A Toy Model of Identity-Continuation Breadth During Symbolic Reuse

Fixed-Frame Pairwise Identity Evaluation over Symbolic-Reuse Trajectories

Kostiantyn Osmolovskyi
Independent Researcher, Ukraine

Version 2.0, 2026

Abstract

This paper presents a typed toy realization of identity-bounded continuation over exact symbolic-reuse trajectories. A fixed 49-member continuation-candidate frame is evaluated at an eligible symbolic root and after each committed reuse step. The current symbolic pair is projected into a narrow toy architecture record, while the symbolic-reuse coordinate q_t conditions only the candidate-to-pair basis through a fixed gain $\gamma_t = 1 + q_t$. Identity evaluation itself uses an independent pair-norm context with a fixed threshold. Candidate-level identity-bounded continuation is partial: missing basis, absent readiness, or undefined identity-continuity evaluation remains undefined rather than negative. Three derived shares, $B_t^{(\mathrm{id},+)}$, $B_t^{(\mathrm{id},-)}$, and $B_t^{(\mathrm{id},\mathrm{undef})}$, retain the fixed denominator of 49 and do not form a recurrent identity state.

Twenty-nine pre-specified conformance tests pass. Five fixed narrative trajectories, four 201×201 deterministic symbolic-reuse surfaces, and three inherited Monte Carlo strata are evaluated without new random draws. In the standard scientific runs all basis-defined candidates are explicitly ready and identity-evaluable, so $B_t^{(\mathrm{id},\mathrm{undef})} = 0$; undefinedness is nevertheless preserved and tested in the conformance suite. The deterministic families produce mean final positive breadths of 0.3028, 0.5728, 0.5881, and 0.5604. The three Monte Carlo strata produce mean final positive breadths of 0.5171, 0.5176, and 0.5185 among exact upstream-eligible trajectories. These outcomes are consequences of the declared fixed candidate geometry combined with the supplied symbolic-reuse trajectories. They are not empirical measurements of selfhood, identity stability, identity-attractor strength, population prevalence, or pathwise identity preservation.

Keywords: identity-bounded continuation; symbolic reuse; continuation candidates; identity continuity; fixed-frame breadth; partial classification; definedness; toy model; cognitive architecture.

1 Analytic Target and Scope

Symbolic reuse and identity-continuity evaluation are distinct architectural operations. The symbolic-reuse layer supplies a temporally ordered sequence of committed symbolic forms and a load-like coordinate q_t, together with trace, reuse, grounding, commit, and source-reference descriptors [2]. Identity-bounded continuation instead concerns whether an explicitly represented continuation candidate can be mapped to an evaluable architectural pair and then classified under an independent identity-continuity context [1]. The present model studies a narrow interface between these layers.

The dependency is

$$
\begin{aligned}
&\text{exact committed symbolic-reuse state}\\
&\quad\longrightarrow \text{toy current-architecture projection}\\
&\quad\longrightarrow \text{fixed continuation candidate}\\
&\quad\longrightarrow \text{IdBasis when provenance and basis conditions are defined}\\
&\quad\longrightarrow \text{identity-evaluation readiness}\\
&\quad\longrightarrow \text{IdCont when its partial map is defined}\\
&\quad\longrightarrow \text{IdBound} \in \{0,1\}\\
&\quad\longrightarrow (B_t^{(\mathrm{id},+)}, B_t^{(\mathrm{id},-)}, B_t^{(\mathrm{id},\mathrm{undef})}).
\end{aligned}
$$

The model is a formal toy realization, not a measurement-validation study. It does not define a selfhood score, a clinical identity measure, a regulatory identity attractor, a configuration-level identity-stabilization region, or a pathwise test that an entire trajectory preserves identity.

2 Exact Symbolic-Reuse State and Temporal Alignment

The model consumes the fixed symbolic-reuse implementation without recreating its role evaluation, modulation bridges, grounding, proposal generation, symbolic-form qualification, trace, reuse, commit, or source-reference assessment. In particular, the identity layer does not use role status, role weight, effective grounding, source-reference distance, distortion, or false-completion labels as direct identity inputs.

Two temporal anchors are distinguished. At $t = 0$, an identity-evaluation frame exists only for an exact eligible symbolic root and retains its case, trajectory, root capture application, positive-form application, candidate-form identity, and initial q_0. At $t = s \geq 1$, evaluation exists only after an exact committed symbolic-reuse step s and retains the corresponding committed form, candidate, q_s, step number, trajectory, and root provenance. A missing, uncommitted, or provenance-detached step is unresolved for identity evaluation; it is not assigned IdBound $= 0$.

