A Toy Model of Symbolic Reuse and Source-Relative Fidelity Change

Trace, Reuse, Source Re-alignment, and Later-Form Comparison

Kostiantyn Osmolovskyi
Independent Researcher, Ukraine

Version 2.0, 2026

Abstract

This paper presents a minimal toy model of symbolic reuse in which initial symbolic capture, later-form trace, reuse, and later source-relative fidelity remain distinct. The initial state is imported from a fixed symbolic-capture model only when a positive symbolic form, positive capture application, applicable fidelity profile, source descriptor record, and profile-specific scalar are all established. At each later step the model computes a local reuse-load proposal, generates a candidate later symbolic form under the same fixed codebook, qualifies that form, establishes an explicit trace from the initial capture through an unbroken committed-form chain, and evaluates a separate reuse relation. The proposed state is committed only after positive reuse. The model does not create a fresh symbolic-capture application at each later endpoint. Instead, the committed later form is compared with the original source descriptor record under a local source-reference policy, yielding independent RefPres, RefRed, and RefDist assessments and a scalar distance D_s^{ref}.

Eighteen pre-specified structural consistency tests pass. Four deterministic 201×201 surfaces show how reuse intensity, source re-alignment, recurrence coefficients, and initial source geometry affect the local trajectory proxy and later source-relative fidelity. Three pre-specified Monte Carlo strata of 20,000 episodes yield eligible initial exports in 47.84%, 72.95%, and 85.52% of generated cases; every eligible case then completes the standard twenty-step reuse route by construction. Among final reference-defined cases, RefDist = 1 occurs in 10.74%, 10.73%, and 10.85%. Both directional changes occur: some initially profile-distortive captures end without later RefDist, while some initially non-distortive captures end with RefDist = 1. These frequencies characterize only the declared equations, codebook, source distributions, and standard positive trace/reuse policy. They are not empirical prevalence estimates, memory-success probabilities, semantic truth judgments, or universal laws of symbolic reuse.

Keywords: symbolic reuse; symbolic trace; source-relative fidelity; symbolic form; preservation; reduction; distortion; re-alignment; toy model; cognitive architecture.

1 Analytic Target and Scope

Symbolic capture and later symbolic reuse answer different architectural questions. A positive symbolic-capture application establishes a source–form relation at one stage, but it does not by itself establish later availability, persistence, a later symbolic endpoint, or reuse. The symbolic-capture architecture therefore keeps positive capture separate from later trace and reuse relations [1]. The present model studies one explicit temporal route through those later relations.

The starting point is a fully typed initial export from the companion symbolic-fidelity toy model [4]. That export retains the original manifestation-source provenance, the positively qualified symbolic form, the positive symbolic-capture application, the fixed fidelity profile, the initial profile-relative statuses SymPres, SymRed, and SymDist, and the local scalar $D_0^{\mathrm{sym,toy},p}$. The

present model does not reinterpret that scalar as a symbolic form, a memory trace, or a general distortion state.

For each later reuse step, the architecture is

$$\begin{aligned}&\text{eligible initial capture/fidelity state}\\ &\quad\longrightarrow \text{later-form proposal}\\ &\quad\longrightarrow \mathrm{SymForm} = 1\\ &\quad\longrightarrow \mathrm{SymTrace} = 1\\ &\quad\longrightarrow \mathrm{SymReuse} = 1\\ &\quad\longrightarrow \text{committed later symbolic endpoint}\\ &\quad\longrightarrow (\mathrm{RefPres}, \mathrm{RefRed}, \mathrm{RefDist}, D^{\mathrm{ref}}).\end{aligned}$$

The trace and reuse applications preserve exact source/form provenance. The final reference statuses are local assessments against the original source descriptor record; they are not fresh symbolic-capture SymPres, SymRed, or SymDist values because no fresh later SymCap application is asserted.

The model is a computational toy realization rather than a measurement-validation study [2]. It does not provide a biological memory mechanism, a theory of semantic truth, a belief-update model, a theory of conscious recall, or an empirical estimate of symbolic-reuse frequency.

2 Initial State and Temporal Route

2.1 Eligible initial state

Every case begins by executing the symbolic-fidelity model, which in turn imports the perceptual-stabilization model rather than reimplementing its equations [3, 4]. Let E_{init} denote the selected initial state. The route is eligible only if the symbolic-fidelity-model result contains all of the following with mutually consistent provenance:

- a positive symbolic-form application under the fixed Q21 symbolic-form context;
- a positive symbolic-capture application tied to the original manifestation application;
- an applicable fidelity profile and profile-qualified fidelity application;
- the original source descriptor record (a, b, c);
- defined initial SymPres, SymRed, SymDist, and $D_0^{\mathrm{sym,toy},p}$;
- a decodable captured symbolic form.

If these conditions are not established, the selected reuse route is not instantiated or remains unresolved. Such a case is not recoded as zero reuse, zero distortion, or absence of symbolic capacity.

