EXECUTABLE SEMANTIC SURVIVAL IN LANGUAGE-MODEL RELAYS

Anonymous authors
Paper under double-blind review

ABSTRACT

Repeated language-model communication can substantially change a message while preserving what a receiver does with it. Surface similarity therefore does not determine operational meaning, and one-step retention may not predict long-horizon semantic survival. We formalize this distinction by defining two messages as equivalent when a frozen receiver selects the same action in the same task context, inducing a context-conditioned executable semantic quotient. We characterize when this quotient evolves as a closed Markov process via strong lumpability and bound the error of approximate closure over finite horizons. To study long-horizon preservation, we stop a trajectory at its first action change and represent the surviving dynamics with a killed operator. Its powers recover the full first-passage survival law, revealing how equal one-step retention can hide different futures because surviving mass may occupy different within-fiber states. We further analyze observable repair using only public messages and receiver outputs, and use two frozen receivers to form survival bounds under evaluator ambiguity. Controlled constructions verify the theory. Empirically, in open-model relays up to 35B parameters, natural messages with identical current actions can still diverge under matched future rewrites, while killed-operator predictions outperform repeated one-step retention, including for held-out receivers. Together, these results show that semantic stability is a trajectory property: a message can be executable now yet be only metastably executable under repeated transmission.

1 INTRODUCTION

When one language model rewrites a request and another acts on it, the transmitted string is not itself the operational object. What matters is the action the receiver takes. A substantial rewrite can preserve the same tool call, intent, or referent, while a lexically similar rewrite can cross an action boundary. Under repeated transmission, this distinction becomes consequential: a message that is harmless now can move the relay toward a different future action. Cultural transmission studies document persistent transformation and amplified model priors (Perez et al., 2024; Ren et al., 2024; Guo et al., 2026), while communication measurements distinguish task use from surface structure (Lowe et al., 2019; Rita et al., 2022). Thus, as Fig. 1(a) illustrates, long-horizon analysis must track executable survival separately from surface drift.

To formalize this distinction, we introduce a *context-conditioned executable semantic quotient*. Specifically, given a fixed receiver and task context, we partition messages according to the action the receiver executes. Transmission still evolves over detailed message states, but the quotient observes only their executable fibers, separating paraphrase variation within a fiber from leakage across an action boundary. This turns language preservation into a precise question: when does the coarsened process itself evolve as a Markov chain, and when does its future still depend on which message within a fiber was reached?

This viewpoint also changes what counts as evidence. High surface similarity can accompany a different executable action, while a substantial rewrite can remain within the same action fiber. Likewise, one-step retention tells us how much probability remains in the current basin after one relay, but not which within-fiber states carry that probability or how stable they are under future rewrites. We therefore separate current action, within-fiber state, and future survival, and evaluate each on held-out trajectories rather than on strings alone.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

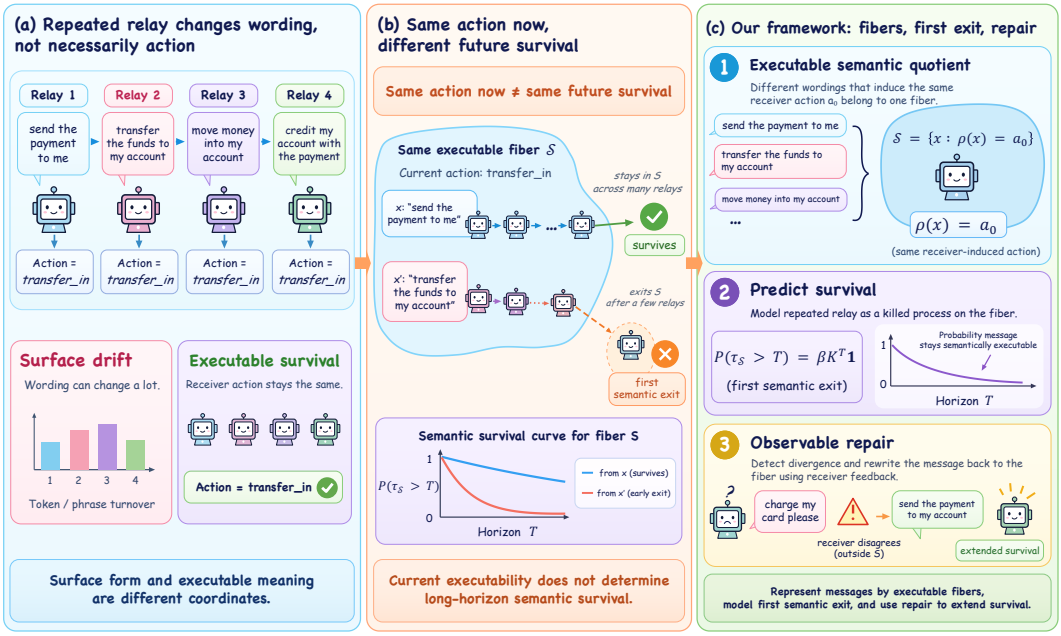

Figure 1: Executable semantic survival in language-model relays. (a) Rewriting can alter surface form while preserving receiver action. (b) Messages in the same current action fiber can induce different future survival laws. (c) We construct executable fibers, predict first passage with killed operators, and intervene through observable repair.

We begin by asking when the executable quotient is itself closed. Exact closure holds precisely under strong lumpability (Kemeny & Snell, 1960; Buchholz, 1994). When closure is only approximate, within-fiber discrepancies and later contraction control the accumulated error. Next, restricting transmission to the initial executable basin yields a killed operator. Its powers give the full survival curve, while its fundamental matrix gives the expected first-exit time. As Fig. 1(b) illustrates, equal current actions or equal one-step retention can still lead to different futures because surviving probability can occupy within-fiber states with different long-horizon stability.

We next formalize the repair mechanism in Fig. 1(c), using only public messages and frozen receiver outputs. If repair makes the killed operator elementwise larger, survival increases at every horizon. False triggers and ambiguous rewrites explain why this monotonicity can fail in practice. We then use two frozen receivers to form intersection and union bounds on semantic survival. Finally, when the transmission channel changes across generations, an ordered product replaces a stationary matrix power, and the measured operator drift bounds the resulting approximation error.

The empirical study follows the progression of the theory while testing all three components of Fig. 1. Models from 2B to 35B relay compositional references, intent and tool requests, and Banking77 utterances under preservation, mixing, and compression (Casanueva et al., 2020; Yang et al., 2025; Abdin et al., 2024). We first verify exact closure and first-passage identities in controlled finite systems, then evaluate quotient-operator predictions on held-out model-generated trajectories. To test whether the current executable state determines the future, we branch natural message pairs with identical current action vectors and expose them to the same future decoding randomness, directly realizing Fig. 1(b). Finally, we intervene with observable repair and test receiver transfer by replaying the trajectories through two frozen receivers excluded from quotient construction. Across these comparisons, the semantic event is fixed before evaluating future survival.

This paper makes three contributions:

- We define executable semantic quotients and characterize exact closure, approximate closure error, killed-operator survival, and time varying quotient products.
- We prove matched-error and surface-semantic separations, then give observable-repair domination and dual-receiver partial-identification results.
- We measure the theory through relays up to 35B, common-future within-fiber branches, and held-out receivers, identifying semantic metastability and its boundary.

2 Related Work

Emergent communication and language drift. Learned communication can be functionally successful without reproducing human linguistic structure, and structural metrics need not predict generalization or causal message use (Graesser et al., 2019; Chaabouni et al., 2019; Resnick et al., 2020; Lowe et al., 2019; Rita et al., 2022). Referential and multimodal games expose variation in how messages encode task distinctions (Evtimova et al., 2017; Bouchacourt & Baroni, 2018; Chaabouni et al., 2020; Kharitonov & Baroni, 2020). Recent work on natural-language grounding and interpretable communication improves partner compatibility, but it still separates surface form from task execution (Lazaridou et al., 2020; Karten et al., 2022; Lipinski et al., 2024; Bullard et al., 2021). In contrast, seeded iterated learning counters drift (Lu et al., 2020); our repair operator isolates an observable intervention inside a fixed transmission lineage and derives its horizon law.

Cultural transmission with language models. Repeated model-mediated transmission can amplify priors, depend on prompts, and exhibit nonmonotone changes in linguistic organization (Perez et al., 2024; Ren et al., 2024; Guo et al., 2026). Classical language evolution establishes error limits for generic transmission systems (Nowak et al., 1999), while iterated learning explains how bottlenecks reshape transmitted representations. Concision pressure can shorten learned messages (Kalinowska et al., 2022), and generative formulations connect language models with distributed world representation (Taniguchi et al., 2025). Building on these observations, our object is the context-conditioned executable partition of an open message space, and the primary outcome is semantic first passage under a fixed lineage.

State aggregation and metastability. Lumpability gives conditions under which a Markov process projects to a closed chain (Kemeny & Snell, 1960; Buchholz, 1994; Feret et al., 2012). Mixing and contraction control perturbation propagation (Levin et al., 2009), while killed chains and quasi-stationary analysis describe survival before absorption (Collet et al., 2013; Seneta, 2006). We connect these tools to executable language fibers, provide an approximate closure bound for transmission, and turn killed quotient products into held-out predictions of semantic survival.

3 Executable Semantic Transmission

We now formalize repeated language-model communication as a stochastic process over detailed message states, observed through the executable actions they induce.

Microstates and executable fibers. Fix a task context c containing the available actions and any information required for execution. A message microstate $x \in \mathcal{X}_c$ includes the public message together with the declared generator, receiver, decoding state, and retained public history. A frozen executable evaluator $E_c : \mathcal{X}_c \to \mathcal{A}_c$ maps the public message in that context to an action. Define

$$x \sim_c x' \quad \Longleftrightarrow \quad E_c(x) = E_c(x'), \tag{1}$$

and let $\phi_c : \mathcal{X}_c \to \mathcal{C}_c = \mathcal{X}_c/\sim_c$ be the quotient map. Appendix A gives the complete notation before the proofs.

Transmission. With the fibers fixed, a possibly generation-dependent kernel $P_t(x, \mathrm{d}y)$ describes one public rewrite. The detailed law at generation t is $\eta_{t+1} = \eta_t P_t$, and its observed semantic law is $\mu_t = \phi_{c\#}\eta_t$. We reserve Q_t for a kernel defined directly on $\mathcal{C}_c$. This distinction matters because two messages in one fiber can induce different distributions over next-generation fibers.

Semantic basin and first passage. To define the path event, let $a_0 = E_c(x_0)$ and let $S = \{x : E_c(x) = a_0\}$ be the executable basin of the initial message. The exit time is

$$\tau_S = \inf\{t \geq 0 : X_t \notin S\}. \tag{2}$$

Survival $\mathbb{P}(\tau_S > T)$ is an absorbing path event. A later return to S is recorded separately and does not erase the first exit. The empirical estimand is evaluator conditioned. With multiple frozen evaluators we report a partial-identification interval rather than collapsing their decisions.