This explicit alignment prevents a clock index from being treated as sufficient evidence that the identity layer has a current symbolic state. It also prevents an identity result from being attached to a form produced by another case, trajectory, root, or step.

3 Fixed Continuation-Candidate Frame

Every eligible evaluation index uses the same Cartesian carrier

$$\Xi = \{(\xi_a, \xi_b) : \xi_a, \xi_b \in \{-1, -2/3, -1/3, 0, 1/3, 2/3, 1\}\},$$

with

$$N_\Xi = 49.$$

Candidate identities, order, offsets, and equal counting weights $1/49$ remain fixed across all trajectories and time indices. A candidate is not removed because its later identity classification is negative or undefined.

The current committed symbolic form supplies a decoded pair

$$z_t = (a_t, b_t),$$

which is wrapped in a provenance-bearing toy architecture record. This projection is intentionally narrow. The symbolic form is not asserted to be a complete cognitive architecture, and a candidate-associated comparison record is not a later symbolic form or an actual successor architecture.

4 Candidate-to-Pair Basis and Independent Identity Context

4.1 The local q_t basis conditioner

The only active symbolic-reuse coordinate entering the identity basis is q_t. The fixed local gain is

$$\gamma_t = 1 + q_t, \qquad r_0 = 0.20.$$

For candidate g_j with offset $\xi_j = (\xi_{a,j}, \xi_{b,j})$, the candidate-associated comparison pair is

$$z'_{t,j} = z_t + r_0 \gamma_t \xi_j.$$

No clipping is applied. The resulting object is a typed comparison architecture, not a predicted or enacted successor.

The basis is partial. It is defined only when the exact symbolic-reuse anchor is aligned, the current form and $q_t \in [0, 1]$ are defined, the candidate belongs to the fixed carrier, and the declared basis context is present. Thus q_t conditions how a fixed candidate offset is represented in the comparison pair; q_t is not itself an identity-continuity value or an IdBound result.

4.2 Readiness and pairwise identity continuity

The identity context is independent of the symbolic-reuse controls. Standard scientific runs explicitly set identity-evaluation readiness to one and define the local IdCont partial map for each basis-defined candidate. The fixed comparison metric is

$$d_{\mathrm{id}}(t, j) = \|z'_{t,j} - z_t\|_2,$$

with

$$\theta_{\mathrm{id}} = 0.20.$$

The local classifier is

$$\mathrm{IdCont}_{t,j} = 1 \quad \Longleftrightarrow \quad d_{\mathrm{id}}(t, j) \le \theta_{\mathrm{id}} + 10^{-12},$$

and otherwise the defined value is zero.

Candidate-level identity-bounded continuation is then the pullback

$$\mathrm{IdBound}_{t,j} = \begin{cases} \mathrm{IdCont}_{t,j}, & \mathrm{IdBasis}_{t,j} \text{ defined, readiness} = 1, \ \mathrm{IdCont}_{t,j} \text{ defined}, \\ \text{undefined}, & \text{otherwise.} \end{cases}$$

A defined negative value and an undefined value are therefore distinct objects.

5 Derived Fixed-Frame Breadth

For each evaluation index, let $N_+(t)$, $N_-(t)$, and $N_{\mathrm{undef}}(t)$ be the counts of positive, negative, and undefined candidate-level IdBound records. The reported breadth profile is

$$B_t^{(\mathrm{id},+)} = \frac{N_+(t)}{49}, \qquad B_t^{(\mathrm{id},-)} = \frac{N_-(t)}{49}, \qquad B_t^{(\mathrm{id,undef})} = \frac{N_{\mathrm{undef}}(t)}{49},$$

so that

$$B_t^{(\mathrm{id},+)} + B_t^{(\mathrm{id},-)} + B_t^{(\mathrm{id,undef})} = 1.$$

These quantities are derived summaries. They are not recurrent state variables and do not feed back into symbolic reuse or identity classification. When at least one candidate is evaluable, an optional evaluable-only positive rate can be computed as $N_+/(N_+ + N_-)$, but it is kept separate from $B_t^{(\mathrm{id},+)}$ so that undefined candidates cannot disappear from the denominator.

Cross-step change is reported only as

$$\Delta B_t^{+} = B_t^{(\mathrm{id},+)} - B_{t-1}^{(\mathrm{id},+)}.$$

No categorical width thresholds are introduced.