2.2 Later symbolic-form endpoint

At step $s \geq 1$, the model proposes a candidate later form and evaluates it with the exact imported Q21 symbolic-form criterion. A candidate becomes a positive later symbolic-form application only when $\mathrm{SymForm} = 1$ in the fixed form context. The trace relation then connects the initial positive capture to this later positive form while retaining the immediately preceding committed positive-form identifier.

The local trace context is `TRACE-ROOT-PROVENANCE-v1`. A positive trace requires the exact root capture, expected previous committed form, expected later form, trajectory identifier, step identifier, fixed form context, and the exact proposal provenance linking the previous candidate to that later endpoint under the declared generator. A detached later form, detached previous endpoint, detached proposal parent, or wrong proposal generator receives a negative trace rather than a positive one.

The local reuse context is `REUSE-STEP-v1`. Positive reuse requires a positive trace application together with the exact current proposal record, its expected previous-candidate parent, the candidate endpoint, trajectory/step identifiers, the source-reference record belonging to the same root and step, and a positive reuse gate. Hence

$$\mathrm{SymTrace} = 1 \not\Rightarrow \mathrm{SymReuse} = 1.$$

The model contains explicit negative and undefined trace/reuse cases in its structural consistency tests.

2.3 No fresh later symbolic capture

The committed later endpoint is a positive symbolic form plus a positive trace/reuse record. The model does not infer

$$\mathrm{SymReuse} = 1 \Rightarrow \text{fresh later } \mathrm{SymCap} = 1.$$

This restriction is central. Without a new source–form capture criterion at the later endpoint, the model has no basis for relabeling its later local comparison as a new symbolic-capture fidelity application.

3 Local Reuse Dynamics and Source Re-alignment

3.1 Reuse-load proxy

The temporal scalar is a model-local deformation-load proxy

$$q_s^{\mathrm{reuse,toy}} \in [0, 1].$$

Initialization is permitted only on an eligible initial export:

$$q_0^{\mathrm{reuse,toy}} := D_0^{\mathrm{sym,toy},p}.$$

This is an initialization bridge, not an identity between formal objects. In particular,

$$q_s^{\mathrm{reuse,toy}} \neq \mathrm{SymDist}, \qquad q_s^{\mathrm{reuse,toy}} \neq \mathrm{RefDist}, \qquad q_s^{\mathrm{reuse,toy}} \neq D_s^{\mathrm{ref}}$$

by definition.

For a declared reuse-intensity coordinate $U_s \in [0, 1]$, the proposed next load is

$$q_{s+1}^{\mathrm{prop}} = \mathrm{clip}_{[0,1]}\left(q_s + \kappa U_s q_s - \delta G_s^{\mathrm{eff}}\right),$$

with baseline

$$\kappa = 0.25, \qquad \delta = 0.20, \qquad S = 20.$$

The recurrence produces only a proposal. The next load is committed to the trajectory only if the generated later form receives positive symbolic-form qualification, positive trace, and positive reuse.

3.2 Source-reference basis

A local source-reference status $\mathrm{GroundRef}_s$ distinguishes an established source-realignment basis from its resolved absence and from undefinedness. A positive basis retains the exact original manifestation-application identifier, the original source-descriptor record, the root capture identifier, and context `GROUND-REF-v1`. Only on this positive basis is a declared $G_s \in [0,1]$ interpreted as source-realignment strength in the toy recurrence.

The effective contribution is

$$G_s^{\mathrm{eff}} = \begin{cases} G_s, & \mathrm{GroundRef}_s = 1, \\ 0, & \mathrm{GroundRef}_s = 0, \\ \text{undefined}, & \mathrm{GroundRef}_s \text{ undefined}. \end{cases}$$

Thus a positive reference basis with $G_s = 0$ is provenance-distinct from a negative basis even though both contribute zero to the recurrence. Neither positive grounding nor positive G_s means semantic truth access, successful real-world correction, or a renewed manifestation application.

4 Later-Form Proposal and Local Reference Fidelity

Let the previous committed symbolic form decode to (a'_s, b'_s, c'_s), and define

$$\Delta q_s = q_{s+1}^{\mathrm{prop}} - q_s.$$

With $\eta = 1$, the pre-quantized protected pair is

$$a^*_{s+1} = \mathrm{clip}_{[0,1]}(a'_s + \eta \Delta q_s), \qquad b^*_{s+1} = \mathrm{clip}_{[0,1]}(b'_s - \eta \Delta q_s).$$

The third symbolic descriptor is carried unchanged:

$$c'_{s+1} = c'_s,$$

including the omission symbol $\varnothing$. The proposed form is then quantized through the Q21-v1 codebook inherited from the symbolic-fidelity model.