Surface change. For any declared surface statistic $d(x, x_0)$, turnover and semantic exit remain different variables. We record exact normalized-string change, token count, type-token ratio, and normalized edit distance. None participates in the definition of ϕ_c.

4 QUOTIENT CLOSURE

Having defined the executable quotient, we next ask when its projected dynamics are self-contained: can the next semantic class be determined from the current class alone, without knowing the underlying message state? For a finite presentation, write $F_i = \phi^{-1}(i)$ and define the aggregate destination row from $x \in F_i$ as $p_\phi(x, j) = \sum_{y \in F_j} P(x, y)$.

Theorem 4.1 (Executable quotient closure). *There is a Markov kernel Q on $\mathcal{C}$ satisfying $\mathbb{P}(\phi(X_{t+1}) = j \mid X_t = x) = Q(\phi(x), j)$ for every microstate x if and only if*

$$\sum_{y \in F_j} P(x, y) = \sum_{y \in F_j} P(x', y), \qquad x, x' \in F_i, \quad i, j \in \mathcal{C}. \tag{3}$$

When it exists, $Q(i, j)$ is the common aggregate row.

Theorem 4.1 makes semantic Markov closure testable. Its proof and a path-law corollary appear in Appendix B. Thus it separates a valid quotient from a projection that merely assigns labels after each step.

For empirical kernels, define the within-fiber defect

$$\delta_t = \max_i \sup_{x, x' \in F_i} \mathrm{TV}\big(p_{\phi,t}(x, \cdot), p_{\phi,t}(x', \cdot)\big). \tag{4}$$

Choose $Q_t(i, \cdot)$ in the convex hull of the aggregate rows in F_i, and let $\alpha(Q_t)$ denote its Dobrushin contraction coefficient.

Theorem 4.2 (Approximate, time dependent closure). *Let μ_T be the true projected law and $\nu_T = \mu_0 Q_0 \cdots Q_{T-1}$. Then*

$$\mathrm{TV}(\mu_T, \nu_T) \le \sum_{s=0}^{T-1} \delta_s \prod_{r=s+1}^{T-1} \alpha(Q_r). \tag{5}$$

For a stationary Q, the bound is $\delta \sum_{r=0}^{T-1} \alpha(Q)^r$.

Consequently, an early defect is attenuated by later contraction, while persistent nonclosure accumulates linearly in the absence of contraction. Appendix C proves the result and gives tight two-fiber constructions.

5 SEMANTIC METASTABILITY

Even when the executable quotient is well defined, the central long-horizon question remains: how long does a trajectory stay within its initial executable basin before first exit?

Killed operator. Restrict P to the basin S and write $K = P_{S,S}$. The missing row mass is cross-fiber leakage. If β is an initial distribution on S, the killed dynamics retain paths only until their first semantic exit.

Theorem 5.1 (First-passage law). *For every finite horizon T,*

$$\mathbb{P}_\beta(\tau_S > T) = \beta K^T \mathbf{1}. \tag{6}$$

If $\rho(K) < 1$, then

$$\mathbb{E}_\beta[\tau_S] = \beta (I - K)^{-1} \mathbf{1}. \tag{7}$$

For irreducible K, the Perron root $\rho(K)$ gives the asymptotic survival rate, subject to periodicity.

Appendix D derives the survival, fundamental-matrix, and quasi-stationary statements. The finite-horizon identity is exact without asymptotics.

Theorem 5.2 (Equal one-step error, different futures). *For any $e \in (0, 1)$ and $T \ge 2$, two killed kernels can have the same initial one-step survival $1 - e$ and different T-step survival. In particular,*

$$K_{\text{fast}} = \begin{bmatrix} 1 - e & 0 \\ 0 & 0 \end{bmatrix}, \qquad K_{\text{store}} = \begin{bmatrix} 0 & 1 - e \\ 0 & 1 \end{bmatrix}. \tag{8}$$

give $(1-e)^T$ *and* $1-e$*, respectively, from the first state.*

The difference is where surviving mass is stored within the semantic basin. Appendix E gives positive irreducible perturbations and the exact separation envelope.

Proposition 5.3 (Surface-semantic separation)**.** *If a semantic fiber contains at least two surface forms, then any declared within-fiber mixing rate and any semantic leakage rate in* $[0,1]$ *can be realized independently by one transmission kernel.*

Consequently, neither lexical turnover nor string preservation identifies semantic survival. The construction and its spectrum are in Appendix E.

6 Observable Repair and Partial Identification

Having characterized semantic survival, we next ask two practical questions: can an observable intervention extend it, and how should survival be reported when different receivers define different executable basins?

Observable intervention. Transmission first proposes a public message. A detector with diagonal trigger operator D observes only that message and the declared receiver outputs. A repair kernel R then acts when triggered. The effective operator is

$$P^R = P\big[(I-D)+DR\big]. \tag{9}$$

This ordering distinguishes detection, repair, and latent semantic exit.

Theorem 6.1 (Repair domination)**.** *Let* K_0 *and* K_R *be the killed restrictions of the unrepaired and repaired operators to the same basin. If* $K_R \geq K_0$ *elementwise, then for every* T *and every nonnegative initial distribution* β*,*

$$\beta K_R^T \mathbf{1} \geq \beta K_0^T \mathbf{1}. \tag{10}$$

Whenever the expectations are finite, $\mathbb{E}_R[\tau_S] \geq \mathbb{E}_0[\tau_S]$.

The condition is sufficient and operational: repair improves every horizon when it redirects at least as much basin mass to every surviving microstate. Appendix F proves the result, characterizes strict gain, and gives minimal counterexamples when domination fails.

Two executable receivers. Let S_1 and S_2 be the basins defined by two frozen receivers. Their intersection and union produce directly observable path events.

Theorem 6.2 (Evaluator sandwich)**.** *If a target basin* $S_\star$ *satisfies* $S_1 \cap S_2 \subseteq S_\star \subseteq S_1 \cup S_2$*, then for every horizon,*

$$\mathbb{P}(\tau_{S_1 \cap S_2} > T) \leq \mathbb{P}(\tau_{S_\star} > T) \leq \mathbb{P}(\tau_{S_1 \cup S_2} > T). \tag{11}$$

Accordingly, Appendix H extends the sandwich to expected exit and time-dependent products. The interval width measures evaluator ambiguity.

Nonstationarity. For generation-indexed killed kernels, survival is $\beta K_0 \cdots K_{T-1}\mathbf{1}$. If K is a stationary reference, then any compatible submultiplicative norm gives

$$\left\| K_0 \cdots K_{T-1} - K^T \right\| \leq \sum_{s=0}^{T-1} \left\| K_s - K \right\|. \tag{12}$$

This converts measured operator drift into a bound on stationary prediction failure. Full proofs and sharper contracted forms appear in Appendix I.

7 Event-Specific State Sufficiency

We next characterize observable states sufficient for horizon-specific survival prediction. The executable quotient determines current action, but first-passage prediction may require additional within-fiber state, independent of any message parameterization.

Fix a horizon T and define the microstate survival value

$$h_T(x) = \mathbb{P}_x(\tau_S > T). \tag{13}$$

Let $\psi : \mathcal{X} \to \mathcal{Z}$ be any observable state representation, such as the current action, receiver agreement, a surface descriptor, or an augmented public history. Its horizon-specific fiber diameter is

$$\Delta_T(\psi) = \sup_{\psi(x)=\psi(x')} |h_T(x) - h_T(x')|. \tag{14}$$

Theorem 7.1 (Survival-state criterion)**.** *There is a function g_T such that $h_T = g_T \circ \psi$ if and only if $\Delta_T(\psi) = 0$. Moreover, under absolute loss,*

$$\inf_g \sup_{x\in\mathcal{X}} |h_T(x) - g(\psi(x))| = \frac{1}{2}\Delta_T(\psi). \tag{15}$$

The theorem distinguishes current semantic sufficiency from future semantic sufficiency. The action quotient ϕ is exact for current execution by definition, yet $\Delta_T(\phi)$ can be positive because different paraphrases in one action fiber lead to different future fibers. Theorem 4.1 supplies a structural route to zero diameter: if the declared quotient is strongly lumpable, its current class determines every future semantic law. Theorem 5.2 supplies the converse certificate for a coarser observation that retains only one-step error.

Corollary 7.2 (One-step non-identification)**.** *Let ψ_1 retain the initial class and its one-step survival probability. For every $e \in (0, 1)$ and $T \geq 2$, there is a transmission family with*

$$\Delta_T(\psi_1) \geq (1 - e) - (1 - e)^T. \tag{16}$$

Every predictor based only on ψ_1 then has worst-case error at least half this separation.

This lower bound gives the exact role of the killed operator. Its rows retain where surviving mass moves, whereas one-step accuracy retains only total surviving mass. A finite-horizon operator prediction can therefore improve even when asymptotic estimation of $\rho(K)$ is difficult.

State augmentation. When $\Delta_T(\phi) > 0$ or the empirical lumpability defect is large, one can augment the quotient with observable variables. Let $\psi'(x) = (\phi(x), u(x))$. Refinement can only shrink the diameter:

$$\Delta_T(\psi') \leq \Delta_T(\phi), \tag{17}$$

because every ψ' fiber is contained in a ϕ fiber. Our empirical microstate adds receiver agreement, exact seed equality, and repair status. Generation indexing handles systematic temporal change, while the remaining defect quantifies unresolved history dependence.

Prediction hierarchy. The theory induces three nested predictors. Repeated one-step retention keeps only a scalar row mass. A stationary killed operator also keeps destination structure within the basin. A time-indexed killed product additionally retains generation order. The first can fail under within-fiber storage, the second under operator drift, and the third under unresolved nonclosure. Equations 16, 43, and 5 give the corresponding error mechanisms. Appendix G proves Theorem 7.1, establishes the refinement property, and relates state diameter to approximate lumpability.

8 Experiments

The experiments proceed from exact verification to model-generated relays. Finite systems test the theoretical identities and boundary cases, while language-model experiments evaluate surface-semantic separation, long-horizon survival prediction, common-future divergence, observable repair, and transfer to held-out receivers.

Exact finite systems. We first enumerate finite killed operators where every quantity is known. The suite varies one-step error, within-basin storage, leakage, surface mixing, detector behavior, repair rows, lumpability defect, contraction, and time variation. It contains 821 parameter cells and evaluates every theorem without sampling error. Appendix J specifies the matrices and grids.

Executable relay tasks. Next, each episode begins with a public English message and eight candidate actions. Compositional reference identifies an object from color, material, shape, and position. Intent routing uses a banking ontology; tool routing uses sixteen common agent actions; Banking77 supplies held-out natural utterances (Casanueva et al., 2020). The receiver must emit one candidate identifier. The original receiver action defines the episode-specific semantic fiber, while the dataset label separately measures initial and final task correctness. Task construction is detailed in Appendix K.

