A Toy Model of Perceptual Stabilization and Profile-Relative Symbolic Fidelity

From Perceptual Application Coordinates to Preservation, Reduction, and Distortion

Kostiantyn Osmolovskyi
Independent Researcher, Ukraine

Version 2.0, 2026

Abstract

This paper presents a minimal toy model connecting a typed perceptual-stabilization application to profile-relative symbolic fidelity without identifying the perceptual application with the source of symbolic capture. The perceptual component separately represents route formation, stabilization level, perceptual-stabilization status, current-signal support, support discordance, and a restricted false-completion flag. The symbolic component independently establishes a synthetic manifestation application, aligns it with the same synthetic episode, qualifies a candidate symbolic form under a fixed codebook, establishes source-specific symbolic capture through encoder provenance, and only then evaluates a declared fidelity profile. Preservation, reduction, and distortion are independent profile-relative Boolean assessments; a scalar $D^{\mathrm{sym,toy},p}$ is computed only after positive capture and is not a definition of symbolic distortion.

Thirteen structural consistency tests pass. Deterministic surfaces show that source geometry and perceptual conditioning affect different fidelity boundaries. In the selector/dominance isolation surface, all 40,401 cases are support-discordant and 30,301 receive the false-completion flag, yet SymDist = 1 never occurs on that surface. Three Monte Carlo strata of 20,000 cases each yield positive profile-relative distortion in 24.03%, 24.14%, and 25.02% of their respective fidelity-defined cases. Distortion occurs in only 45.08%, 44.57%, and 44.80% of support-discordant cases, and in 72.12%, 72.14%, and 70.61% of false-completion-positive cases. These are properties of the declared toy equations, source distributions, codebook, and fidelity profile; they are not empirical prevalence estimates, semantic truth judgments, or universal cognitive laws.

Keywords: perceptual stabilization; symbolic capture; symbolic fidelity; preservation; reduction; distortion; manifestation; support discordance; toy model; overload.

1 Analytic Target and Scope

Perceptual stabilization and symbolic capture answer different architectural questions. A route-qualified perceptual application may be formed and evaluated for stabilization and current-signal support without thereby becoming a symbolic object. Conversely, a symbolic form does not count as a capture merely because it exists in the same episode. The present model studies a narrower interface: how selected coordinates of a formed perceptual application may condition one explicitly constructed symbolic encoding of an independently established manifestation source, and how that positive symbolic capture may then be assessed under one declared fidelity profile.

The perceptual component follows the distinction among recruitment, formation, perceptual stabilization, and current-signal support developed in the perceptual-stabilization architecture and instantiated in the companion toy model [2, 5]. In particular,

$$\mathrm{PercStab}_t \neq \mathrm{SigSup}_t,$$

and a support-discordant application with $\mathrm{PercStab}_t = 1$ and $\mathrm{SigSup}_t = 0$ is formally distinct from the narrower simulation-level false-completion interpretation.

The symbolic component follows the typed architecture in which a candidate form must first receive a positive symbolic-form qualification, a positive capture relation must then link that form to a specific manifestation application, and fidelity is assessed only on the resulting positive capture under a well-formed comparison profile [1, 3]. Thus the computational dependency is

$$\begin{aligned}&\text{formed perceptual-application coordinates} + \text{independent synthetic manifestation application}\\ &\quad\longrightarrow \text{same-episode alignment}\\ &\quad\longrightarrow \mathrm{SymForm}^{\mathrm{toy}}\\ &\quad\longrightarrow \mathrm{SymCap}^{\mathrm{toy}}\\ &\quad\longrightarrow (\mathrm{SymPres}, \mathrm{SymRed}, \mathrm{SymDist}, D^{\mathrm{sym,toy},p}).\end{aligned}$$

The construction is an illustrative computational realization rather than a measurement-validation model [4].

The model does not provide a theory of language, semantic truth, belief, consciousness, hallucination, pathology, communication, symbolic persistence, reuse, or cross-system readability. It also does not claim that overload universally increases symbolic distortion or that perceptual support discordance universally produces distorted symbolic capture.

2 Joint Synthetic Episode

2.1 Imported perceptual application

The perceptual side uses the computational realization described in the companion toy model [5]. No perceptual equation is redefined by the symbolic layer. For each episode, the perceptual result supplies the route status, recruited field, selected and historically dominant indices, stabilization score $J_t^{\mathrm{stab,toy}}$, stabilization level, $\mathrm{PercStab}_t$, SigSup_t, support-discordance status, and the restricted $\mathrm{FalseComp}_t^{\mathrm{toy}}$ flag.