6 A Design-Theorem Relation Between q and Breadth

Under the standard all-defined premises,

$$d_{\mathrm{id}}(t, j) = 0.20(1 + q_t)\|\xi_j\|_2.$$

Hence candidate j is positive exactly when

$$\|\xi_j\|_2 \leq \frac{1}{1 + q_t}.$$

For fixed Ξ,

$$q_1 \leq q_2 \quad \Longrightarrow \quad \{j : \mathrm{IdBound}_j(q_2) = 1\} \subseteq \{j : \mathrm{IdBound}_j(q_1) = 1\},$$

and therefore

$$B^{(\mathrm{id},+)}(q_2) \leq B^{(\mathrm{id},+)}(q_1).$$

This monotonic relation is a property of the present local basis gain and pair-norm classifier. It is a design theorem, not a simulation discovery and not a general claim that symbolic reuse universally narrows identity continuity.

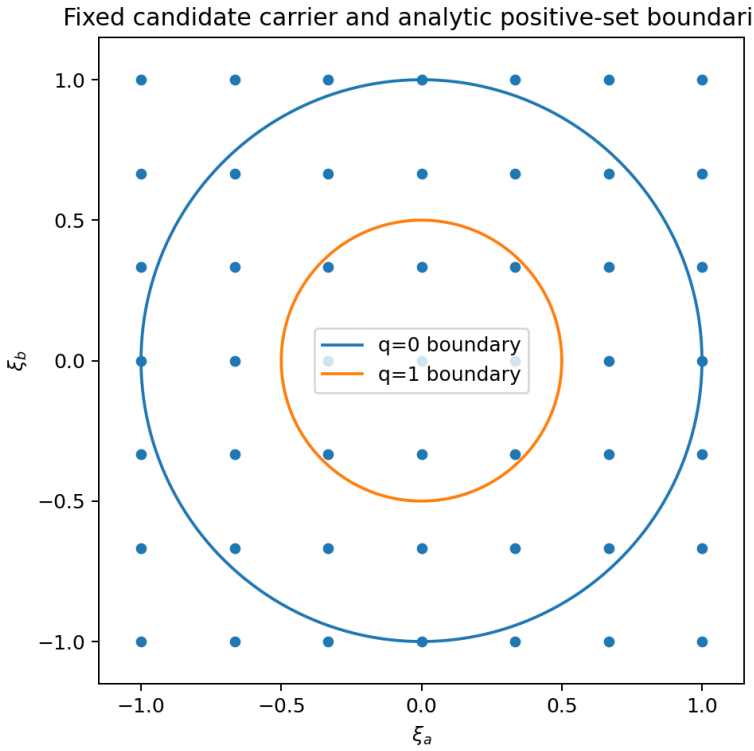


Figure 1: Fixed 49-candidate carrier in offset space. The circles show the analytic positive-set boundaries implied by the local design at $q = 0$ and $q = 1$. This panel visualizes the model geometry rather than a simulated empirical distribution.

7 Definedness and Structural Consistency Checks

Twenty-nine structural consistency tests are specified independently of the deterministic and Monte Carlo surfaces, and all twenty-nine pass. The suite includes upstream ineligibility, uncommitted-step handling, provenance detachment, missing form or q, missing basis context, readiness zero,

undefined IdCont, defined-negative versus undefined status, the closed threshold boundary, a negative outside-threshold case, retention of an undefined candidate in the fixed denominator, fixed carrier identity across steps, breadth calculation, and separation of evaluable-only rates from fixed-frame breadth.

The tests also enforce architectural boundaries. Role payloads, effective grounding, source-reference distance, reference distortion, and false-completion labels are not direct arguments of the identity basis or classifier. The identity context and threshold remain fixed when q changes. Candidate comparison architectures are typed separately from symbolic later forms and actual successors. The identity computation does not modify upstream symbolic-reuse records. The candidate-level IdBound object is kept separate from configuration identity-stabilization regions, regulatory identity attractors, and pathwise identity preservation.

In standard scientific families every aligned candidate is explicitly ready and IdCont-defined, so $B_t^{(\mathrm{id,undef})} = 0$. This zero is a declared standard-run premise, not evidence that identity evaluation is always defined. The consistency suite separately demonstrates that undefined cases remain undefined and continue to occupy the fixed denominator.

8 Simulation Design

8.1 Narrative trajectories

Five twenty-step trajectories use the same eligible symbolic root and fixed base inputs $U^{\mathrm{base}} = 0.50$ and $G^{\mathrm{base}} = 0.30$. T0 has three defined-negative modeled roles. T1 has burden-like organization positive with unit local weight; T2 has threat-like organization positive with unit weight; T3 has relief-like organization positive with unit weight; and T4 has simultaneous positive burden-like, threat-like, and relief-like roles with weights 0.8, 0.6, and 0.4. These fixtures illustrate how the already-defined symbolic-reuse layer supplies different q_t trajectories to the same identity-breadth mapping.