Transmission conditions. We then let a fixed lineage repeatedly rewrite its current public message for 16 or 32 generations. Conservative transmission requests clear preservation, mixing requests substantially different wording and syntax, and compression requests the shortest unambiguous form. Rewriters span Qwen-family configurations at 2B, 4B, 8B, 14B, and 35B scale plus Phi-3.5-mini (Yang et al., 2025; Qwen Team, 2026; Abdin et al., 2024). Decoding temperature is 0.4. Main experiments use frozen 14B and 35B receivers and a 35B repairer. The smaller panels use independent 2B and 4B receivers. Prompts and declared decoding fields appear in Appendix L.

Observable repair. After every proposed rewrite, the two receivers act independently. Repair triggers exactly when their public actions disagree or either action is invalid. The repairer receives the current message and the stored initial message, then produces a message rather than an action. It never receives the gold candidate or a semantic-exit indicator. We record the proposed and repaired messages and action pairs, enabling immediate mediation and paired policy analyses in Appendix R.

Estimation. Finally, task instances are assigned deterministically to design and held-out halves, with the split shared by every transmission condition. Design trajectories estimate quotient rows, killed operators, generation-indexed retention, and approximate-lumpability defects. Held-out trajectories evaluate survival at generations 1, 2, 4, 8, 16, and 32. We report intersection and union events for the receiver sandwich. Repair contrasts use paired task instances and seed bootstrap intervals. Exact one-step matches require equality of integer survivor counts, so no matching tolerance or threshold is selected. Appendix M defines every statistic and Appendix N gives the estimators.

Within-fiber common futures. For each 35B mixing trajectory, we select the largest-edit natural paraphrase that preserves the seed's action vector. Seed and paraphrase are rechecked under an independent 2B and 4B receiver pair, then each is continued by the same 35B rewriter under eight matched future decoding tapes. The analysis retains 160 message pairs whose action vectors agree on every repetition. This intervention estimates the survival diameter in Equation 14 while holding task, current executable state, model, and future seeds fixed.

Held-out receivers. Qwen3 8B and Phi-3.5 Mini replay every message from the main and 35B long panels. These receivers do not define the original fibers, repair triggers, or model selection. Their 9,984 complete trajectory-receiver records test whether the operator hierarchy transfers to new executable partitions. Appendix Q gives both protocols and complete results.

9 Results and Implications

Observation 1: Exact certificates. We begin with the finite controls, where every quantity is known. All exact systems agree with the theory. At one-step error $e = 0.10$, the repeated-exposure and storage kernels both retain 0.900 after one generation, but their survival values are 0.430 and 0.900 at generation 8, then 0.034 and 0.900 at generation 32. Fig. 2 shows the complete horizon curve. All 30 approximate-lumpability cells and all 32 time-varying cells satisfy their respective bounds. Every one of the 60 repair-domination cells has nonnegative survival gain; outside domination, 91 cells have negative gain. Appendix J and Appendix U report the complete exact sweeps.

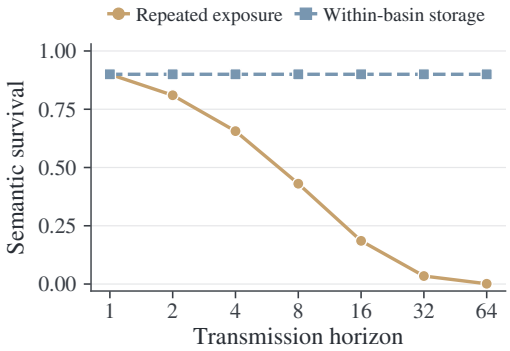

Figure 2: Equal one-step retention does not determine later survival.

Observation 2: Surface change and semantic survival separate. We next turn to model-generated relays. The language-model relays contain 7,776 complete trajectories, including 2,688 trajectories at horizon 32. Surface rewriting is often extensive while executable behavior remains stable. In the main 35B mixing condition without repair, attribute messages have turnover 0.907 and intersection survival 0.812 at generation 16. Repair changes these values to 0.919 and 0.958. At generation 32, attribute survival remains 0.828 despite turnover 0.810, whereas intent survival is 0.609 at turnover 0.635. Consistent with Fig. 1(a), Fig. 3 shows that surface change cannot order the semantic basins. Appendix O gives the complete curves and task-level measurements.

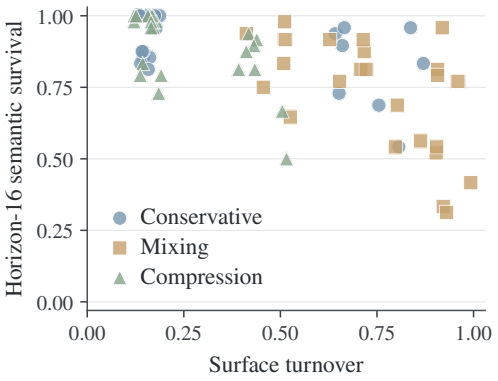

Figure 3: Surface turnover and executable survival are separate coordinates.



Figure 4: Held-out survival prediction error for one-step, stationary, and time-indexed operators.

Observation 3: Killed operators predict future survival. Across 720 held-out operator groups, repeating the design one-step retention has mean absolute error 0.129. The stationary killed operator reduces this to 0.057, a paired reduction of 0.072 with interval $[0.064, 0.080]$. The generation-indexed product reaches 0.050, improving on the stationary power by 0.007 with interval $[0.005, 0.009]$. In the 32-generation 8B panel, one-step repetition has error 0.232, compared with 0.078 for the stationary operator and 0.059 for the time product. Fig. 4 summarizes these comparisons. The threshold-free matching analysis finds 2,366 condition pairs with identical design sample size and one-step survivor count, yet later held-out survival differs in 1,473 pairs. In the 35B horizon-32 panel, two intent conditions each retain 32 of 33 design trajectories at generation 1, but held-out survival is 0.968 for repaired compression and 0.484 for unrepaired mixing. These measurements instantiate the storage mechanism of Theorem 5.2 in neural transmission. Appendix P reports all operator estimates; Appendix S reports the exact empirical matches.

Observation 4: Equal current actions lead to different common futures. We next remove the remaining group-comparison ambiguity with a paired intervention. For 160 natural message pairs, the seed and a high-edit paraphrase produce exactly the same action vector under both frozen receivers at the branch point. Each message then receives eight matched future decoding tapes from the same 35B rewriter. Mean absolute survival separation grows from 0.077 after one future step to 0.144 after four and eight steps. At horizon 8, 35.6% of pairs have different empirical survival laws, and four pairs separate completely at 1.0 versus 0.0. The mean signed seed-minus-paraphrase contrast is 0.056 with interval $[0.013, 0.102]$, while pair-level reversals show that canonical wording is not uniformly favored. The held-out receiver replay independently changes the quotient map. Across Qwen3 8B and Phi-3.5 Mini, repeated one-step retention has prediction error 0.113, while the stationary killed operator has error 0.048 and the time product 0.050. Fig. 5 combines the intervention and receiver-transfer results. Appendix Q reports all message pairs, horizons, models, and tasks.

Observation 5: Observable repair extends metastable survival. We then test the intervention predicted by the theory. The effect is strongest under mixing, where receiver disagreement exposes ambiguous messages. At horizon 32 with the 35B rewriter, repair changes intersection survival by 0.125 on attribute reference, 0.219 on Banking77, 0.266 on intent routing, and 0.156 on tool routing. Their paired intervals are respectively $[0.031, 0.219]$, $[0.094, 0.344]$, $[0.156, 0.391]$, and $[0.063, 0.250]$. The independently seeded 8B long panel gives positive effects on attribute, Banking77, and intent. Across all panels, 3,287 observable triggers yield 1,313 immediate rescues and 116 induced exits under the intersection event. This asymmetry explains the positive horizon effect, while the induced exits locate departures from elementwise domination. Appendix R reports every paired contrast and mediation cell.

Observation 6: Receiver ambiguity remains measured. Finally, we retain evaluator uncertainty instead of selecting one receiver after seeing the outcomes. The mean union-minus-intersection width is 0.060 in the main panel and 0.056 in the 35B long panel. The interval is therefore informative rather than vacuous, while preserving receiver disagreement as an explicit uncertainty quantity. Appendix H defines the event construction; Appendix O reports all horizon curves; Appendix T reports task behavior; Appendix V gives complete public-message examples.

Table 1: Core empirical quantities from the main experiments.

Test	Condition	Primary quantity	Value
Matched kernels	$e = 0.10$, horizon 8	survival (fast / stored)	0.430 / 0.900
Surface / event	35B mixing, gen. 16	turnover / intersection survival	0.907 / 0.812
Operator prediction	held-out groups	one-step / stationary / time MAE	0.129 / 0.057 / 0.050
Common futures	160 pairs, horizon 8	mean absolute survival separation	0.144
Observable repair	35B, $h = 32$	intent survival effect	+0.266
Held-out receivers	Qwen3 / Phi-3.5	one-step / stationary / time MAE	0.113 / 0.048 / 0.050

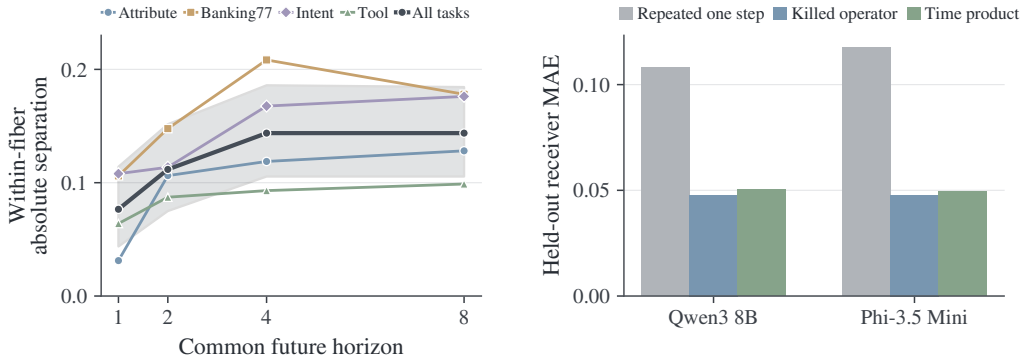

Figure 5: Common-future separation and held-out-receiver prediction. (a, left) Common-future test: within-fiber survival separation across matched horizons. (b, right) Held-out receivers: prediction error for repeated one-step retention, the stationary killed operator, and the time-indexed product.

Observation 7: Preservation is event specific. The executable quotient replaces undifferentiated language similarity with a declared event. Tool execution, intent routing, and referent selection induce different fibers on the same message space. Preservation means remaining inside the initial action basin; naturalness, compression, and lexical continuity remain separate coordinates.