The normalized coordinates ℓ_t^{toy} and r_t^{toy} remain local realization coordinates. Their use below as symbolic-conditioning inputs is an additional toy assumption, not a consequence of the symbolic-capture construction [3].

2.2 Independent manifestation source

The symbolic branch does not use the perceptual application as the source of symbolic capture. Each standard case independently establishes a synthetic manifestation application [1, 3]

$$e_t^{\mathrm{man,toy}} = (\lambda_t^{\mathrm{man,toy}}, c_t^{\mathrm{man,toy}}, m_t^{\mathrm{toy}}),$$

with its own source/application identifier and context. The manifestation output token remains formally unstructured. A separate model-local descriptor registry

$$q_t^{\mathrm{src}} = (a, b, c)$$

is keyed to that manifestation application solely for the encoder and fidelity profile. The tuple is not asserted to be intrinsic structure of the manifestation output carrier.

The standard source domain is

$$0 \le a < b \le 1, \qquad 0 \le c \le 1.$$

2.3 Same-episode alignment

A local alignment status is evaluated only when both a formed perceptual application and an established manifestation application are present. Positive alignment requires the same synthetic episode identifier, the same synthetic signal token, the declared signal-to-perceptual and signal-to-manifestation token bridges, and the pre-specified alignment context. Positive alignment means only that the two distinct applications may participate in the declared conditioning rule. It does not identify them.

If the perceptual route is absent or unresolved, if the manifestation source is not established, or if the required conditioning coordinates are undefined, the standard symbolic branch is not generated. None of these cases is converted to symbolic zero.

3 Symbolic Form and Capture

3.1 Codebook and symbolic-form qualification

The candidate alphabet Q21-v1 contains tokens

$$q_{00}, q_{01}, \ldots, q_{20}$$

with decoded centroids $0.00, 0.05, \ldots, 1.00$. The reserved omission symbol $\varnothing$ is allowed only in the third payload position. Quantization uses deterministic half-up rounding:

$$Q_{0.05}(x) = \min\left(20, \max\left(0, \left\lfloor 20x + \frac{1}{2} \right\rfloor\right)\right).$$

A candidate form is the record

$$(\texttt{SYM1}, \texttt{Q21-v1}, (s_a, s_b, s_c)).$$

The local SymForm$^{\text{toy}}$ criterion is positive exactly when the fixed form context is `SYMFORM-Q21-v1` and the header, alphabet, payload length, and token positions satisfy this grammar. A candidate evaluated in a different declared form context, or a malformed but otherwise assessable record, receives zero under this criterion; absent candidate or absent form context leaves SymForm undefined.

3.2 Pre-capture conditioning drives

Two simulation-local drives condition the encoder. They are not symbolic-fidelity measures. On a positive alignment with all required perceptual coordinates defined,

$$\begin{aligned} u_t^{\text{dist}} = \text{clip}_{[0,1]}\,(&0.05 + 0.25(1 - J_t^{\text{stab,toy}}) + 0.15\mathbf{1}[\text{SigSup}_t = 0] \\ &+ 0.10\mathbf{1}[\text{support discordant}] + 0.15\mathbf{1}[\text{FalseComp}_t^{\text{toy}} = 1] \\ &+ 0.15\ell_t^{\text{toy}} - 0.10r_t^{\text{toy}}), \end{aligned}$$

and

$$\begin{aligned} u_t^{\text{red}} = \text{clip}_{[0,1]}\,(&0.05 + 0.35(1 - J_t^{\text{stab,toy}}) + 0.10\mathbf{1}[\text{SigSup}_t = 0] \\ &+ 0.10\ell_t^{\text{toy}} - 0.10r_t^{\text{toy}}). \end{aligned}$$

The third source descriptor is omitted when

$$u_t^{\text{red}} \geq 0.27.$$

These equations deliberately make selected perceptual coordinates relevant to the local encoder, but they do not make any perceptual status sufficient for any symbolic-fidelity status.

3.3 Encoder and source-specific capture

For $q_t^{\mathrm{src}} = (a, b, c)$, the pre-quantization protected descriptors are

$$a^* = \mathrm{clip}_{[0,1]}(a + 0.50u_t^{\mathrm{dist}}), \qquad b^* = \mathrm{clip}_{[0,1]}(b - 0.50u_t^{\mathrm{dist}}).$$

The third descriptor is either quantized from c or encoded as $\varnothing$ according to the reduction gate. The encoder record stores the exact manifestation application identifier, candidate-form identifier, conditioning-record identifier, encoder identifier, and both drives.