8.2 Deterministic surfaces and Monte Carlo strata

Four endpoint-inclusive 201×201 deterministic surfaces are inherited exactly from the symbolic-reuse design. S1 varies burden and threat weights; S2 varies relief and threat weights; S3 varies burden and relief weights; and S4 varies burden-role qualification and its weight candidate. All 40,401 trajectories in each deterministic family have eligible roots and complete twenty-step symbolic-reuse paths.

Three Monte Carlo strata contain 20,000 generated upstream episodes each, with the predefined random draws, ordering, and seeds. The identity layer adds no random draw. The numbers of upstream-eligible trajectories are 9,584, 14,416, and 16,948 for MC-M1, MC-M2, and MC-M3. The strata are reported separately and are not pooled into an unqualified population estimate.

9 Results

9.1 Narrative trajectories

The five fixtures begin from the same root breadth $B_0^{(\mathrm{id},+)} = 21/49 \approx 0.4286$. T0 and T3 finish with $q_{20} = 0$ and $B_{20}^{(\mathrm{id},+)} = 29/49 \approx 0.5918$. T1, T2, and T4 finish with $q_{20} = 1$ and $B_{20}^{(\mathrm{id},+)} = 9/49 \approx 0.1837$. These changes are consistent with the analytic q-to-breadth relation; the simulation contribution is the upstream trajectory of q, not the monotonic mapping itself.

Fixture	Final q	B_0^+	B_{20}^+	ΔB^+
T0	0.000000	0.4286	0.5918	0.1633
T1	1.000000	0.4286	0.1837	-0.2449
T2	1.000000	0.4286	0.1837	-0.2449
T3	0.000000	0.4286	0.5918	0.1633
T4	1.000000	0.4286	0.1837	-0.2449

Table 1: Final symbolic-reuse coordinate and fixed-frame positive breadth for the five declared narrative trajectories.

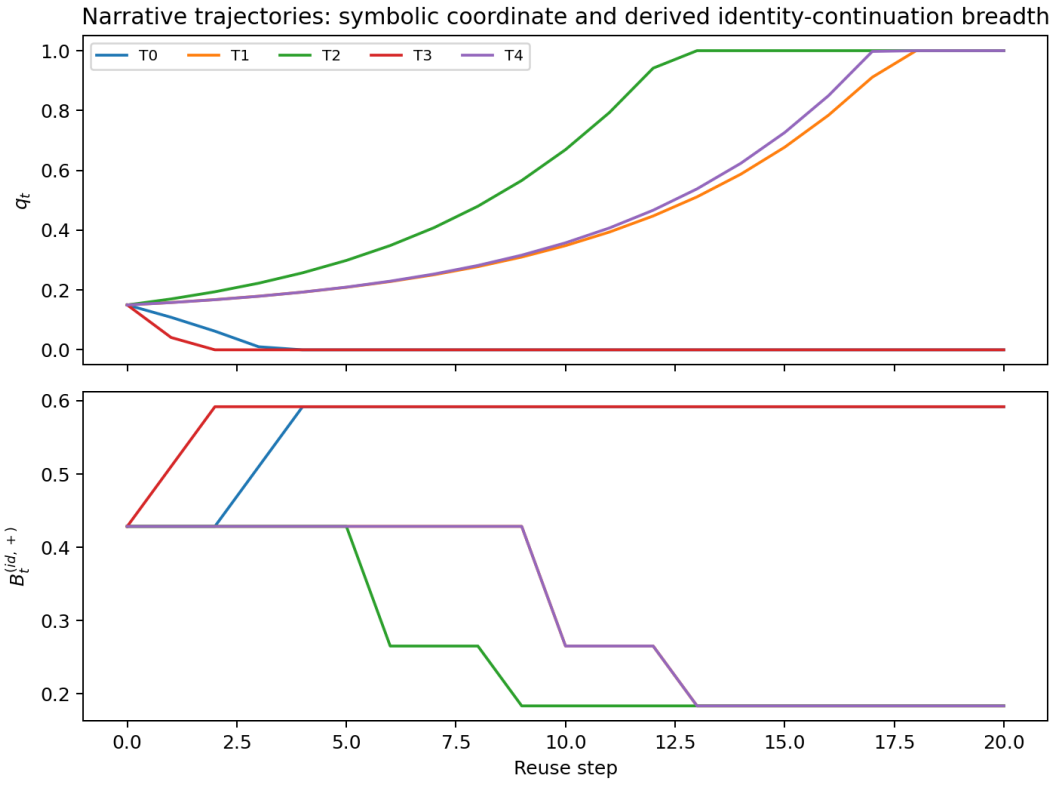


Figure 2: Upstream symbolic coordinate q_t (top) and derived positive identity-continuation breadth (bottom) across the five fixed narrative trajectories. Breadth changes occur only through changes in the supplied q_t trajectory under the fixed candidate geometry.