Observation 8: Metastability is a trajectory property. The experiments show why a one-generation evaluation can appear stable even when the process has a short remaining lifetime. Surviving probability mass may concentrate on messages that are fragile under the next rewrite, or on messages that remain deep inside the basin. The killed operator represents this occupancy. Its finite powers predict the full survival curve, while the Perron value summarizes persistent stationary regimes. Exact matched-error constructions and 2,366 empirical matches expose the same distinction at different levels of control. Moreover, repair changes the next relay input, so its path effect extends beyond immediate mediation.

Observation 9: Nonstationarity is measurable. Finally, repeated rewriting changes within-fiber occupancy under a fixed model and prompt. A stationary quotient averages this movement, while a generation-indexed product exposes temporal drift in the remaining prediction error.

10 Conclusion

We develop executable semantic survival as a framework for studying meaning across repeated language-model communication. Semantic equivalence is defined by the action of a frozen receiver in a fixed context, separating surface rewriting from executable behavior. Strong lumpability characterizes when the induced quotient has closed Markov dynamics, while killed operators recover the first-passage law of semantic escape and explain why equal current actions or one-step retention can still lead to different futures. Observable repair provides an intervention on these trajectories, and dual-receiver bounds retain evaluator uncertainty rather than collapsing it to a single semantic judgment. Across finite controls and language-model relays, extensive surface change can coexist with stable execution, messages with identical current actions can diverge under matched future rewrites, and killed-operator predictions improve on repeated one-step retention, including under held-out receivers. These results show that semantic stability is a trajectory property and motivate evaluating communication protocols through long-horizon executable survival, destination structure, and sensitivity to repair rather than through surface similarity or one-generation accuracy alone.

AI USAGE STATEMENT

Generative AI (GPT-6 Astra) assists parts of manuscript preparation, including text editing, code drafting and debugging, figure layout, and consistency checks for equations. The authors design the study, derive the claims, prove the theorems, and perform the substantive analyses; AI assistance does not generate experimental evidence or determine the scientific conclusions. Every suggested change and computational aid is reviewed against the intended implementation, and all reported measurements are checked against the recorded outputs. The authors revise all assisted material and remain responsible for the paper's text, mathematics, results, code, and released artifacts.

REFERENCES

Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat Behl, Alon Benhaim, et al. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*, 2024.

Diane Bouchacourt and Marco Baroni. How agents see things: On visual representations in an emergent language game. *arXiv preprint arXiv:1808.10696*, 2018.

Peter Buchholz. Exact and ordinary lumpability in finite markov chains. *Journal of Applied Probability*, 31(1):59–75, 1994. doi: 10.2307/3215235.

Kalesha Bullard, Douwe Kiela, Franziska Meier, Joelle Pineau, and Jakob Foerster. Quasi-equivalence discovery for zero-shot emergent communication. *arXiv preprint arXiv:2103.08067*, 2021.

Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. Efficient intent detection with dual sentence encoders. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pp. 38–45. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.nlp4convai-1.5.

Rahma Chaabouni, Eugene Kharitonov, Emmanuel Dupoux, and Marco Baroni. Anti-efficient encoding in emergent communication. *arXiv preprint arXiv:1905.12561*, 2019.

Rahma Chaabouni, Eugene Kharitonov, Diane Bouchacourt, Emmanuel Dupoux, and Marco Baroni. Compositionality and generalization in emergent languages. *arXiv preprint arXiv:2004.09124*, 2020.

Pierre Collet, Servet Martínez, and Jaime San Martín. *Quasi-Stationary Distributions: Markov Chains, Diffusions and Dynamical Systems*. Springer, 2013. doi: 10.1007/978-3-642-33131-2.

Katrina Evtimova, Andrew Drozdov, Douwe Kiela, and Kyunghyun Cho. Emergent communication in a multi-modal, multi-step referential game. *arXiv preprint arXiv:1705.10369*, 2017.

Jérôme Feret, Thomas A. Henzinger, Heinz Koeppl, and Tatjana Petrov. Lumpability abstractions of rule-based systems. *Theoretical Computer Science*, 431:137–164, 2012. doi: 10.1016/j.tcs.2011.12.059.

Laura Graesser, Kyunghyun Cho, and Douwe Kiela. Emergent linguistic phenomena in multi-agent communication games. *arXiv preprint arXiv:1901.08706*, 2019.

Dongxin Guo, Jikun Wu, and Siu Ming Yiu. Model collapse as cultural evolution. *arXiv preprint arXiv:2605.23054*, 2026.

Aleksandra Kalinowska, Elnaz Davoodi, Florian Strub, Kory W. Mathewson, Ivana Kajić, Michael Bowling, Todd D. Murphey, and Patrick M. Pilarski. Over-communicate no more: Situated rl agents learn concise communication protocols. *arXiv preprint arXiv:2211.01480*, 2022.

Seth Karten, Mycal Tucker, Huao Li, Siva Kailas, Michael Lewis, and Katia Sycara. Interpretable learned emergent communication for human-agent teams. *arXiv preprint arXiv:2201.07452*, 2022.

John G. Kemeny and J. Laurie Snell. *Finite Markov Chains*. Van Nostrand, 1960.

Eugene Kharitonov and Marco Baroni. Emergent language generalization and acquisition speed are not tied to compositionality. *arXiv preprint arXiv:2004.03420*, 2020.

Angeliki Lazaridou, Anna Potapenko, and Olivier Tieleman. Multi-agent communication meets natural language: Synergies between functional and structural language learning. *arXiv preprint arXiv:2005.07064*, 2020.

David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2009.

Olaf Lipinski, A. J. Sobey, Federico Cerutti, and Timothy J. Norman. Speaking your language: Spatial relationships in interpretable emergent communication. *arXiv preprint arXiv:2406.07277*, 2024.

Ryan Lowe, Jakob Foerster, Y-Lan Boureau, Joelle Pineau, and Yann Dauphin. On the pitfalls of measuring emergent communication. *arXiv preprint arXiv:1903.05168*, 2019.

Yuchen Lu, Soumye Singhal, Florian Strub, Olivier Pietquin, and Aaron Courville. Countering language drift with seeded iterated learning. *arXiv preprint arXiv:2003.12694*, 2020.

Martin A. Nowak, David C. Krakauer, and Andreas Dress. An error limit for the evolution of language. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 1999. doi: 10.1098/rspb.1999.0898.

Jeremy Perez, Corentin Leger, Marcela Ovando-Tellez, Chris Foulon, Joan Dussauld, Pierre-Yves Oudeyer, and Clement Moulin-Frier. Cultural evolution in populations of large language models. *arXiv preprint arXiv:2403.08882*, 2024.

Qwen Team. Qwen3.5: Towards native multimodal agents, February 2026.

Yi Ren, Shangmin Guo, Linlu Qiu, Bailin Wang, and Danica J. Sutherland. Bias amplification in language model evolution: An iterated learning perspective. *arXiv preprint arXiv:2404.04286*, 2024.

Cinjon Resnick, Abhinav Gupta, Jakob Foerster, Andrew M. Dai, and Kyunghyun Cho. Capacity, bandwidth, and compositionality in emergent language learning. *arXiv preprint arXiv:1910.11424*, 2020.

Mathieu Rita, Corentin Tallec, Paul Michel, Jean-Bastien Grill, Olivier Pietquin, Emmanuel Dupoux, and Florian Strub. Emergent communication: Generalization and overfitting in lewis games. *arXiv preprint arXiv:2209.15342*, 2022.

Eugene Seneta. *Non-negative Matrices and Markov Chains*. Springer, 2006.

Tadahiro Taniguchi, Ryo Ueda, Tomoaki Nakamura, Masahiro Suzuki, and Akira Taniguchi. Generative emergent communication: Large language model is a collective world model. *arXiv preprint arXiv:2501.00226*, 2025.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.

A Notation and Semantic Conventions

Table 2 lists the objects used throughout the theory and measurements. The context c is fixed within an episode and contains the candidate action set. A microstate contains every declared variable used by the next transmission kernel. The message is public, while the semantic fiber is induced by executing a frozen receiver in the fixed context.

Path-event convention. The event $\{\tau_S > T\}$ requires every state at generations 0 through T to remain in S. Once a path leaves, later reentry does not change first-passage survival. We separately retain the pointwise action at every generation, so recovery after exit remains measurable.

Table 2: Notation used in the theory and experiments.

Symbol	Type	Meaning
c	context	Fixed candidate actions and execution information for one episode
$\mathcal{X}_c$	state space	Declared message microstates under context c
E_c	map	Frozen executable receiver from a message state to an action
ϕ_c	map	Quotient map induced by equality of receiver actions
F_i	set	Fiber of all microstates assigned to semantic class i
P_t	kernel	Detailed generation t transmission operator on microstates
Q_t	kernel	Closed or approximate quotient operator on semantic classes
η_t	distribution	Detailed microstate law at generation t
μ_t	distribution	Projected semantic-class law $\phi_{c\#}\eta_t$
S	set	Basin that executes the initial action
K_t	operator	Killed restriction $(P_t)_{S,S}$
τ_S	stopping time	First generation outside S
δ_t	scalar	Maximum within-fiber aggregate-row defect
$\alpha(Q_t)$	scalar	Dobrushin contraction coefficient of Q_t
D	operator	Observable diagonal repair trigger
R	kernel	Public-message repair operator
$\rho(K)$	scalar	Perron survival value of a stationary killed operator

Receiver convention. Every receiver is frozen for an entire run. Semantic fibers can vary across contexts because candidate actions vary. They are therefore indexed by c even when the index is suppressed. For two receivers, we preserve both induced partitions and report intersection and union path events.

B Proofs for Exact Quotient Closure

We give the full proof of Theorem 4.1 and record its consequence for semantic path laws.

Proof of Theorem 4.1. Suppose first that a kernel Q with the stated conditional law exists. Fix classes i, j and two microstates $x, x' \in F_i$. Since $\phi(x) = \phi(x') = i$,

$$\begin{aligned}\sum_{y\in F_j} P(x,y) &= \mathbb{P}(\phi(X_{t+1}) = j \mid X_t = x) \\ &= Q(i,j).\end{aligned} \tag{18}$$

Applying the same equality at x' gives

$$\sum_{y\in F_j} P(x',y) = Q(i,j), \tag{19}$$

which proves Equation 3.

Conversely, assume Equation 3. For any i, j, select an arbitrary $x_i \in F_i$ and define

$$Q(i,j) = \sum_{y\in F_j} P(x_i,y). \tag{20}$$

The aggregate-row equality makes this definition independent of the representative. Nonnegativity follows from P, and summing over the partition yields

$$\sum_j Q(i,j) = \sum_j \sum_{y\in F_j} P(x_i,y) = 1. \tag{21}$$

Thus Q is a Markov kernel. For any $x \in F_i$, Equation 3 and the definition give the required conditional law. This proves sufficiency and uniqueness of every aggregate row. □

Corollary B.1 (Semantic path law)**.** *Under the conditions of Theorem 4.1, if $Z_t = \phi(X_t)$ then (Z_t) is Markov with transition kernel Q for every initial microstate distribution.*

Proof. For any semantic history with positive probability, condition on the detailed state at time t. Theorem 4.1 makes the distribution of Z_{t+1} equal to the same row $Q(Z_t, \cdot)$ for every detailed state compatible with that semantic history. Averaging over the conditional detailed-state law leaves that row unchanged. □

Why a label sequence is not sufficient. Without Equation 3, one can still record $Z_t = \phi(X_t)$ after every transition. The resulting observation need not be a Markov chain because the conditional distribution of Z_{t+1} can depend on which surface form inside F_{Z_t} was reached. The matched-error construction in Appendix E is the smallest instance of this hidden within-fiber memory.