Positive symbolic capture is then evaluated separately. A positive $\mathrm{SymCap}^{\mathrm{toy}}$ requires a positive symbolic-form application in the fixed form context and complete encoder provenance linking that exact form candidate to the exact manifestation application under the fixed capture context. Complete source mismatch gives $\mathrm{SymCap}^{\mathrm{toy}} = 0$; missing capture context or provenance leaves the relation undefined. Hence

$$\mathrm{SymForm}^{\mathrm{toy}} = 1 \not\Rightarrow \mathrm{SymCap}^{\mathrm{toy}} = 1.$$

Positive capture also does not imply persistence, later reuse, expression, communication, or semantic truth [3].

4 Profile-Relative Fidelity

The model uses one fixed fidelity profile, `FID-REL3-v1`. It compares the source descriptor record (a, b, c) with decoded symbolic descriptors (a', b', c'), where c' may be absent. The indexwise alignment is profile-local:

$$a \leftrightarrow a', \qquad b \leftrightarrow b', \qquad c \leftrightarrow c'.$$

No universal manifestation-to-symbol metric is asserted. The fidelity application retains the exact positive-capture identifier, manifestation-source identifier, positive symbolic-form application, candidate-form identifier, source-descriptor record, and profile identifier; fidelity is therefore undefined if a detached candidate or source record is substituted for the captured pair.

4.1 Preservation

The protected relation is the source ordering and gap of (a, b). With tolerance $\varepsilon_{\mathrm{gap}} = 0.20$,

$$\mathrm{SymPres} = 1$$

iff

$$b' > a' \quad \text{and} \quad \big|(b' - a') - (b - a)\big| \le 0.20.$$

Otherwise, on the same defined fidelity application, $\mathrm{SymPres} = 0$.

4.2 Reduction

Reduction tracks omission of the optional third descriptor:

$$\mathrm{SymRed} = 1 \quad \Longleftrightarrow \quad c' \text{ is absent.}$$

This is profile-relative symbolic reduction and does not imply structural compression [7].

4.3 Distortion

Distortion tracks loss or reversal of the protected source ordering:

$$\text{SymDist} = 1 \quad \Longleftrightarrow \quad b' \leq a'.$$

It is not, by definition, semantic falsity, false belief, deception, hallucination, pathology, or externally incorrect report.

The three statuses are not encoded as an exclusive symbolic mode. In this particular profile SymPres and SymDist cannot both be positive because their local criteria are incompatible, while SymRed may coexist with either. The profile therefore permits, among others,

$$(1,0,0), \quad (1,1,0), \quad (0,0,1), \quad (0,1,1), \quad (0,0,0).$$

This profile-specific combination structure is not a universal complement law for symbolic fidelity.

4.4 Post-capture scalar proxy

Only after positive capture and formation of the profile-qualified fidelity application is the local scalar

$$D_t^{\text{sym,toy},p} = \frac{1}{2}\left(|a' - a| + |b' - b|\right)$$

computed. It ranges over $[0, 1]$ for this profile. Zero means exact equality of the two protected descriptor values after decoding. It does not imply absence of reduction or global fidelity. No scalar threshold defines SymPres or SymDist.

5 Definedness and Structural Consistency Tests

The implementation keeps distinct the absence or unresolvedness of the perceptual route, absence or unresolvedness of the manifestation source, undefined or negative SymForm, undefined or negative SymCap, profile inapplicability, defined negative fidelity statuses, numeric zero, and positive scalar divergence. No missing status is replaced with zero.

Thirteen structural tests cover these boundaries and the non-identification claims. All pass under the declared implementation. Selected tests are especially diagnostic:

- CT-1 keeps all symbolic outputs undefined when no formed perceptual route exists, rather than assigning symbolic zero.
- CT-5 yields SymForm $= 1$ but SymCap $= 0$ under complete source-provenance mismatch.
- CT-9 yields $(\text{SymPres}, \text{SymRed}, \text{SymDist}) = (1, 1, 0)$ with $D^{\text{sym,toy},p} = 0$, showing that zero protected-pair divergence does not mean absence of reduction.
- CT-11 yields $(0, 0, 0)$ on a fully defined fidelity application, showing that negative preservation does not force distortion.
- CT-12 uses a perceptual case with $\text{FalseComp}^{\text{toy}} = 1$ but produces SymDist $= 0$ under the fixed source/profile configuration.
- CT-13 keeps the symbolic conditioning drives undefined when required perceptual support-derived coordinates are undefined.