After positive reuse, only the form named as committed by that positive reuse application is eligible for comparison with the original source record (a, b, c) under policy `REF-REL3-v1`. A detached form or detached reuse application leaves the local reference assessment undefined. Define

$$\mathrm{RefPres}_s = 1$$

iff

$$b'_s > a'_s \quad \text{and} \quad \left|(b'_s - a'_s) - (b - a)\right| \le 0.20,$$

$$\mathrm{RefRed}_s = 1 \quad \Longleftrightarrow \quad c'_s = \varnothing,$$

and

$$\mathrm{RefDist}_s = 1 \quad \Longleftrightarrow \quad b'_s \le a'_s.$$

The local protected-pair distance is

$$D_s^{\mathrm{ref}} = \frac{1}{2}\left(|a'_s - a| + |b'_s - b|\right).$$

These outputs remain distinct. RefPres = 0 does not force RefDist = 1: a later form may preserve the positive ordering while changing its gap beyond the preservation tolerance. Likewise, $D^{\mathrm{ref}} = 0$ does not force RefRed = 0, because the protected-pair distance does not inspect the third descriptor. Finally, RefRed is not Structural Compression [5].

5 Definedness and Structural Consistency Tests

The implementation preserves separate states for initial-route eligibility, source-reference definedness, candidate-form presence, SymForm, SymTrace, SymReuse, commit status, and later reference assessment. Undefined status is never replaced with a defined negative value or numeric zero.

Eighteen structural consistency tests cover route eligibility, definedness, trace/reuse separation, source-reference status, and non-identification of key coordinates. All eighteen pass. Among the diagnostic cases:

- CT-1 keeps q_0, trace, reuse, and commit absent when no eligible initial capture/fidelity state exists.
- CT-7 yields SymTrace $= 1$ but SymReuse $= 0$ under reuse-provenance mismatch.
- CT-10 yields $q_1 = 0$ together with $(\text{RefPres}, \text{RefRed}, \text{RefDist}) = (1, 1, 0)$ and $D_1^{\text{ref}} = 0$, showing that perfect protected-pair realignment does not undo an omitted third descriptor.
- CT-11 begins with SymDist $= 1$ but yields later RefDist $= 0$ after strong source re-alignment.
- CT-12 begins with SymDist $= 0$ but yields later RefDist $= 1$ after four high-reuse, weak-realignment steps.
- CT-17 confirms that positive reuse creates a later positive symbolic form but no fresh later symbolic-capture application.

The public implementation records trajectory-level and step-level states with explicit source descriptors, symbolic-form qualification, trace, reuse, commitment, and source-reference fields. The reported tables and figures are computed from these typed records. This representation preserves the model distinctions used in the analysis without introducing additional state variables.

6 Simulation Design

All deterministic surfaces contain 201×201 endpoint-inclusive grid points. The pre-specified simulation families are:

S1: state-level reuse/re-alignment
$U \times G$ for an initially non-distortive symbolic-fidelity state, with baseline recurrence coefficients.

S2: initially distortive root
$U \times G$ for an initially distortive symbolic-fidelity state with SymDist $= 1$.

S3: recurrence coefficients
$\kappa \times \delta$ at fixed $U = 0.65$ and $G = 0.35$ using the initially non-distortive symbolic-fidelity state.

S4: initial source geometry
source gap $\times U$ with the ambiguous-signal perceptual episode, $a = 0.10$, $c = 0.50$, $G = 0.05$, and exact symbolic-fidelity-model execution independently at every grid point.

The three Monte Carlo strata use 20,000 generated episodes each with 1, 2, or 3 perceptual blocks and fixed seeds. The exact sampling order is ℓ, r, blockwise M, P, and signal support, followed by source a, source gap, c, U, and G. The source construction guarantees $0.10 \leq a < b \leq 0.90$. No missingness, malformed form, trace failure, reuse failure, perceptual-model parameter randomization, symbolic-fidelity-model parameter randomization, or recurrence-parameter randomization is

injected. Those definedness cases belong to the structural consistency tests rather than to arbitrary probability assignments.

Consequently, Monte Carlo initial-export fractions characterize how often the declared sampled episodes satisfy the specified initial route. They are not empirical probabilities. Conditional on an eligible initial export, the standard Monte Carlo path deliberately supplies positive trace/reuse provenance at all twenty steps, so completion is one by construction and is not a reuse-success estimate.

7 Results

7.1 Deterministic state and coefficient surfaces

S1 begins from an initially non-distortive, reductive capture with $q_0 = 0.15$. After twenty standard reuse steps, RefPres $= 1$ in 84.85% of the grid and RefDist $= 1$ in 7.84%. The final mean source-reference distance is 0.078. Relative to the initial protected-pair distance, 34,278 cells decrease, 2,862 are unchanged, and 3,261 increase. Figure 1 shows that reuse intensity and source re-alignment act on a continuous load/reference geometry rather than on one exclusive symbolic mode.