C Approximate Lumpability and Contraction

For probability measures p, q on a finite state space, write $\mathrm{TV}(p, q) = \frac{1}{2}\|p - q\|_1$. The Dobrushin coefficient of a stochastic kernel is

$$\begin{aligned} \alpha(Q) &= \sup_{p \neq q} \frac{\mathrm{TV}(pQ, qQ)}{\mathrm{TV}(p, q)} \\ &= \max_{i,i'} \mathrm{TV}\big(Q(i, \cdot), Q(i', \cdot)\big). \end{aligned} \tag{22}$$

Lemma C.1 (One-step aggregation error)**.** *If every aggregate microstate row in fiber F_i lies within δ_t of $Q_t(i, \cdot)$, then*

$$\mathrm{TV}(\mu_{t+1}, \mu_t Q_t) \leq \delta_t. \tag{23}$$

Proof. Disintegrate the microstate law as $\eta_t(\mathrm{d}x) = \mu_t(i)\eta_t(\mathrm{d}x \mid F_i)$. Its next projected row from class i is the convex mixture of $p_{\phi,t}(x, \cdot)$ under $\eta_t(\cdot \mid F_i)$. Convexity of total variation keeps this row within δ_t of $Q_t(i, \cdot)$. Averaging across i preserves the same bound. □

Proof of Theorem 4.2. Let $e_t = \mathrm{TV}(\mu_t, \nu_t)$. Add and subtract $\mu_t Q_t$, use Lemma C.1, and apply Dobrushin contraction:

$$e_{t+1} \leq \mathrm{TV}(\mu_{t+1}, \mu_t Q_t) + \mathrm{TV}(\mu_t Q_t, \nu_t Q_t) \tag{24}$$

$$\leq \delta_t + \alpha(Q_t) e_t. \tag{25}$$

Since $e_0 = 0$, iterating Equation 25 gives

$$e_T \leq \delta_{T-1} + \alpha(Q_{T-1})\delta_{T-2} + \cdots + \delta_0 \prod_{r=1}^{T-1} \alpha(Q_r), \tag{26}$$

which is Equation 5 after reindexing. □

Stationary special cases. If $\alpha(Q) = 1$, Equation 5 becomes $\mathrm{TV}(\mu_T, \nu_T) \leq T\delta$. If $\alpha(Q) < 1$, it becomes

$$\mathrm{TV}(\mu_T, \nu_T) \leq \delta \frac{1 - \alpha(Q)^T}{1 - \alpha(Q)} \leq \frac{\delta}{1 - \alpha(Q)}. \tag{27}$$

Tightness. Take two semantic states and two microstates inside the first fiber whose aggregate rows differ by 2δ in signed mass across the two destinations. A detailed-state process can select the extreme row at each generation. When Q has no contraction, the signed deviations align and the first-order error grows as $T\delta$ until probability constraints bind. The exact sweep in Appendix J includes both this regime and contracted variants.

D Killed Operators and First Passage

Let P be a Markov kernel on a finite microstate space and S the semantic basin. The restriction $K = P_{S,S}$ is substochastic because transitions from S to its complement are removed.

Proof of Theorem 5.1. For $T = 1$, the retained mass is $\beta K\mathbf{1}$, which is exactly the probability that the first transition remains in S. Suppose βK^T is the unnormalized state distribution at time T of paths that have never left S. Multiplying by K retains exactly those paths whose next transition also remains in S. Induction gives

$$\mathbb{P}_\beta(\tau_S > T) = \beta K^T \mathbf{1}. \tag{28}$$

For a nonnegative integer-valued exit time, the tail-sum identity gives

$$\mathbb{E}_\beta[\tau_S] = \sum_{T=0}^{\infty} \mathbb{P}_\beta(\tau_S > T). \tag{29}$$

When $\rho(K) < 1$, the Neumann series converges and Equations 28 and 29 yield

$$\mathbb{E}_\beta[\tau_S] = \beta \sum_{T=0}^{\infty} K^T \mathbf{1} = \beta(I - K)^{-1}\mathbf{1}. \tag{30}$$

For irreducible K, Perron-Frobenius theory supplies a positive eigenvalue $\rho(K)$ and nonnegative left and right eigenvectors. The leading term of K^T has exponential rate $\rho(K)^T$, with the standard periodic refinement when the peripheral spectrum is not simple. □

Proposition D.1 (Hazard representation)**.** *If $s_t = \mathbb{P}(\tau_S > t)$ is positive, define $h_t = 1 - s_{t+1}/s_t$. Then*

$$s_T = \prod_{t=0}^{T-1} (1 - h_t). \tag{31}$$

For a one-state killed quotient, $1 - h_t$ is the generation-indexed killed row sum.

Proof. The definition implies $s_{t+1} = s_t(1 - h_t)$. Multiplying the recursion and using $s_0 = 1$ gives the claim. □

Return after exit. The killed operator intentionally removes every path at its first exit. A return probability instead requires the full operator P and an event such as $\mathbb{P}(X_t \in S, \tau_S \leq t)$. Our trajectory records permit both quantities, but all survival theorems use the first-passage convention.

E Matched Error and Surface Separation

Proof of Theorem 5.2. Let the initial distribution be $\beta = (1, 0)$. Both matrices in Equation 8 have row sum $1 - e$ in the initial state, so

$$\beta K_{\text{fast}}\mathbf{1} = \beta K_{\text{store}}\mathbf{1} = 1 - e. \tag{32}$$

For the fast kernel, surviving mass remains in the first state and loses a factor $1 - e$ at every step, hence

$$\beta K_{\text{fast}}^T \mathbf{1} = (1 - e)^T. \tag{33}$$

For the storage kernel, the first step places all surviving mass in the second state, whose row sum is one, hence

$$\beta K_{\text{store}}^T \mathbf{1} = 1 - e \quad \text{for every } T \geq 1. \tag{34}$$

The difference is positive for $T \geq 2$. □

Strictly positive perturbation. For $\epsilon > 0$, add ϵ to the zero within-basin entries and adjust the positive entries and exit probabilities continuously so that both initial row sums remain $1 - e$. The finite-horizon survival functions are polynomials in the matrix entries. Therefore, for sufficiently small ϵ, the separation in Equations 33 and 34 remains positive while both killed kernels become irreducible and strictly positive.

Proof of Proposition 5.3. Let the target fiber contain surface states x_1 and x_2, and let z be outside the fiber. Choose any two-state stochastic kernel W_m with the declared mixing rate m. For example, a symmetric kernel can be parameterized by its nontrivial eigenvalue. Define

$$\begin{aligned} P_{m,e}(x_i, x_j) &= (1-e)W_m(i,j), \\ P_{m,e}(x_i, z) &= e. \end{aligned} \tag{35}$$

Make z absorbing. The conditional movement among surviving surface forms is exactly W_m, while the row mass leaving the semantic fiber is exactly e. Since m appears only in W_m and e only in the multiplicative survival factor, they can be selected independently. □

Consequences. At $e = 0$, a kernel can alternate deterministically between distinct paraphrases while executable survival remains one. At maximal self-retention inside the fiber, surface turnover can be zero while any $e > 0$ creates semantic leakage. The statement applies to any surface statistic that distinguishes two forms in the same executable fiber.

Finite-horizon envelope. For a fixed one-step survival $1 - e$, subsequent survival lies between zero and $1 - e$ over unrestricted substochastic kernels. The storage construction attains the upper endpoint. A kernel that sends surviving mass to a state with zero row sum attains zero at horizon two. The matched finite systems in Appendix J use the smoother fast-decay comparison because it separates every later horizon.

F Observable Repair Theory

The detector and repairer act after transmission. Let D be diagonal with $D(x, x)$ equal to the trigger probability for public post-transmission state x. Conditional on no trigger, the state remains x; conditional on a trigger, R produces the repaired public state. This yields Equation 9.

Lemma F.1 (Order preservation)**.** *For nonnegative square matrices $A \geq B$ elementwise, $A^T \geq B^T$ for every integer $T \geq 1$.*

Proof. The claim is immediate for $T = 1$. Suppose $A^T \geq B^T$. Then

$$A^{T+1} = A^T A \geq B^T A \geq B^T B = B^{T+1}, \tag{36}$$

where each multiplication preserves elementwise order because all factors are nonnegative. □

Proof of Theorem 6.1. Apply Lemma F.1 to $K_R \geq K_0$. Left multiplication by $\beta \geq 0$ and right multiplication by $\mathbf{1} \geq 0$ preserve the order, giving Equation 10. Summing the ordered survival probabilities gives the expected-exit result whenever both sums are finite. □

Corollary F.2 (Strict horizon gain)**.** *If $K_R \geq K_0$ and some positive-mass path reachable from β uses an entry where the inequality is strict before horizon T, then $\beta K_R^T \mathbf{1} > \beta K_0^T \mathbf{1}$ provided the remaining repaired path has positive basin retention.*

Domination failure. Consider two safe microstates. An unrepaired kernel can retain each state with probability one. A detector that always triggers and a repairer that maps the first safe state outside the basin gives a repaired killed row of zero at that state. The repaired operator does not dominate, and survival decreases. A second construction begins from an unsafe proposed message and maps it into the basin, increasing immediate survival. These examples show that the intervention sign is determined by the repair kernel relative to the basin, not by the presence of a repair label.

Immediate mediation counts. Our empirical decomposition records whether a triggered repair changes a pre-repair exit into a post-repair survivor, changes a survivor into an exit, preserves survival, or preserves exit. These four cells estimate the local direction of the intervention. The paired trajectory contrast estimates the policy effect after subsequent rewrites have adapted to the repaired message.

G Event-Specific Sufficiency Proofs

Proof of Theorem 7.1. If $h_T = g_T \circ \psi$, then any x, x' in the same ψ fiber have equal h_T values, so $\Delta_T(\psi) = 0$. Conversely, if the diameter is zero, h_T is constant on every fiber. Define $g_T(z)$ to be that common value for any nonempty fiber $\psi^{-1}(z)$, with an arbitrary value on empty fibers. Then $h_T = g_T \circ \psi$.