Table 1: Pre-specified structural consistency suite.

Test	Result	Test	Result
CT-1	PASS	CT-8	PASS
CT-2	PASS	CT-9	PASS
CT-3	PASS	CT-10	PASS
CT-4	PASS	CT-11	PASS
CT-5	PASS	CT-12	PASS
CT-6	PASS	CT-13	PASS
CT-7	PASS		

6 Computational Design

The analyses use one fixed computational specification comprising the perceptual equations, symbolic-conditioning equations, codebook, fidelity profile, deterministic grid ranges, Monte Carlo distributions, seeds, sample sizes, denominator rules, and structural-test expectations. The reported deterministic and stochastic results are evaluated under this single specification.

Four deterministic 201×201 families are evaluated:

1. S1 varies $\ell_t^{\text{toy}} \times r_t^{\text{toy}}$ under a fixed three-block episode and source $(0.20, 0.80, 0.50)$.
2. S2 varies source gap $g = b - a \in [0.20, 0.80]$ against selected current-signal support, holding the remaining one-block perceptual inputs fixed.
3. S3 varies one-block match against current-signal support, exposing route, stabilization/support, and fidelity boundaries on the same surface.
4. S4 varies two-block selector/dominance offsets $\delta_M \times \delta_P$, while both blocks remain recruited and the source/profile remain fixed.

Three Monte Carlo episode/source strata use the perceptual-model parameters reported in the companion toy model [5]. Each contains 20,000 cases, with one, two, or three blocks. For every case,

$$\ell, r \sim U(0,1), \qquad M_k, P_k, S_k^{\text{sup}} \sim U(0,1),$$

and the independent source geometry is sampled as

$$a \sim U(0.10, 0.25), \quad g \sim U(0.35, 0.65), \quad b = a + g, \quad c \sim U(0.10, 0.90).$$

The main Monte Carlo policy supplies an established manifestation source, standard positive same-episode alignment, the declared encoder, correct capture provenance, and the fixed fidelity profile whenever the perceptual route is formed. Missingness, malformed forms, capture mismatch, and profile inapplicability are evaluated by the structural consistency suite rather than assigned arbitrary random probabilities.

The use of fixed parameter ranges and denominator-conditioned reporting follows the sensitivity-analysis framework developed for the companion perceptual model [6].

All reporting is denominator-conditioned. In particular, D_{all} differs from the formed-route, positive-capture, and fidelity-defined denominators even when several downstream denominators happen to coincide under the standard symbolic policy.

7 Results

7.1 Illustrative fixed episodes

Five fixed episodes provide compact examples of the full pipeline. Under the standard source $(0.20, 0.80, 0.50)$, the clear-stabilization and resource-recovery cases are profile-preservative and non-reductive; the ambiguous-signal case is reductive but not distortive; the high-overload dominant-block case is both reductive and distortive; and the no-route case leaves the symbolic branch undefined. These examples are descriptive fixtures rather than independent tests of the chosen coefficients.

7.2 Deterministic surfaces

Table 2 summarizes the four deterministic families on their fidelity-defined domains.

Table 2: Deterministic families. Fidelity proportions are conditioned on the profile-defined domain.

Family	$D_{\text{fid}}/D_{\text{all}}$	SymPres +	SymRed +	SymDist +	mean $D^{\text{sym,toy},p}$
S1	23,316/40,401	0.00%	66.13%	61.42%	0.272
S2	40,401/40,401	8.73%	79.10%	57.71%	0.260
S3	22,713/40,401	23.38%	33.90%	59.67%	0.242
S4	40,401/40,401	0.00%	0.00%	0.00%	0.238

S2 demonstrates that source-side organization matters to profile-relative fidelity. The same symbolic transformation can preserve or reverse the protected relation depending on the initial source gap, while reduction is controlled by a separate omission criterion. Figure 1 shows the profile-specific scalar together with preservation and distortion surfaces.