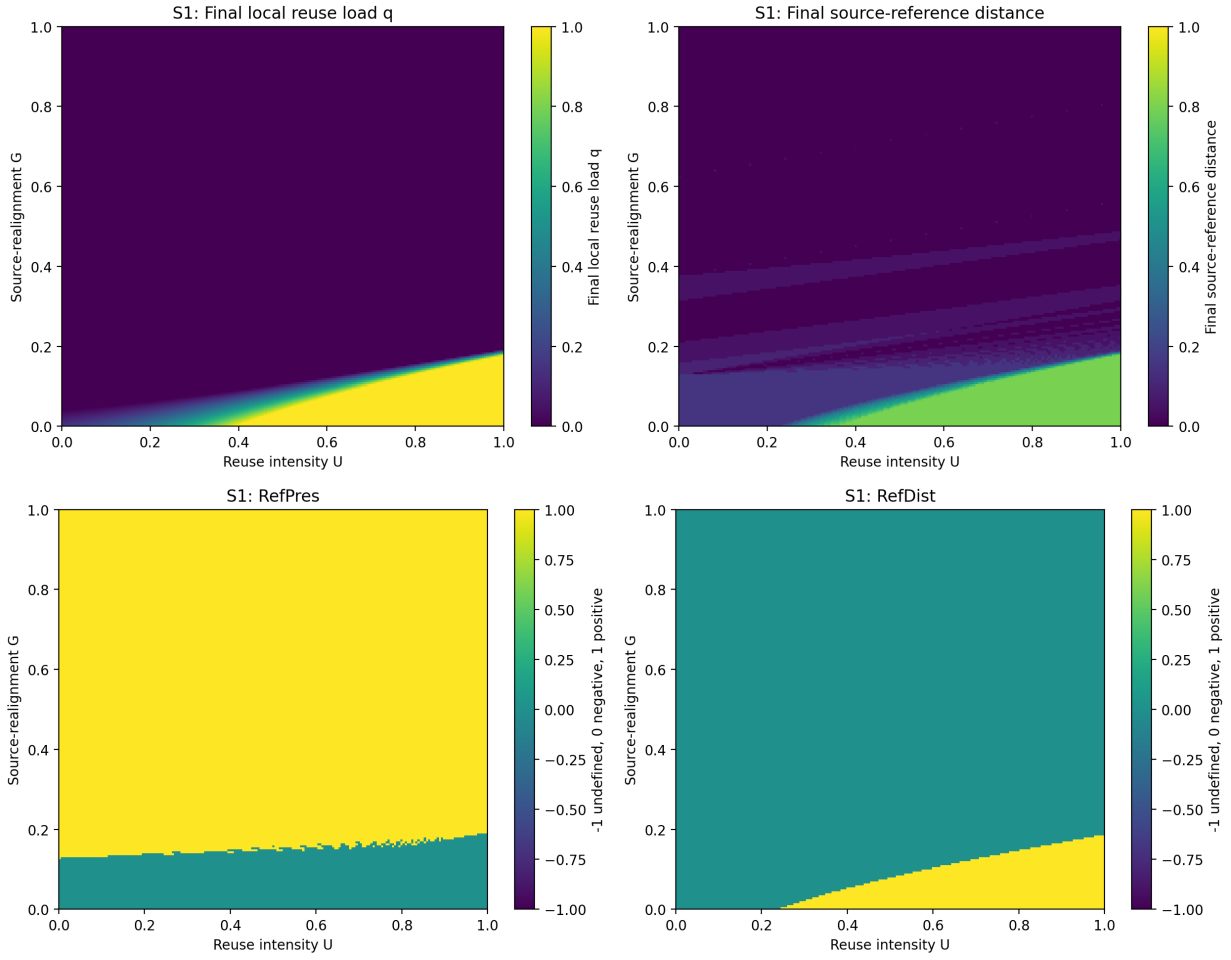


Figure 1: S1, initially non-distortive symbolic-fidelity state. Final local reuse load q, final source-reference distance D^{ref}, and the separate RefPres and RefDist surfaces over reuse intensity U and source-realignment strength G. The storage-order axis correction is applied in the renderer: horizontal is U, vertical is G.

S2 starts from an initially profile-distortive symbolic-fidelity state with $q_0 = 0.30$. Nevertheless, strong enough source re-alignment can end outside the later RefDist region: final RefDist = 1 occupies 21.90% of the grid, while final RefPres = 1 occupies 76.28%. Thus an initial SymDist = 1 does not force a later local RefDist = 1 trajectory endpoint.