9.2 Deterministic surfaces

Across S1, mean final positive breadth is 0.3028, and 69.92% of trajectories finish below their common baseline breadth. S2 has mean final breadth 0.5728 and 4.34% finish below baseline. S3 has mean final breadth 0.5881 and 0.75% finish below baseline. S4 has mean final breadth 0.5604 and 7.00% finish below baseline. These surface differences arise because the fixed upstream role-modulation and symbolic-reuse equations generate different q_{20} values before the identity layer applies the same fixed candidate geometry.

Family	D_{all}	D_{elig}	Mean B_0^+	Mean B_{20}^+	$P(B_{20}^+ < B_0^+)$
S1	40,401	40,401	0.4286	0.3028	69.92%
S2	40,401	40,401	0.4286	0.5728	4.34%
S3	40,401	40,401	0.4286	0.5881	0.75%
S4	40,401	40,401	0.4286	0.5604	7.00%

Table 2: Trajectory denominators and final positive breadth summaries for the four deterministic families.

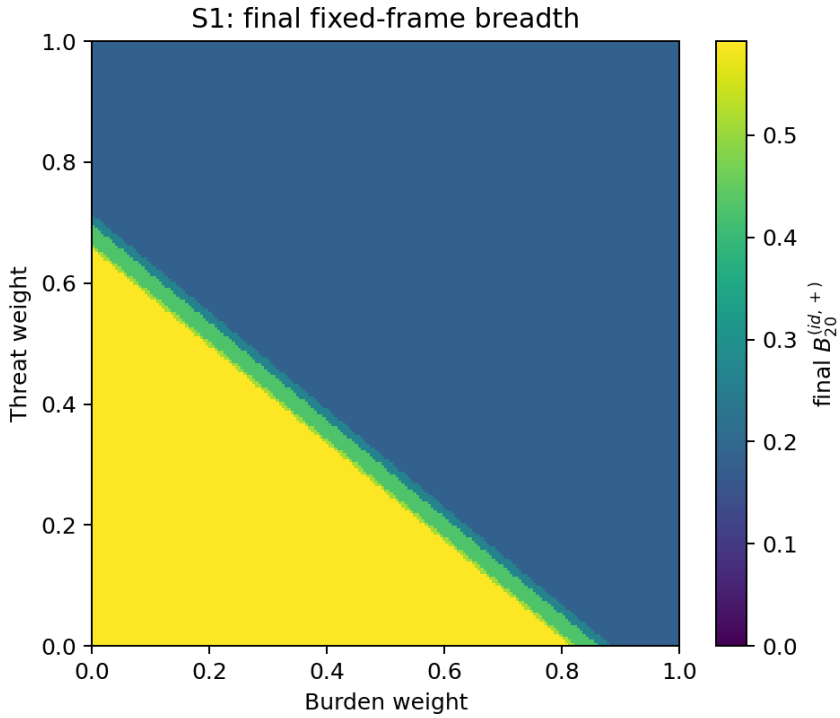


Figure 3: S1 final positive breadth over burden-weight and threat-weight inputs.