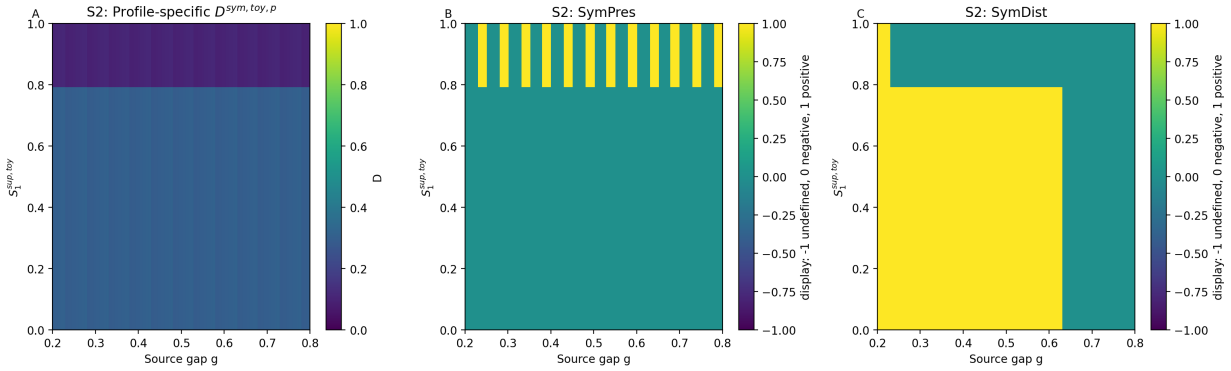


Figure 1: S2: source-gap × signal-support analysis. (A) profile-specific protected-pair divergence; (B) profile-relative preservation; (C) profile-relative distortion. The fidelity outcomes depend on both the symbolic transformation and the protected source geometry.

S3 makes the multi-coordinate organization especially visible. No fidelity status is assigned in the no-route region. Within the formed region, preservation, reduction, and distortion have different boundaries rather than occupying a single ordered symbolic axis (Figure 2).