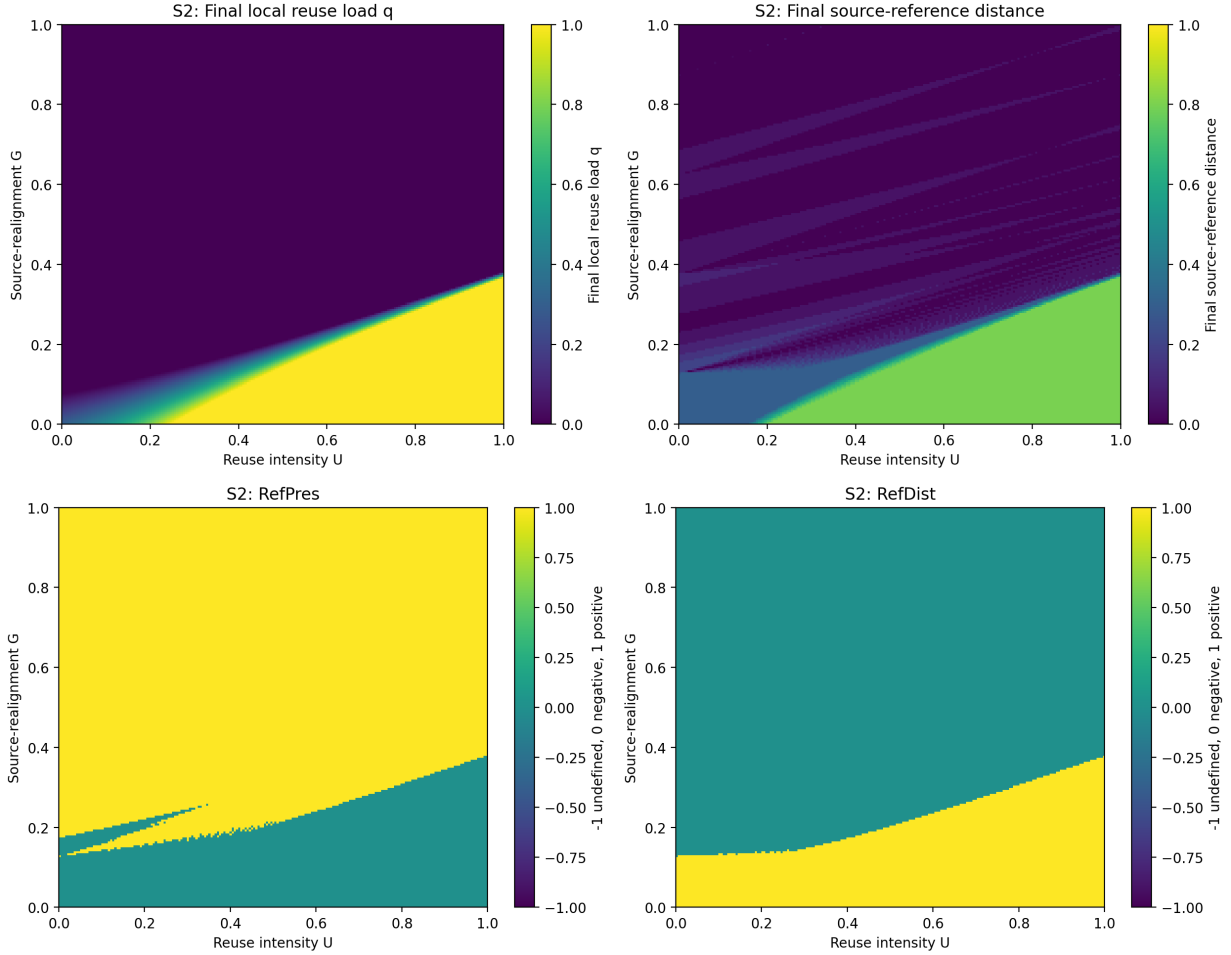


Figure 2: S2, initially profile-distortive symbolic-fidelity state. The same initial capture can end in locally preserved, distorted, or neither-preserved-nor-distorted source-relative regions depending on the declared $U \times G$ path.

S3 varies only the recurrence coefficients at fixed episode coordinates. Final RefDist = 1 occupies 15.64% of the grid and RefPres = 1 occupies 76.25%. These boundaries are properties of the local recurrence and quantized form update, not parameters supplied by the symbolic-capture construction.

