costs explicit is a prerequisite for designing institutions that remain responsive to error under load.

A Default parameters

Tables 3 and 4 list the default simulation parameters used for the illustrative run.

Parameter	Value
Tree branching k	3
Tree depth D	4
Timesteps T	250
World noise σ_W	0.05
Observation base $\sigma_{\text{obs},0}$	0.25
Observation distance scale s_{obs}	2.0
Edge base $\sigma_{\text{edge},0}$	0.10
Edge distance scale s_{edge}	2.0
Capacity scale s_κ	1.0
Projection bin Δ	0.50
Projection shrink ρ	0.08
Projection floor $v_{\min}$	0.01
$\lambda_{\text{root}} \to \lambda_{\text{leaf}}$	$0.85 \to 0.15$
$\kappa_{\text{root}} \to \kappa_{\text{leaf}}$	$0.25 \to 0.75$
$\bar{p}_{\text{root}} \to \bar{p}_{\text{leaf}}$	$0.80 \to 0.20$
Depth jitter (std)	0.05
Random seeds (world / params)	2 / 7

Table 3: Default parameter values for the base signal-routing simulation.

EH parameter	Value
Headroom bias h_0	0.50
Headroom distance weight a_λ	0.70
Headroom capacity weight b_κ	0.70
Headroom noise std σ_H	0.05
Pressure noise std σ_p	0.15
Absorption slack threshold s_{abs}	0.20
Absorption capacity threshold $\kappa_{\min}$	0.60
Absorption distance threshold $\lambda_{\max}$	0.60

Table 4: Default parameter values for epistemic headroom routing.

B Reproducibility

The simulation is implemented in Python and produces per-scenario time series, action counts, summary statistics, and diagnostic plots. Code is available upon request (see the title page contact address).

References

Justin Parten. *Structural Distance Theory: A Mechanism-Level Account of Epistemic Drift in Complex Systems*. Preprint, Version 4 (uploaded December 7, 2025). Available via PhilArchive: `https://philpapers.org/archive/PARVDT-2.pdf`.

Justin Parten. *Truth Stabilization Theory: A Minimal Drift Dynamic with Networked Simulation*. Preprint, December 2025. Available via PhilArchive: `https://philpapers.org/archive/PARTST-8.pdf`.

Justin Parten. *Relational Systems Theory: Drift Dynamics and Thresholds in Coupled Relational Ecologies*. Preprint, January 8, 2026. Available via PhilArchive: `https://philpapers.org/archive/PARRST-2.pdf`.

Justin Parten. *Load-Bearing Node Theory (LBNT): Identifying Structurally Decisive Nodes in Drift-Amplifying Systems*. Preprint, Version 1.01, January 20, 2026. Available via PhilArchive: `https://philpapers.org/archive/PARLNT.pdf`.