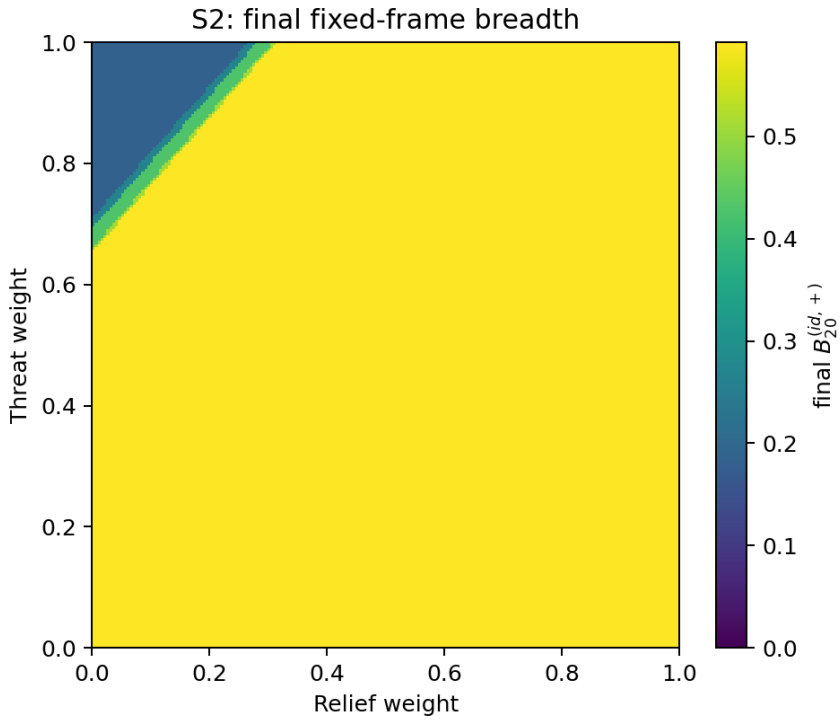


Figure 4: S2 final positive breadth over relief-weight and threat-weight inputs.

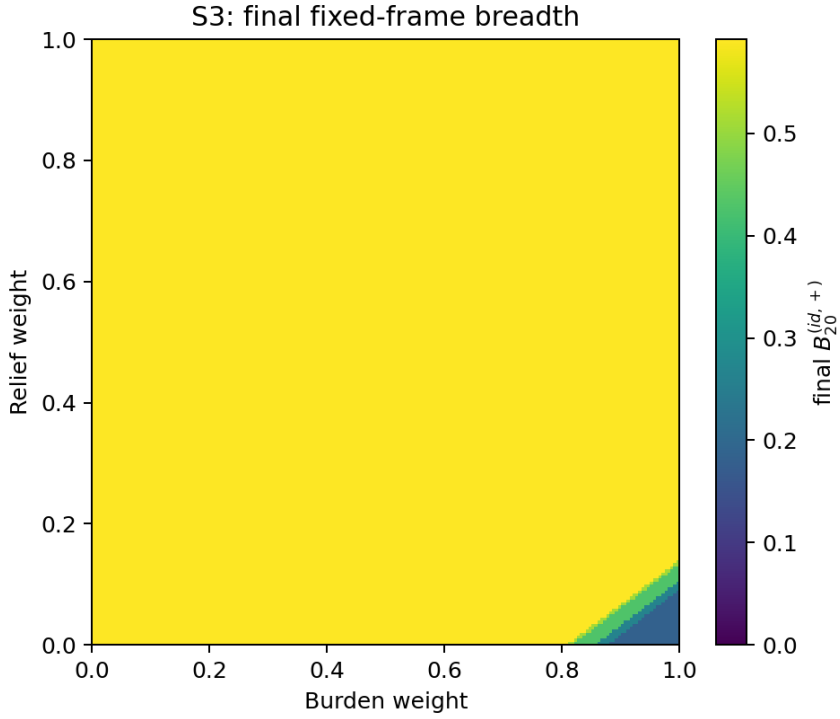


Figure 5: S3 final positive breadth over burden-weight and relief-weight inputs.

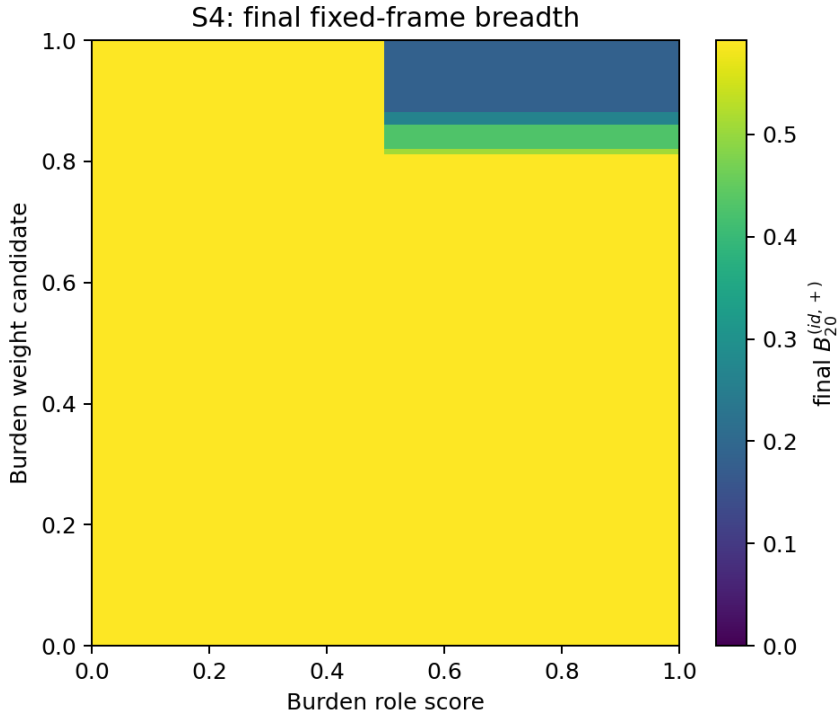


Figure 6: S4 final positive breadth over burden-role qualification and burden-weight candidate. The role-qualification boundary belongs to the upstream symbolic-reuse input design; the identity classifier itself does not consume the role variable.