For the minimax statement, fix a fiber F_z and write

$$m_z = \inf_{x \in F_z} h_T(x), \qquad M_z = \sup_{x \in F_z} h_T(x). \tag{37}$$

Any scalar prediction $g(z)$ has maximum error on that fiber at least $(M_z - m_z)/2$. Choosing the midpoint $g(z) = (M_z + m_z)/2$ attains this radius, with the usual limiting interpretation when extrema are not attained. Taking the supremum over fibers gives Equation 15. □

Proof of Corollary 7.2. The two kernels in Equation 8 share the initial class and one-step survival $1 - e$, so the observation ψ_1 is identical. Equations 33 and 34 give horizon-T values $(1 - e)^T$ and $1 - e$. Their difference is a lower bound on the fiber diameter. The minimax statement follows from Theorem 7.1. □

Proposition G.1 (Refinement monotonicity)**.** *If ψ' refines ψ, meaning $\psi'(x) = \psi'(x')$ implies $\psi(x) = \psi(x')$, then $\Delta_T(\psi') \le \Delta_T(\psi)$.*

Proof. Every pair over which the supremum for $\Delta_T(\psi')$ is taken is also an admissible pair for the supremum defining $\Delta_T(\psi)$. □

Connection to approximate closure. Let two microstates share a quotient class and suppose their one-step aggregate rows differ by at most δ in total variation. Applying Theorem 4.2 from each initial point bounds the difference between their horizon-T quotient laws by twice the approximation radius around the same quotient prediction. For any basin event,

$$|h_T(x) - h_T(x')| \le 2 \sum_{s=0}^{T-1} \delta_s \prod_{r=s+1}^{T-1} \alpha(Q_r). \tag{38}$$

Thus approximate lumpability controls event-specific diameter, while an observed positive diameter certifies that the proposed state representation is not exact for that horizon.

Event dependence. A representation can be sufficient for one basin event and insufficient for another. For example, a receiver action may fully determine survival of a coarse tool family while leaving uncertainty among fine-grained tools. The criterion in Theorem 7.1 applies separately to each declared event and horizon.

H Dual-Receiver Partial Identification

Proof of Theorem 6.2. The set inclusion $S_1 \cap S_2 \subseteq S_\star \subseteq S_1 \cup S_2$ lifts to path-event inclusion:

$$\begin{aligned} \{X_{0:T} \in S_1 \cap S_2\} &\subseteq \{X_{0:T} \in S_\star\} \\ &\subseteq \{X_{0:T} \in S_1 \cup S_2\}. \end{aligned} \tag{39}$$

Taking probabilities gives Equation 11. Summing the survival probabilities over T gives the same order for finite expected exit times. □

Empirical events. For each receiver $r \in \{1, 2\}$, let $A_{r,t}$ indicate that its generation t action equals its generation zero action. We compute

$$L_T = \prod_{t=0}^{T} (A_{1,t} A_{2,t}), \tag{40}$$

$$U_T = \prod_{t=0}^{T} \max(A_{1,t}, A_{2,t}). \tag{41}$$

Then $\mathbb{E}[L_T]$ and $\mathbb{E}[U_T]$ are the intersection and union survival probabilities. Their difference is an observed ambiguity width.

Context specificity. Receiver agreement is evaluated within the same task context and candidate ordering. If a context changes, the induced action maps and quotient fibers change as well. The experiment therefore holds context fixed along each trajectory and randomizes new contexts only across episodes.

Fig. 6 shows the empirical interval widths in the main panel.

More than two receivers. For m frozen receivers, the intersection of all basins and union of all basins give the analogous outer interval. Intermediate vote-defined basins can tighten the interval under additional assumptions, but they are not required for the set-inclusion result.

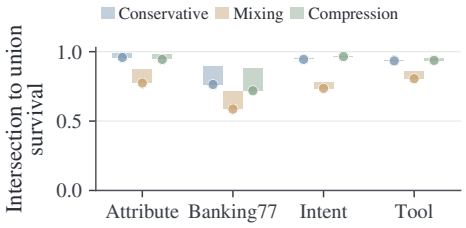

Figure 6: Intersection and union survival bounds by task and transmission pressure. Each translucent bar spans the two observable endpoints.

I Generation-Indexed Operators

For a nonstationary lineage, the exact killed survival is

$$s_T = \beta K_0 K_1 \cdots K_{T-1}\mathbf{1}. \tag{42}$$

This expression permits changing prompts, receiver states, or occupancy inside an executable fiber, provided the generation-indexed microstate remains conditionally closed.

Proposition I.1 (Stationary approximation error)**.** *Suppose K_t and K have operator norm at most one. Then*

$$\left\|K_0 \cdots K_{T-1} - K^T\right\| \le \sum_{s=0}^{T-1} \|K_s - K\|. \tag{43}$$

Proof. Insert one reference factor at a time to obtain

$$K_0 \cdots K_{T-1} - K^T = \sum_{s=0}^{T-1} K_0 \cdots K_{s-1} \cdot (K_s - K)K^{T-s-1}. \tag{44}$$

Submultiplicativity and the unit norm bounds give Equation 43. □

Proposition I.2 (Contracted perturbation)**.** *If suffix products contract signed mass by factors $\gamma_{s+1:T-1}$, the contribution of generation s is bounded by $\gamma_{s+1:T-1}\|K_s - K\|$.*

The proof repeats Equation 44 while retaining the suffix norm rather than replacing it by one. This is the killed-operator analogue of the contraction weighting in Theorem 4.2.

One-state time product. At the coarsest surviving versus exited resolution, define the empirical conditional retention $r_t = s_{t+1}/s_t$. Its generation-indexed product equals the design survival curve by Proposition D.1. We evaluate that product on held-out trajectories and compare it with both the stationary killed estimate and repeated one-step retention.

J Exact Finite-System Sweeps

The finite experiments instantiate every operator directly, without language-model sampling. They serve three roles: verifying algebraic identities, locating the range of each separation, and testing that the implementation respects the stated event convention.

Table 3: Exact Matched One-Step Constructions

e	T	Fast	Storage
0.10	1	0.900	0.900
0.10	8	0.430	0.900
0.10	32	0.034	0.900

Table 4: Exact observable-repair sweep.

Property	**Cells**
All repair cells	324
Domination satisfied	60
Negative gain without domination	91

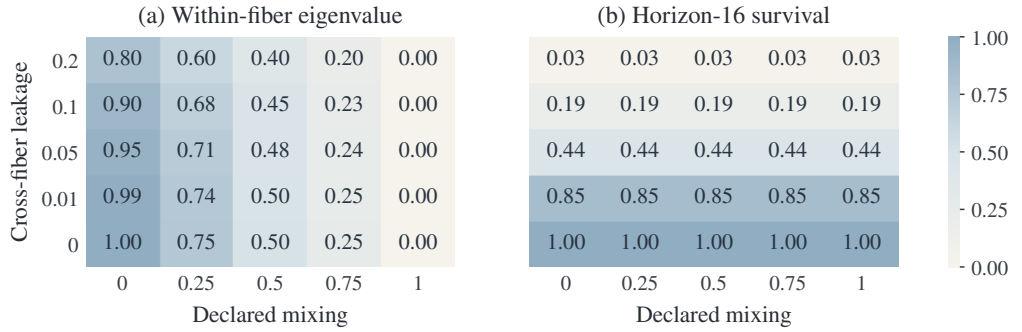

Figure 7: Independent control of within-fiber surface mixing and cross-fiber semantic leakage in the exact construction. The left panel changes with mixing; the right panel changes with leakage.

J.1 Matched one-step retention

For $e \in \{0.02, 0.05, 0.10, 0.20, 0.35\}$ and $T \in \{1, 2, 4, 8, 16, 32, 64\}$, we evaluate the two kernels in Equation 8. Both one-step values are equal by construction. The storage process preserves the initial surviving mass, while the fast process repeatedly exposes it to the same leakage. Table 3 reports representative cells, and the full 35-cell grid is included in the generated result tables.

J.2 Independent mixing and leakage

We cross twenty within-fiber mixing values with twenty leakage values. For each cell, the code verifies that conditional surface mixing depends only on its assigned coordinate and killed row mass depends only on leakage. This produces 400 exact cells. The sweep also reports the second within-fiber eigenvalue and the killed Perron value.

Fig. 7 displays the resulting factorization at fiber size eight and horizon 16.

J.3 Observable repair

The repair grid crosses detector true-trigger behavior, false triggers, repaired retention, and repaired destination rows. Sixty cells satisfy elementwise killed-operator domination. All sixty obey the finite-horizon ordering in Theorem 6.1. The remaining cells include both helpful and harmful interventions, confirming that an observable trigger alone does not determine the direction.

Fig. 8 shows that every domination cell lies on or above the equal-survival line, while the broader grid spans both sides.

Reading the repair grid. The domination label is a row-wise statement about the repaired and unrepaired killed operators; it is stronger than observing a positive endpoint difference. A cell can improve at horizon 16 while failing the elementwise order, because a temporary loss of mass may be followed by a favorable redistribution. We therefore retain both the operator-level label and the finite-horizon survival gain in the generated table.

The grid also separates trigger opportunity from repair quality. True-trigger and false-trigger rates determine how often the intervention is invoked, while repaired retention and destination rows determine what happens after invocation. The 324 cells cross these mechanisms independently. The

60 domination cells provide exact certificates, whereas the 91 negative-gain cells outside domination show why an observable trigger alone cannot guarantee improvement.

For practical comparisons, the safest order is to read the domination label first, then the repaired and unrepaired rows, and finally the horizon-specific gain. This order prevents a large gain in one horizon from being mistaken for a uniform policy guarantee. The language-model mediation results in Appendix R use the same distinction: immediate rescue is a local transition, while paired survival is the downstream policy effect.

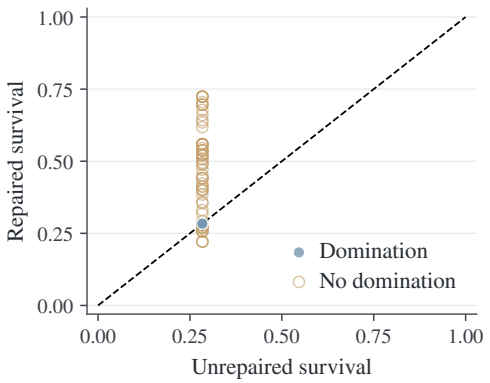

Figure 8: Exact repair sweep at horizon 16. Elementwise domination guarantees nonnegative survival gain.

The diagonal in Fig. 8 marks equal survival from the declared initial distribution at horizon 16. Points above it show a favorable policy effect at that horizon, but do not by themselves establish row-wise domination. Likewise, positive gains at several horizons remain weaker evidence than the operator inequality: they can reflect redistribution among surviving states. The domination label certifies the sufficient condition, while the plotted gain records its consequence for this particular initial distribution and horizon.

Table 4 summarizes the exact intervention grid before the boundary cells are expanded in Appendix U.

J.4 Approximate closure and time variation

Thirty approximately lumpable systems vary the aggregate-row defect and quotient contraction. Observed total-variation error remains below Equation 5 in every cell. Thirty-two time-varying systems vary alternating leakage and compare an ordered killed product with a stationary reference. Equation 43 holds in all cases. Complete cells appear in Appendix U.