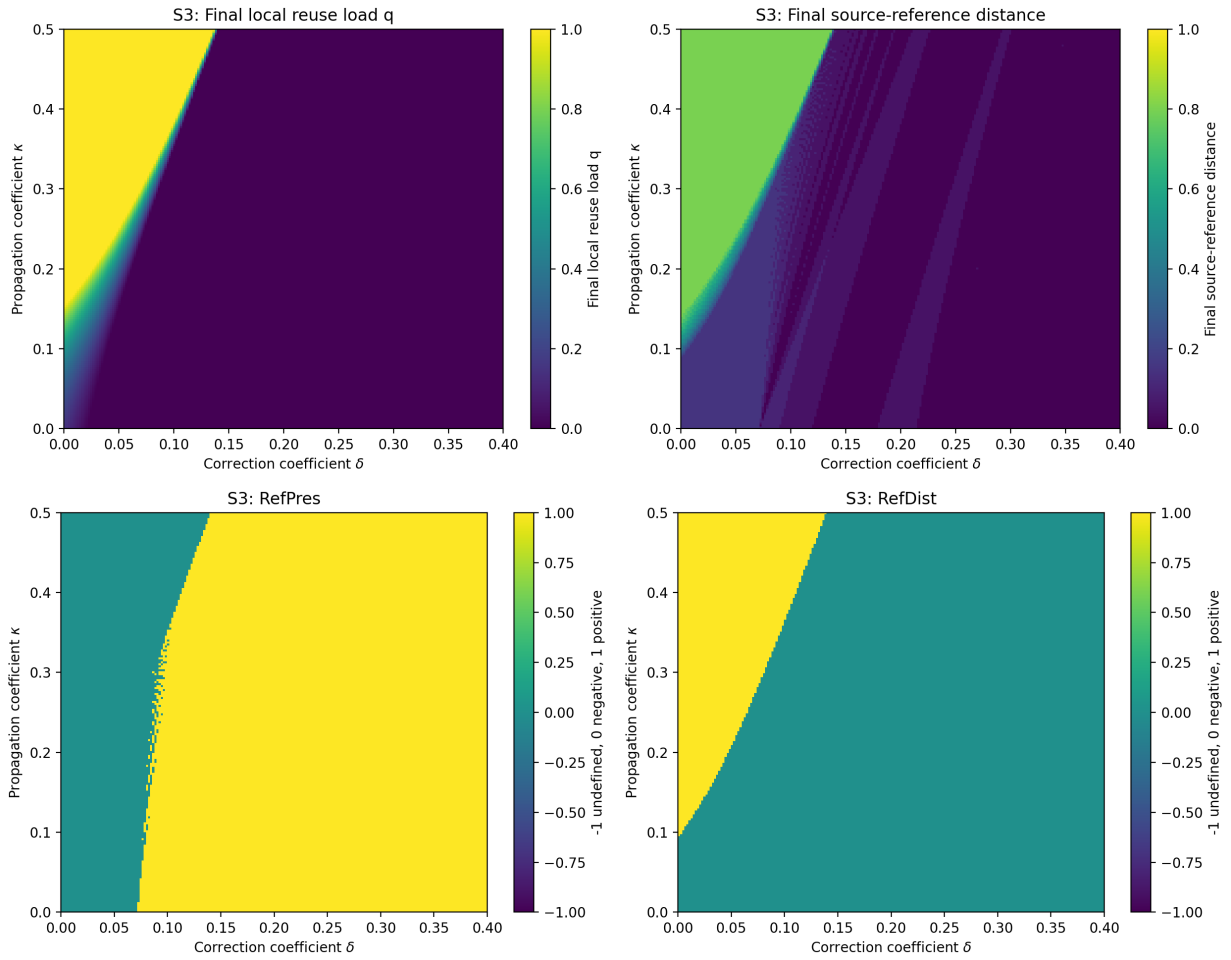


Figure 3: S3 coefficient surface. Horizontal is the source-realignment coefficient δ and vertical is the reuse-amplification coefficient κ. The figure separates the scalar trajectory state from the later reference classifiers.

S4 recomputes the complete symbolic-fidelity initial state at every source-gap value before applying the present model reuse dynamics. No cell is assigned an arbitrary initial scalar. Across this surface final RefDist $= 1$ occurs in 65.17% of cells, while RefPres $= 1$ occurs in 0.00%. The zero preservation rate is specific to the specified weak-realignment/high-reuse geometry of S4; it is not a universal property of reuse. The stepped boundaries reflect Q21 quantization.