9.3 Monte Carlo strata

Among upstream-eligible trajectories, mean final positive breadth is 0.5171 in MC-M1, 0.5176 in MC-M2, and 0.5185 in MC-M3. The fractions finishing below their own baseline breadth are 18.01%, 17.93%, and 17.72%, respectively. Because the candidate carrier is discrete, final breadth occupies a small set of rational values. In all three strata, the largest final mass is at $29/49 \approx 0.5918$ and a smaller mass occurs at $9/49 \approx 0.1837$.

These proportions describe outcomes under the declared upstream sampling laws and the present toy bridge. They are not prevalence estimates for identity narrowing or widening in persons, organisms, organizations, or artificial systems.

Family	D_{all}	D_{elig}	Mean B_0^+	Mean B_{20}^+	$P(B_{20}^+ < B_0^+)$
MC-M1	20,000	9,584	0.4353	0.5171	18.01%
MC-M2	20,000	14,416	0.4361	0.5176	17.93%
MC-M3	20,000	16,948	0.4368	0.5185	17.72%

Table 3: Stratum-specific denominators and final positive-breadth summaries for the three inherited Monte Carlo strata. The strata are not pooled.

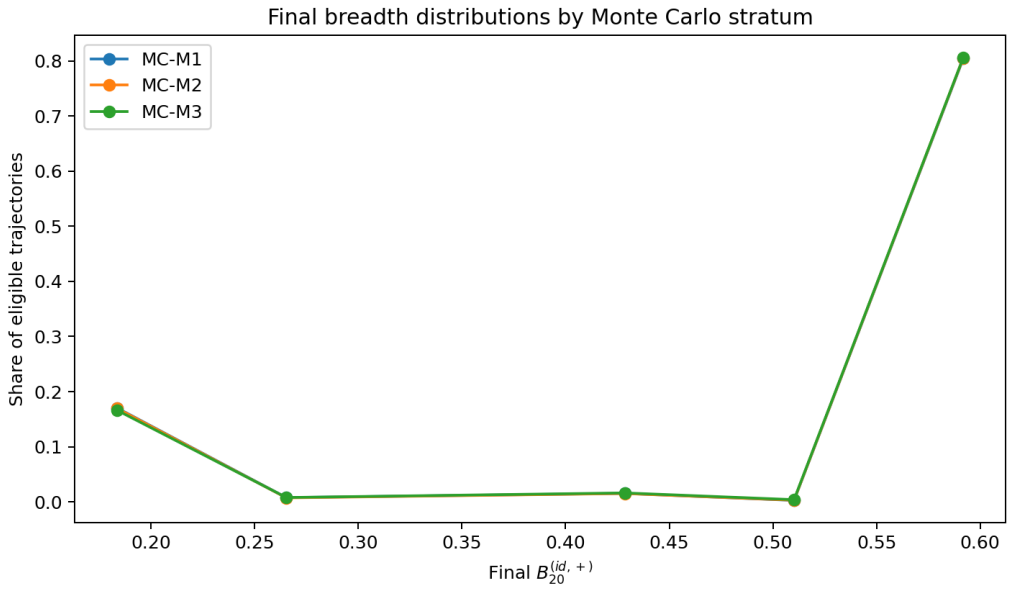


Figure 7: Final positive-breadth distributions for the three predefined Monte Carlo strata, conditioned on upstream symbolic-reuse eligibility. The strata are shown separately; no pooled prevalence interpretation is intended.

9.4 Denominator accounting

The deterministic families each contribute 848,421 aligned identity-evaluation indices: one baseline plus twenty committed steps for 40,401 trajectories. The corresponding candidate-level denominator is 41,572,629 per family. MC-M1, MC-M2, and MC-M3 contribute 201,264, 302,736, and 355,908 aligned evaluation indices, respectively. Every aligned standard-run evaluation uses all 49 candidates. Across these scientific runs, candidate-level undefined count is zero because readiness and IdCont definedness are explicit standard premises. The fixed-denominator storage and conformance tests ensure that an undefined candidate would be retained rather than reclassified or dropped.