K Executable Task Construction

Every episode presents exactly eight candidate actions. Candidate identifiers are randomly assigned after the target and distractors are selected, preventing a fixed identifier prior from defining the task. The same generated task instance is reused across relay, repair, and model conditions.

K.1 Compositional attribute reference

An object has color, shape, material, and position attributes. The inventory contains eight values per attribute. A target tuple and seven distinct distractors are sampled, then assigned identifiers. The initial public message names all four target attributes in ordinary English. This task makes semantic loss executable: a receiver crosses the fiber exactly when it selects a different object.

K.2 Manual intent routing

The ontology contains sixteen banking operations, including freezing and replacing a card, cash deposit and withdrawal, domestic and international transfer, beneficiary problems, refund and cashback status, PIN actions, disputes, failed and pending transfers, and account closure. Each operation has a short description and a manually written utterance. An episode presents the target plus seven sampled distractors.

K.3 Tool routing

The tool ontology contains calendar creation, movement and cancellation; email sending and search; file search and sharing; weather lookup; route planning; restaurant search and booking; calculation;

Table 5: Transmission conditions used in every language-model relay.

Condition	Instruction
Conservative	Rewrite the message clearly while preserving every operationally relevant detail. Prefer ordinary public English.
Mixing	Paraphrase the message with substantially different wording and syntax while preserving the exact intended action.
Compression	Rewrite the message as briefly as possible while preserving the exact intended action and all disambiguating details.

currency conversion; contact lookup; reminders; and translation. The task is action selection rather than free-form answer generation, making receiver behavior exactly observable.

K.4 BANKING77

Banking77 provides natural user utterances and 77 fine-grained banking intents (Casanueva et al., 2020). We sample a test utterance, pair its label with seven distractor labels, randomize identifier assignment, and retain the original dataset label for external task correctness. Semantic survival remains defined relative to each frozen receiver's original action, so an initially incorrect receiver still induces a well-defined executable fiber.

K.5 COMMON-INSTANCE SCHEDULE

The task seed depends only on the global run seed, task type, and episode index. It does not depend on relay mode, repair mode, or rewriter identity. This creates matched task instances across all conditions and supports paired repair contrasts. Candidate assignment and source wording are therefore fixed before any model output is observed. Comparisons across relay conditions change only the induced message trajectory, not the executable choice set presented to either receiver.

L TRANSMISSION, RECEIVER, AND REPAIR PROMPTS

The system and user messages below are the semantic content of the prompts. Endpoint-specific chat templates only serialize these declared messages.

L.1 TRANSMISSION INSTRUCTION

All rewriters receive the system instruction: "You are one link in a message transmission chain. Return only the transmitted message, with no analysis." The user message supplies the incoming message and one of the following condition instructions. Table 5 records the three instructions verbatim.

L.2 EXECUTABLE RECEIVER

The receiver system instruction is: "Execute the message by selecting one candidate. Return strict JSON `{"choice":"C0"}` and no explanation." The user message contains the full candidate list and current transmitted message. Parsing first attempts strict JSON, then accepts a unique valid candidate identifier found in the output. Multiple identifiers or no valid identifier are recorded as invalid.

L.3 OBSERVABLE REPAIR

The repair system instruction is: "You repair a transmitted message after two public interpreters disagreed. Return only a clear public message. Do not choose an action." The user message contains the stored public seed and current transmitted message, followed by a request to make the intended action unambiguous and faithful to the seed. The repairer never receives candidate labels, gold actions, or a semantic-survival indicator.

L.4 Decoding

Rewriters use temperature 0.4 and at most 96 output tokens. Receivers and repairers use deterministic decoding and at most 32 and 96 output tokens, respectively. Each request carries a deterministic episode and generation seed. Full model aliases and horizons are tabulated in Appendix T.

M Measurement Definitions

This appendix defines the trajectory quantities from which all tables and figures are computed.

M.1 Receiver-specific survival

Let $a_{r,0}$ be receiver r's action on the initial public message and $a_{r,t}$ its action at generation t. Receiver-specific survival through horizon T is

$$S_{r,T} = \prod_{t=0}^{T} \mathbf{1}\{a_{r,t} = a_{r,0}\}. \tag{45}$$

The intersection and union variables are

$$S_{\cap,T} = \prod_{t=0}^{T} \prod_{r} \mathbf{1}\{a_{r,t} = a_{r,0}\}, \tag{46}$$

$$S_{\cup,T} = \prod_{t=0}^{T} \max_{r} \mathbf{1}\{a_{r,t} = a_{r,0}\}. \tag{47}$$

M.2 First exit and survival area

First exit is the minimum generation with a zero receiver-specific or combined indicator. A path without exit receives the censored value $H + 1$ for a run of horizon H. Survival area is the mean of S_T over recorded post-initial generations. These two quantities retain more temporal information than final survival alone.

M.3 Surface statistics

Messages are normalized by lowercasing and collapsing whitespace. Exact turnover is one when the normalized message differs from the normalized seed. Normalized edit distance is one minus the SequenceMatcher ratio. We also record character count, whitespace-delimited word count, and type-token ratio. No surface quantity changes receiver labels or the basin definition.

M.4 Quotient microstates

The measured microstate records receiver agreement, exact equality to the seed message, and whether repair occurred at the current generation. Aggregate killed rows are estimated only from design trajectories that have survived through their source generation. Missing row mass is empirical semantic leakage.

M.5 Stationary and time-indexed predictions

The stationary prediction raises the design killed operator to the requested horizon. The generation-indexed prediction estimates each conditional retention on design trajectories and multiplies them in temporal order. Repeating only the first design retention gives the one-step baseline. All three predict held-out survival without refitting on held-out outcomes.

N Statistical Analysis

N.1 Design and held-out split

The split is a deterministic hash of task type and episode index. It is therefore common across model, transmission, and repair conditions. Operator estimation uses only design episodes; reported predictive error uses only held-out episodes. Because the split key is independent of generated outputs, changing the relay cannot reassign an instance.

N.2 Paired repair inference

For each task, transmission condition, rewriter, receiver event, and metric, we pair the repair and no-repair trajectory with the same episode index. Let d_i be the paired difference. We report $\bar{d}$ and a two-sided percentile interval from 10,000 resamples of the episode pairs. The resampling unit is a complete trajectory, preserving dependence across generations.

N.3 Exact one-step matching

For every condition, design one-step retention is represented by its integer number of surviving episodes. Two conditions match only when both the design sample size and integer survivor count are exactly equal within a task and receiver event. Later held-out survival is then compared across the matched conditions. This discrete rule introduces no distance tolerance.

N.4 Prediction errors

At horizons $1, 2, 4, 8, 16$, and 32 when available, we compute absolute deviation between predicted and held-out survival. Mean absolute error averages these predeclared horizons within a condition. Prediction comparisons retain the same held-out observations and therefore admit paired condition-level differences.

N.5 Repair mediation

At every triggered generation, the pre-repair and post-repair receiver actions produce four exhaustive cells: rescued exit, induced exit, preserved survival, and preserved exit. The net immediate rescue rate is rescued minus induced exits divided by the number of triggers. This measurement is descriptive of the local operator. The paired trajectory analysis measures the complete downstream policy effect.

N.6 No outcome-defined partition changes

The executable partition is fixed before transmission by the initial frozen receiver actions. We do not redefine equivalence from later surface similarity, gold correctness, or observed experimental separation. Initial and final gold accuracy remain external task checks.

This fixed partition is also what makes the survival comparisons auditable. Every trajectory carries the same candidate set and receiver action vector from its seed, so a later change can be attributed to the relay kernel or repair rule rather than to a moving definition of success. Surface similarity is reported as a diagnostic, while the event itself remains executable and receiver-defined.

The statistical split follows the same principle. Design episodes estimate the operator and held-out episodes evaluate its predictions, with the hash split applied before any condition-specific filtering. No endpoint outcome is used to redefine a basin, select a horizon, or remove a difficult task from the atlas.

O Complete Survival Curves

This section reports every predeclared horizon for the intersection and union receiver events. The curves retain first-passage semantics: once a trajectory exits, its survival indicator remains zero even if a later message returns to the initial action.

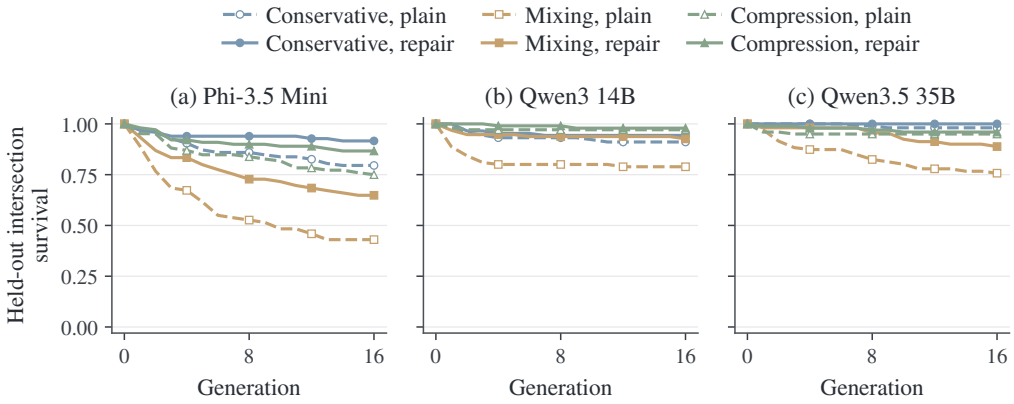

Figure 10: Held-out intersection survival across the main relay panel, averaged over four executable tasks. Solid curves use observable repair; dashed curves use unmodified transmission.

Fig. 10 gives the full 16-generation main panel, and Fig. 9 compares the 32-generation scale panels under strong paraphrase mixing. Table 6 summarizes the principal long-horizon cells before the following tables provide every task-level condition.

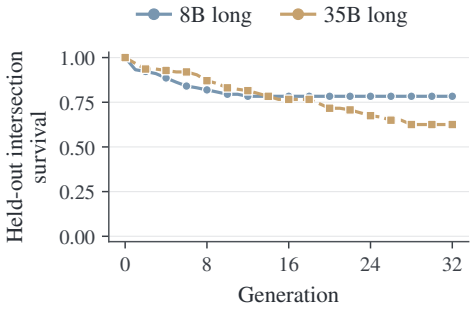

Figure 9: Held-out semantic survival through 32 relay generations, averaged across tasks under mixing without repair.