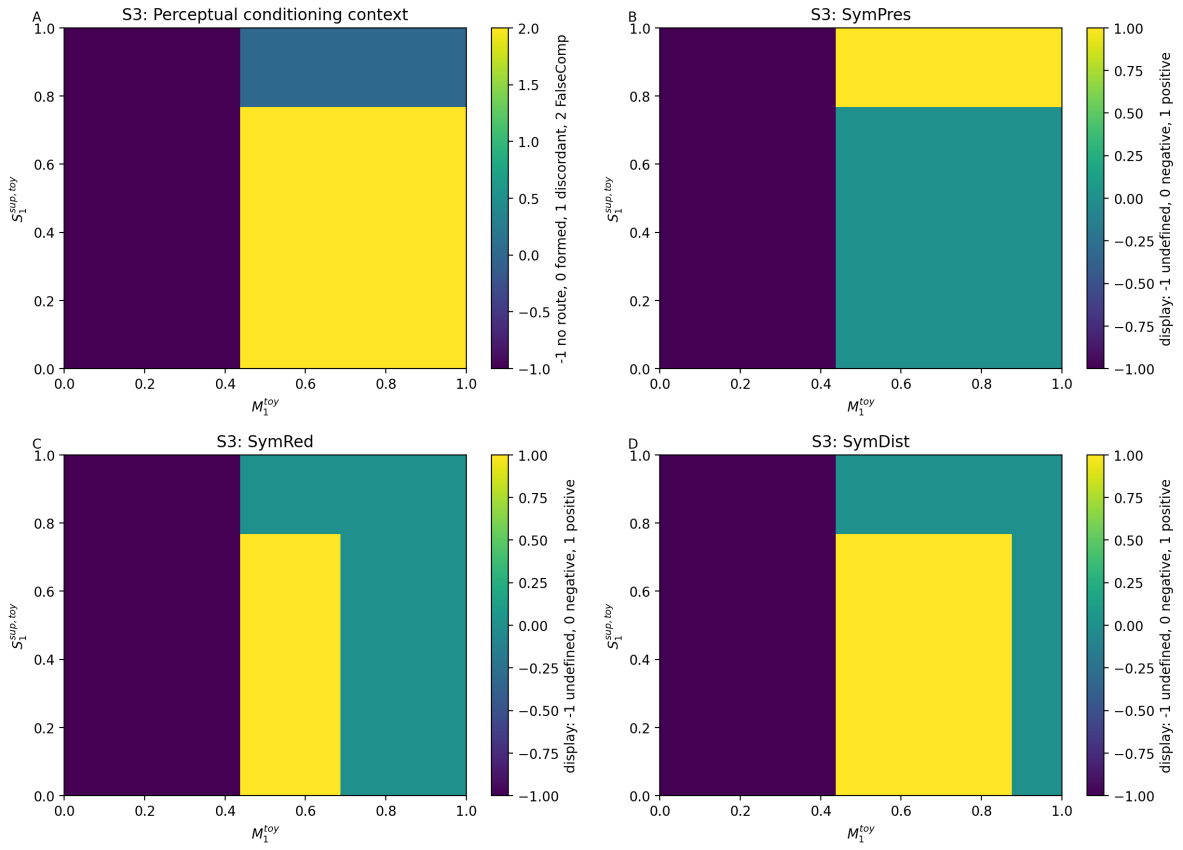


Figure 2: S3: match × current-signal-support analysis. (A) perceptual conditioning context, with undefined/no-route cases separated from formed cases and support-discordant/false-completion cases; (B) SymPres; (C) SymRed; (D) SymDist. Integer values in panel A are display codes only.

S4 provides the clearest non-entailment witness. All 40,401 cells are formed, fully stabilized, and support-discordant under the selected perceptual inputs. The formation selector coincides with the historically dominant block in 30,301 cells, producing $\text{FalseComp}^{\text{toy}} = 1$ there. The extra false-completion term raises u_t^{dist} and changes $D^{\text{sym,toy},p}$ from 0.20 to 0.25, yet the protected source ordering remains intact throughout the entire grid, so SymDist = 0 in all 40,401 cells (Figure 3). Thus, even inside this deliberately coupled toy encoder,

$$\text{FalseComp}^{\text{toy}} = 1 \not\Rightarrow \text{SymDist} = 1.$$

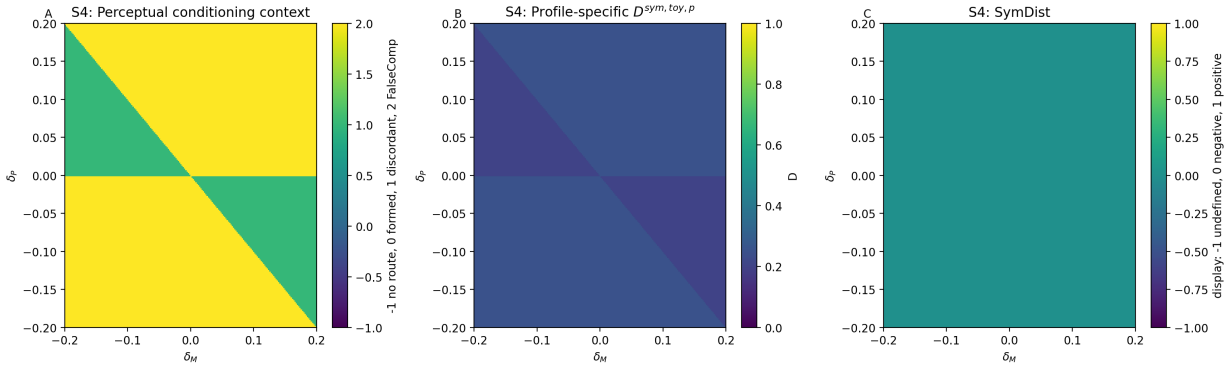


Figure 3: S4 selector/dominance isolation. (A) perceptual model context: all cases are support-discordant, while the false-completion flag occupies the selector–dominance coincidence region; (B) the profile-specific scalar differs between those regions; (C) SymDist remains negative everywhere.

S1, by contrast, yields no preservation-positive cases under its fixed narrow source/profile configuration and selected perceptual inputs; 66.13% of fidelity-defined cases are reductive and 61.42% are distortive. This zero-frequency result follows from the specified anchors and does not imply a general absence of preservation.

7.3 Monte Carlo episode/source coverage

The fidelity-defined denominators increase with block count because positive route formation becomes more frequent under the declared independent sampling of block-level match and dominance coordinates: 9,501, 14,441, and 17,007 cases for the one-, two-, and three-block strata respectively. These are properties of the sampling design, not estimates of cognitive prevalence.

Table 3: Monte Carlo fidelity summaries. SymPres, SymRed, and SymDist proportions use their common defined fidelity domain. The two conditional distortion columns use support-discordant and false-completion-positive subsets of that domain.

Stratum	$D_{\text{fid}}/D_{\text{all}}$	Pres+	Red+	Dist+	Dist+ |Disc+	Dist+ |FC+	mean D
MC-E1	9,501/20,000	19.10%	50.08%	24.03%	45.08%	72.12%	0.173
MC-E2	14,441/20,000	20.42%	45.41%	24.14%	44.57%	72.14%	0.174
MC-E3	17,007/20,000	21.17%	42.35%	25.02%	44.80%	70.61%	0.175

Across the three strata, profile-relative preservation occurs in approximately 19–21% of fidelity-defined cases, reduction in approximately 42–50%, and distortion in approximately 24–25%. The scalar distribution remains similar across block counts, with means 0.173, 0.174, and 0.175 (Figure 4).