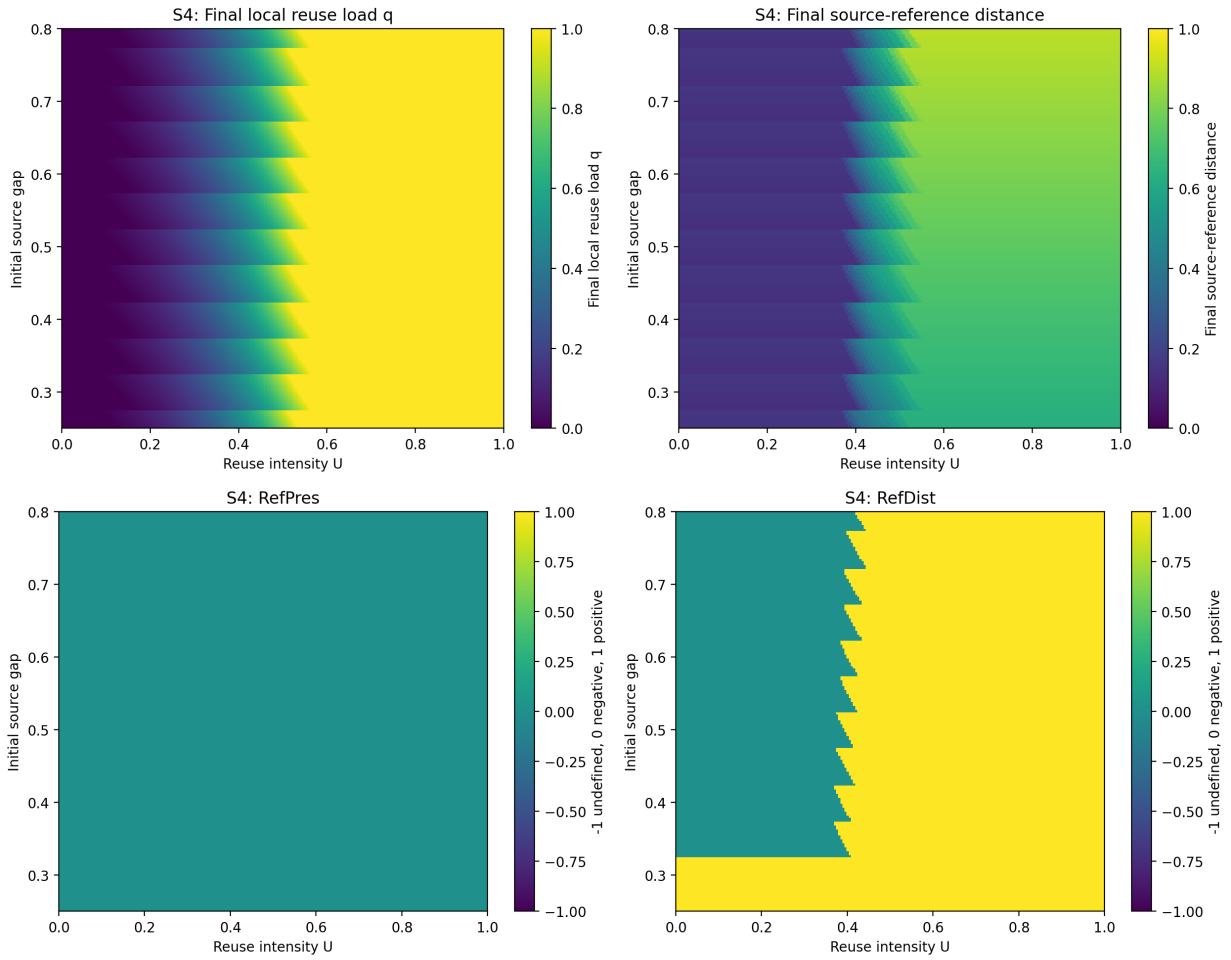


Figure 4: S4 source-geometry surface. Each cell obtains its initial state through exact symbolic-fidelity-model execution. The final load, source-reference distance, preservation, and distortion surfaces therefore remain tied to a valid typed root rather than to an independently sampled q_0.

7.2 Monte Carlo strata

Table 1 reports only the declared sampling laws. Initial eligibility rises from 47.84% in the one-block stratum to 85.52% in the three-block stratum because the specified initial route has more sampled block opportunities; this is not a prevalence claim about real systems. Every eligible root completes the standard twenty-step trace/reuse path because the Monte Carlo design fixes valid later-form, trace, and reuse provenance rather than sampling failure probabilities.

Table 1: Monte Carlo results. Reference-status proportions are conditioned on the final reference-defined denominator.

Family	N	Init. elig.	Complete	RefPres +	RefRed +	RefDist +	mean D^{ref}
MC-R1	20,000	9,569	9,569	84.13%	50.62%	10.74%	0.099
MC-R2	20,000	14,589	14,589	84.15%	46.47%	10.73%	0.098
MC-R3	20,000	17,104	17,104	84.51%	41.62%	10.85%	0.098

Final RefDist is close to eleven percent across the three declared eligible-root populations: 10.74%, 10.73%, and 10.85%. Mean final D^{ref} is approximately 0.099, 0.098, and 0.098, respectively. In roughly 81% of reference-defined cases in each stratum, the final protected-pair distance is lower than the initial symbolic-fidelity protected-pair distance under the same source descriptors. That

result follows from the particular uniform U/G laws and recurrence coefficients; it is not evidence that reuse generally improves fidelity.

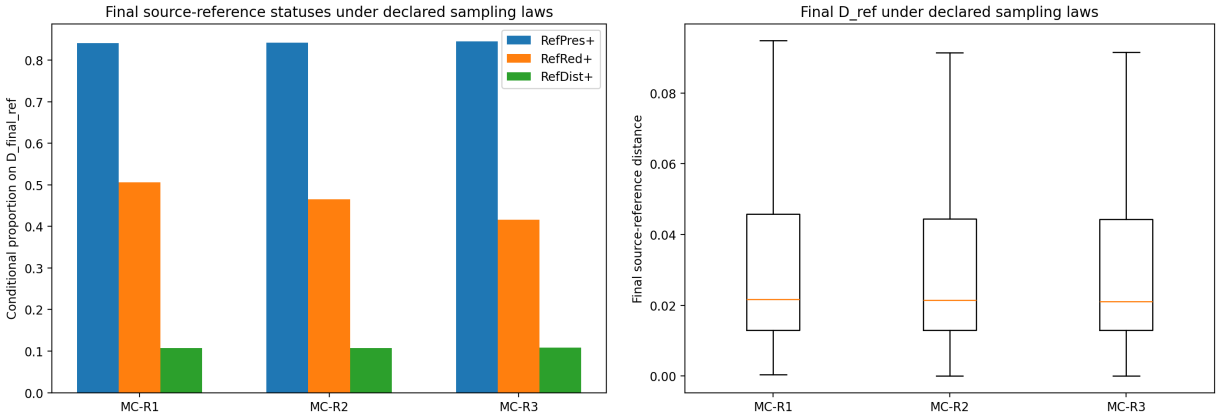


Figure 5: Monte Carlo summaries under the declared sampling laws. Left: final local reference-status proportions, which are independent coordinates rather than an exclusive mode distribution. Right: final D^{ref} distributions for final-reference-defined cases.

The temporal distinction is visible directly in the joint tables. Among initially SymDist $= 1$ eligible cases, 80.88%, 81.20%, and 80.24% end with RefDist $= 0$. Conversely, among initially SymDist $= 0$ eligible cases, 8.29%, 8.13%, and 7.76% end with RefDist $= 1$. These are not transitions of a single status variable from SymDist to RefDist; they compare an initial profile-relative symbolic-fidelity assessment with a later source-reference assessment under the same original source record.