10 Interpretation

The model demonstrates a narrow compositional possibility. An exact symbolic-reuse trajectory can supply a current symbolic pair and a load-like coordinate to a separate candidate-based identity-continuation evaluation without allowing role, grounding, reference distance, or distortion statuses to become identity criteria. A fixed candidate frame then makes breadth comparisons across steps well-defined because a change in the denominator cannot masquerade as a change in breadth.

The model also shows why positive breadth must not be interpreted as a standalone identity state. $B_t^{(\mathrm{id},+)}$ is a summary of 49 pairwise candidate classifications under one fixed local basis and one fixed identity context. It does not identify an attractor, define a stable configuration region, certify the identity of the whole trajectory, or establish whether any candidate will be admissible, enacted, or realized. The same applies to ΔB_t^{+}: it describes a change in this local candidate share, not a transition in a general identity phase space.

The deterministic and Monte Carlo results should likewise be read as propagation results. Once the upstream trajectory supplies q_t, the positive set is analytically constrained by the fixed geometry. The surfaces therefore expose how variation already present in symbolic reuse is expressed through the declared identity-basis bridge; they do not independently discover a causal effect of affect-like roles or symbolic distortion on identity continuity.

11 Computational Reproducibility

The identity-continuation layer is deterministic once an aligned symbolic-reuse state, its current ordered pair, and the symbolic coordinate q_t are supplied. The 49-member candidate frame, candidate order, local gain, pair projection, distance metric, identity threshold, readiness convention, and fixed-denominator breadth definitions are held constant across all reported evaluations. The implementation therefore separates two reproducibility surfaces: generation of the upstream symbolic-reuse trajectories, described in the cited symbolic-reuse model, and deterministic evaluation of those trajectories by the present identity-continuation layer.

The accompanying implementation includes the candidate geometry, partial candidate evaluation, fixed-denominator breadth calculation, declared simulation design, reported summaries, and structural consistency checks. The checks cover alignment and definedness, the closed threshold boundary, defined-negative versus undefined status, fixed candidate order, denominator preservation, independence from excluded symbolic-reuse fields, and the distinction between candidate comparison architectures and actual successors. The reported figures and tables are generated from the declared model outputs without introducing additional random draws in the identity-continuation layer.

12 Limitations

The candidate carrier, basis gain, pair projection, distance metric, and identity threshold are declared toy choices. Their usefulness here is architectural: they make typing, definedness, provenance, and denominator semantics explicit. They are not empirically calibrated and should not be treated as psychologically or biologically validated scales.

Only q_t is selected as an active symbolic-reuse basis conditioner. This selection is local to the model and does not establish that q is the unique or correct variable for identity-continuity analysis. Direct use of source-reference distance, distortion, grounding strength, or affect-like role variables is deliberately absent; alternative bridges would require separate assumptions and separate evaluation rather than being inferred from the present results.

The standard scientific runs do not inject identity-readiness or IdCont missingness. Consequently, the undefined breadth share is zero in the reported surfaces. The implementation supports and tests undefinedness, but empirical or richer theoretical work would be needed to specify when candidate identity evaluation should be unavailable.

Finally, pairwise candidate continuity is not pathwise identity. A sequence can have high positive candidate breadth at multiple time points without that fact alone constituting a proof of identity preservation across the sequence. Conversely, a defined negative candidate does not imply impossibility, pathology, failure of symbolic reuse, or inability to enact another continuation.

13 Conclusion

This paper establishes a typed toy interface from exact symbolic-reuse states to candidate-level identity-bounded continuation. The interface keeps temporal provenance explicit, preserves partial definedness, uses one fixed continuation-candidate denominator across trajectories and steps, and derives breadth only after candidate-level identity evaluation. Under the declared pair-norm geometry, positive breadth is analytically non-increasing in the supplied q coordinate, while the symbolic-reuse layer remains solely responsible for the actual trajectory of q.

The resulting deterministic surfaces and Monte Carlo strata show how different exact symbolic-reuse trajectories are expressed through this fixed local identity-basis bridge. They do not convert breadth into an identity attractor, a configuration-level stability region, a pathwise identity criterion, or an empirical prevalence measure. The principal result is therefore architectural: identity-continuation breadth can be represented as a fixed-frame, provenance-complete, definedness-aware summary without collapsing the distinctions that separate symbolic reuse, candidate-level identity continuity, and broader identity organization.

References

[1] K. Osmolovskyi. *Identity-Bounded Continuation in Cognitive Systems*. 2026. doi:10.5281/zenodo.19789471.

[2] K. Osmolovskyi. *A Toy Model of Affect-like Role Modulation in Symbolic Reuse: Role-Qualified Modulation of Reuse Intensity and Source Re-alignment*. 2026. doi:10.5281/zenodo.20186054.