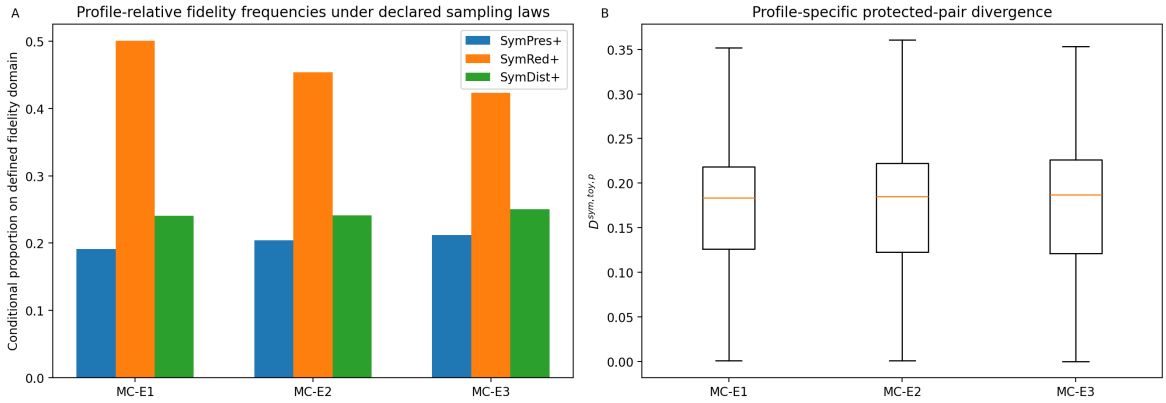


Figure 4: Monte Carlo summaries. (A) independent profile-relative fidelity-status proportions on the fidelity-defined domain; (B) distribution of the post-capture protected-pair scalar $D^{\text{sym,toy},p}$.

The conditional results also delimit the interpretation. Support discordance raises the distortion rate under the selected encoder, but distortion is present in only 45.08%, 44.57%, and 44.80% of support-discordant cases. The narrower false-completion flag is associated with a stronger local conditioning term, yet distortion still occurs in only 72.12%, 72.14%, and 70.61% of false-completion-positive cases. These values describe the specified encoder and source distributions, not a general architectural implication.

The full joint fidelity table also confirms that reduction is not an exclusive middle mode. It co-occurs with distortion in substantial subsets and with preservation in a small number of sampled cases; CT-9 supplies an explicit zero-divergence preservation-plus-reduction case even where a finite Monte Carlo stratum contains few or no such draws.

Table 4: Joint profile-relative fidelity combinations on the common defined domain. Cells report count and within-stratum proportion.

SymPres	SymRed	SymDist	MC-E1	MC-E2	MC-E3
0	0	0	2,164 (22.78%)	3,547 (24.56%)	4,408 (25.92%)
0	0	1	768 (8.08%)	1,389 (9.62%)	1,796 (10.56%)
0	1	0	3,239 (34.09%)	4,459 (30.88%)	4,744 (27.89%)
0	1	1	1,515 (15.95%)	2,097 (14.52%)	2,459 (14.46%)
1	0	0	1,811 (19.06%)	2,947 (20.41%)	3,600 (21.17%)
1	0	1	0 (0.00%)	0 (0.00%)	0 (0.00%)
1	1	0	4 (0.04%)	2 (0.01%)	0 (0.00%)
1	1	1	0 (0.00%)	0 (0.00%)	0 (0.00%)

The zeros in combinations with both SymPres = 1 and SymDist = 1 are consequences of the mutually incompatible preservation and distortion criteria in this single declared profile. They do not establish a universal complement relation between preservation and distortion.

8 Interpretation

The model establishes a typed route from perceptual application coordinates to symbolic-fidelity outcomes without turning the perceptual state into the symbolic source. The source of capture remains a separately established manifestation application; perceptual coordinates enter only through a declared same-episode conditioning record. This separation matters because a formed or stabilized perceptual application alone supplies neither a positive symbolic form nor a positive source-form capture relation.