8 Interpretation

The model establishes four separations that are easy to lose in a scalar recurrence.

First, symbolic reuse is not produced by recurrence alone. The recurrence proposes q_{s+1}, but the state is committed only after positive later-form qualification, positive trace, and positive reuse. A scalar sequence without those applications would therefore be an arithmetic trajectory, not a typed symbolic-reuse trajectory.

Second, the initial symbolic-fidelity statuses are not propagated as state labels. The later endpoint receives only local RefPres, RefRed, and RefDist under the source-reference policy. The model therefore permits both an initially SymDist $= 1$ root followed by RefDist $= 0$ and an initially SymDist $= 0$ root followed by RefDist $= 1$ without treating them as one Boolean state propagated through time.

Third, q and D^{ref} are non-identical coordinates. In every eligible Monte Carlo trajectory the final values differ under the declared realization, and deterministic grids also contain many such cases. The recurrence variable controls the proposed deformation of the previous form, whereas D^{ref} evaluates the committed form against the original source. Quantization and repeated updates make these relations nontrivial.

Fourth, reduction is deliberately static in the present model. The third Q21 payload is carried exactly at every step, so final RefRed equals the initial omission status by construction in every standard trajectory. The equality of initial and final reduction proportions is therefore a design invariant, not a discovered robustness result. A future model that allows detail restoration or further omission would require a separately declared update rule and would constitute a different simulation.

The reference classifiers are also not exhaustive complements. RefPres $= 0$ and RefDist $= 0$ occur

together whenever the later pair remains positively ordered but its gap falls outside the preservation tolerance. Across the pre-specified families this neither-preserved-nor-distorted region is nonempty. The result prevents replacing the independent assessments with a single preservative/distortive mode variable.

9 Limitations

The model is intentionally narrow. The Q21 codebook, source descriptors, recurrence coefficients, reuse-intensity coordinate, source-realignment coordinate, reference policy, and twenty-step horizon are realization-specific. They are not canonical scales or empirical parameters.

The Monte Carlo studies sample synthetic episode geometry, not populations of people, organisms, artificial agents, memories, beliefs, or communications. Initial eligibility proportions are properties of the imported symbolic-fidelity route under the specified distributions. Conditional completion is fixed by the standard positive trace/reuse policy and must not be interpreted as memory reliability or reuse success.

The source-reference basis is also deliberately modest. $G > 0$ means only that the declared local recurrence receives a source-realignment contribution after a valid reference record has been established. It does not mean access to objective truth, successful learning, corrected belief, or renewed environmental evidence.

No fresh later SymCap is constructed. Consequently RefPres, RefRed, and RefDist are not new symbolic-capture fidelity statuses. They are local comparisons of a trace/reuse-qualified later form with the original source descriptor record. Likewise, RefRed = 1 is not Structural Compression, and RefDist = 1 is not false belief, deception, hallucination, pathology, or semantic falsity.

Finally, the present model keeps the omission dimension fixed. It therefore cannot study the emergence, recovery, or propagation of symbolic reduction itself. Its temporal variation is restricted to the protected pair and the resulting local preservation/distortion geometry.

10 Conclusion

A symbolic-reuse toy model can represent temporal change without treating an initial symbolic-capture status or scalar as a self-propagating object. The present architecture begins from an eligible typed capture/fidelity state, constructs each later symbolic endpoint explicitly, establishes trace and reuse as separate relations, commits the local recurrence only after positive reuse, and compares the committed later form with the original source under a separate reference policy.

Under this realization, reuse intensity and source re-alignment can move the same initial symbolic organization toward or away from later source-relative distortion. An initially profile-distortive capture need not remain locally distorted after later reuse, and an initially non-distortive capture can end in a locally distorted later form. These outcomes do not represent transitions of one universal distortion state: initial SymDist and later RefDist belong to different typed assessments. The local recurrence proxy q is likewise distinct from the later source-reference distance D^{ref}.

The resulting model therefore supplies a controlled temporal interface among positive symbolic capture, later symbolic form, trace, reuse, and source-relative fidelity change while preserving the distinction between computational possibility and empirical or semantic interpretation.

References

[1] K. Osmolovskyi, *Symbolic Capture of Manifestation in Cognitive Systems*, Zenodo, 2026. DOI: 10.5281/zenodo.19895497.

[2] K. Osmolovskyi, *Operationalizing Coherence-Related Manifestation in Cognitive Systems*, Zenodo, 2026. DOI: 10.5281/zenodo.20183831.

[3] K. Osmolovskyi, *A Toy Model of Perceptual Stabilization under Historical Compression and Overload*, Zenodo, 2026. DOI: 10.5281/zenodo.20184038.

[4] K. Osmolovskyi, *A Toy Model of Perceptual Stabilization and Profile-Relative Symbolic Fidelity*, Zenodo, 2026.

[5] K. Osmolovskyi, *Structural Compression in Cognitive Architectures*, Zenodo, 2026. DOI: 10.5281/zenodo.19481795.