The simulation also shows why fidelity is more informative when represented by several profile-relative coordinates rather than one exclusive symbolic category. Reduction tracks omission of an optional source descriptor, preservation tracks the protected ordering and gap, and distortion tracks loss or reversal of that ordering. These relations can diverge. Reduction can coexist with preservation or distortion, and a fully defined capture can satisfy neither preservation nor distortion. A scalar displacement of the protected pair supplies additional local information but cannot reconstruct these distinctions by thresholding alone.

The perceptual-conditioning effects are similarly conditional. In the declared encoder, support-negative and support-discordant states increase the pre-capture drive, and the false-completion flag contributes an additional increment. Consequently these perceptual coordinates alter the distribution of fidelity outcomes. But source geometry remains decisive: a wide protected gap can remain ordered even under strong conditioning. S4 and CT-12 provide direct model-level witnesses that false completion does not entail profile-relative symbolic distortion.

These results should be read as constructive model properties. They show that the declared architectural distinctions can coexist in one executable model and that nontrivial downstream relations can be studied without collapsing those distinctions. They do not establish that real cognitive systems implement this codebook, these coefficients, this source geometry, or this fidelity profile.

9 Limits

First, the manifestation application is a synthetic construction used for this computational realization. The descriptor tuple (a, b, c) is a local registry attached to that application for this encoder/profile; it is not claimed to reveal intrinsic structure of manifestation itself.

Second, Q21-v1 is an intentionally minimal symbolic codebook. Symbolic-form qualification here demonstrates a non-vacuous local criterion, not a general theory of symbolicity or language.

Third, the source-specific capture criterion is provenance-based. It shows how a positive capture

relation can be made explicit in a toy implementation, but it does not establish that provenance matching is a universal criterion for symbolic capture.

Fourth, the fidelity profile is selective. It protects one ordered pair and treats one additional descriptor as optional. Another well-formed profile could protect different relations and produce different preservation, reduction, or distortion statuses for the same positive capture. No cross-profile ordering or global fidelity score is defined.

Fifth, the conditioning coefficients are simulation assumptions. The observed association between support discordance, false completion, and distortion is generated partly because the local encoder gives those coordinates explicit weights. The simulation therefore supports conditional statements about this implementation, not a universal psychological or architectural law.

Sixth, the Monte Carlo frequencies are conditioned on declared uniform distributions over toy episode inputs and source descriptors. They are not prevalence estimates, calibrated probabilities, or predictions for biological, social, or artificial cognitive systems.

Finally, the model ends at current profile-relative fidelity. It does not infer current symbolic availability, persistence, trace continuity, later reactivation, reuse, external expression, communication, cross-system readability, or structural updating from positive capture.

10 Conclusion

This paper shows that perceptual stabilization and symbolic fidelity can be coupled in a minimal computational model while preserving their architectural boundaries. A formed perceptual application and an independently established manifestation application are aligned within one synthetic episode; the perceptual coordinates condition a local encoder, but the manifestation application remains the source of symbolic capture. Symbolic-form qualification, source-specific capture, and profile qualification are explicit stages rather than consequences of perceptual stabilization.

Within the fixed fidelity profile, preservation, reduction, and distortion constitute distinct evaluative coordinates. Reduction can coexist with preservation or distortion, a fully defined capture can satisfy neither preservation nor distortion, and zero protected-pair divergence can coexist with reduction. The post-capture scalar therefore supplements rather than replaces the profile-relative Boolean assessments.

The deterministic and Monte Carlo results further show that perceptual support discordance and the restricted false-completion flag can condition symbolic outcomes without determining them. In the S4 isolation surface, 30,301 false-completion-positive cases occur with no positive symbolic distortion; across the three Monte Carlo strata, only about 71–72% of false-completion-positive cases are distortive under the declared profile. Symbolic fidelity depends jointly on the local encoder, perceptual conditioning, source geometry, capture provenance, and the selected comparison profile.

The resulting model provides an executable basis for studying later symbolic persistence or reuse only if those later relations are introduced through their own typed interfaces. Positive capture and current fidelity, by themselves, do not supply those downstream functions.

References

[1] K. Osmolovskyi, *Inner Manifestation Beyond Admissible Processing: A Formal Distinction in Cognitive Systems*, Zenodo, 2026. DOI: 10.5281/zenodo.19583268.

[2] K. Osmolovskyi, *Perceptual Stabilization through Historically Compressed Processing Organization*, Zenodo, 2026. DOI: 10.5281/zenodo.19814865.

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

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

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

[6] K. Osmolovskyi, *Sensitivity Analysis of Perceptual Stabilization under Historical Compression*, Zenodo, 2026. DOI: 10.5281/zenodo.20184845.

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