How to compare a survival curve. The first point is the probability of preserving the declared event after one relay step; later points retain only trajectories that have never crossed the event boundary. A steep initial drop therefore identifies immediate ambiguity, while a gradual decline indicates that initially safe messages become fragile under repeated composition. Comparing only the final point loses this timing information. The same horizon grid is used for plain and repaired transmission, and the task split is held fixed across conditions. Differences can therefore be read as changes in the transmission policy rather than changes in sampled tasks. For each row, the intersection event is a conservative joint guarantee across receivers; the union event is an outer envelope. Their separation records evaluator dependence along the path and should be reported alongside the curve. The long panels also provide a scale check. They extend the horizon while preserving the task definitions and relay pressures, so a persistent plateau suggests stored executable mass rather than a short-lived one-step coincidence. The tables below retain the task-level rows needed to distinguish that plateau from averaging across heterogeneous conditions. We preserve the full path for auditability: every row is generated from the same task-level trajectory records used in operator estimation, and the long-horizon appendix does not substitute a different endpoint sample. This makes the horizon-by-horizon contrast reproducible and allows later state refinements to be evaluated without rebuilding the task split.

The pilot appears in Tables 7, 8, and 9. The 8B horizon-16 panel appears in Tables 10, 11, 12, and 13. The main panel appears in Tables 14, 15, 16, and 17. The 8B long panel appears in Tables 18, 19, 20, and 21; the 35B long panel appears in Tables 22, 23, 24, and 25. Together they provide the complete condition-level survival atlas. They are organized by task so that model scale, relay pressure, and repair can be compared without changing the executable event. The intersection column requires both frozen receivers to preserve their respective initial actions at every generation. The union column requires at least one receiver to do so. Their width is therefore an observed measure of receiver ambiguity along the full path, not merely disagreement at the endpoint.

O.1 How to read the survival atlas

Each table combines semantic and surface measurements while keeping their definitions independent. Turnover is the fraction of distinct normalized strings along a trajectory, and edit distance compares

Table 6: Thirty-two-generation intersection survival under paraphrase mixing. Turnover is averaged over all messages in a trajectory.

Rewriter	Task	Plain	Repair	Plain turn.	Repair turn.	Effect
Qwen3 8B	Attribute	0.708	0.812	0.374	0.442	0.104
Qwen3 8B	Banking77	0.688	0.833	0.222	0.217	0.146
Qwen3 8B	Intent	0.771	0.938	0.277	0.280	0.167
Qwen3 8B	Tool	0.792	0.854	0.219	0.220	0.062
Qwen3.5 35B	Attribute	0.828	0.953	0.810	0.860	0.125
Qwen3.5 35B	Banking77	0.531	0.750	0.755	0.709	0.219
Qwen3.5 35B	Intent	0.609	0.875	0.635	0.661	0.266
Qwen3.5 35B	Tool	0.688	0.844	0.706	0.681	0.156

Table 7: Complete attribute survival panel for 2B and 4B pilot, 16 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Qwen3-VL 2B	Compression	Yes	0.833	1.000	0.167	0.167	0.336
Qwen3-VL 2B	Compression	No	0.875	1.000	0.125	0.162	0.344
Qwen3-VL 2B	Conservative	Yes	0.917	1.000	0.083	0.152	0.077
Qwen3-VL 2B	Conservative	No	0.917	1.000	0.083	0.157	0.083
Qwen3-VL 2B	Mixing	Yes	0.833	1.000	0.167	0.439	0.331
Qwen3-VL 2B	Mixing	No	0.875	1.000	0.125	0.456	0.335
Qwen3-VL 4B	Compression	Yes	0.875	1.000	0.125	0.167	0.417
Qwen3-VL 4B	Compression	No	0.875	1.000	0.125	0.172	0.433
Qwen3-VL 4B	Conservative	Yes	0.917	1.000	0.083	0.407	0.338
Qwen3-VL 4B	Conservative	No	0.875	1.000	0.125	0.426	0.364
Qwen3-VL 4B	Mixing	Yes	0.792	0.917	0.125	0.699	0.484
Qwen3-VL 4B	Mixing	No	0.667	0.875	0.208	0.662	0.478

Table 8: Complete intent survival panel for 2B and 4B pilot, 16 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Qwen3-VL 2B	Compression	Yes	1.000	1.000	0.000	0.179	0.324
Qwen3-VL 2B	Compression	No	1.000	1.000	0.000	0.181	0.312
Qwen3-VL 2B	Conservative	Yes	1.000	1.000	0.000	0.167	0.265
Qwen3-VL 2B	Conservative	No	1.000	1.000	0.000	0.167	0.264
Qwen3-VL 2B	Mixing	Yes	0.875	0.917	0.042	0.520	0.543
Qwen3-VL 2B	Mixing	No	0.875	0.917	0.042	0.534	0.550
Qwen3-VL 4B	Compression	Yes	1.000	1.000	0.000	0.194	0.497
Qwen3-VL 4B	Compression	No	0.917	1.000	0.083	0.189	0.507
Qwen3-VL 4B	Conservative	Yes	1.000	1.000	0.000	0.238	0.363
Qwen3-VL 4B	Conservative	No	1.000	1.000	0.000	0.238	0.357
Qwen3-VL 4B	Mixing	Yes	1.000	1.000	0.000	0.571	0.575
Qwen3-VL 4B	Mixing	No	0.833	0.875	0.042	0.613	0.576

the final message with the seed. Neither statistic enters the receiver event. A row can consequently have large turnover, large edit distance, and high survival. Such rows are the empirical form of Proposition 5.3. Conversely, a small amount of surface motion can cross an executable boundary when the changed phrase is disambiguating.

The three relay instructions alter different parts of the transition kernel. Conservative rewriting tends to keep occupancy near familiar forms. Mixing distributes mass across a wider collection of paraphrases. Compression removes tokens and can either preserve a compact action description or delete a distinguishing attribute. The survival columns show the resulting cross-fiber leakage after repeated composition. They should not be interpreted as rankings of writing quality, since each condition deliberately imposes a different transmission pressure.

Repair changes the transition law only after an observable disagreement or invalid output. As a result, repair and plain rows are often identical when both receivers agree throughout. Larger differences

Table 9: Complete tool survival panel for 2B and 4B pilot, 16 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Qwen3-VL 2B	Compression	Yes	0.958	1.000	0.042	0.154	0.175
Qwen3-VL 2B	Compression	No	0.958	1.000	0.042	0.157	0.175
Qwen3-VL 2B	Conservative	Yes	1.000	1.000	0.000	0.176	0.169
Qwen3-VL 2B	Conservative	No	1.000	1.000	0.000	0.179	0.157
Qwen3-VL 2B	Mixing	Yes	0.917	0.917	0.000	0.478	0.352
Qwen3-VL 2B	Mixing	No	0.875	0.958	0.083	0.495	0.353
Qwen3-VL 4B	Compression	Yes	1.000	1.000	0.000	0.181	0.320
Qwen3-VL 4B	Compression	No	1.000	1.000	0.000	0.176	0.308
Qwen3-VL 4B	Conservative	Yes	1.000	1.000	0.000	0.196	0.232
Qwen3-VL 4B	Conservative	No	1.000	1.000	0.000	0.189	0.233
Qwen3-VL 4B	Mixing	Yes	0.958	0.958	0.000	0.331	0.371
Qwen3-VL 4B	Mixing	No	0.875	0.958	0.083	0.341	0.392

Table 10: Complete attribute survival panel for 8B relay, 16 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Qwen3 8B	Compression	Yes	0.969	1.000	0.031	0.132	0.347
Qwen3 8B	Compression	No	1.000	1.000	0.000	0.131	0.355
Qwen3 8B	Conservative	Yes	1.000	1.000	0.000	0.169	0.140
Qwen3 8B	Conservative	No	0.938	1.000	0.062	0.180	0.138
Qwen3 8B	Mixing	Yes	0.812	0.938	0.125	0.722	0.385
Qwen3 8B	Mixing	No	0.750	0.938	0.188	0.653	0.403

Table 11: Complete banking77 survival panel for 8B relay, 16 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Qwen3 8B	Compression	Yes	0.969	1.000	0.031	0.143	0.160
Qwen3 8B	Compression	No	0.969	0.969	0.000	0.147	0.168
Qwen3 8B	Conservative	Yes	1.000	1.000	0.000	0.178	0.159
Qwen3 8B	Conservative	No	1.000	1.000	0.000	0.169	0.158
Qwen3 8B	Mixing	Yes	0.844	0.938	0.094	0.371	0.359
Qwen3 8B	Mixing	No	0.719	0.938	0.219	0.364	0.401

Table 12: Complete intent survival panel for 8B relay, 16 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Qwen3 8B	Compression	Yes	1.000	1.000	0.000	0.132	0.202
Qwen3 8B	Compression	No	1.000	1.000	0.000	0.131	0.202
Qwen3 8B	Conservative	Yes	1.000	1.000	0.000	0.136	0.156
Qwen3 8B	Conservative	No	1.000	1.000	0.000	0.136	0.156
Qwen3 8B	Mixing	Yes	0.969	0.969	0.000	0.434	0.476
Qwen3 8B	Mixing	No	0.781	0.938	0.156	0.443	0.477

occur where the detector encounters proposed messages near a receiver boundary. The 35B horizon-32 mixing rows give the clearest instance: repair increases intersection survival in all four tasks, with the largest change on intent routing. The 8B long panel reproduces the same direction on three tasks and retains a smaller positive tool effect.

O.2 Horizon reveals stored fragility

Endpoint survival at generation 16 or 32 summarizes the entire first-passage path. The underlying curves reveal when the loss occurs. Some conditions fall early and then flatten, indicating that a subset of task instances is immediately ambiguous while the remaining instances occupy persistent

Table 13: Complete tool survival panel for 8B relay, 16 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Qwen3 8B	Compression	Yes	0.844	0.906	0.062	0.153	0.193
Qwen3 8B	Compression	No	0.844	0.875	0.031	0.147	0.191
Qwen3 8B	Conservative	Yes	0.938	1.000	0.062	0.199	0.186
Qwen3 8B	Conservative	No	0.938	0.969	0.031	0.200	0.184
Qwen3 8B	Mixing	Yes	0.750	0.812	0.062	0.371	0.337
Qwen3 8B	Mixing	No	0.688	0.906	0.219	0.379	0.392

Table 14: Complete attribute survival panel for 14B, 35B, and Phi main panel, 16 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Phi-3.5 Mini	Compression	Yes	0.938	0.958	0.021	0.418	0.512
Phi-3.5 Mini	Compression	No	0.812	0.917	0.104	0.434	0.525
Phi-3.5 Mini	Conservative	Yes	0.958	0.979	0.021	0.837	0.476
Phi-3.5 Mini	Conservative	No	0.833	0.958	0.125	0.870	0.489
Phi-3.5 Mini	Mixing	Yes	0.771	0.792	0.021	0.960	0.597
Phi-3.5 Mini	Mixing	No	0.417	0.625	0.208	0.993	0.653
Qwen3 14B	Compression	Yes	0.979	1.000	0.021	0.170	0.378
Qwen3 14B	Compression	No	1.000	1.000	0.000	0.168	0.381
Qwen3 14B	Conservative	Yes	1.000	1.000	0.000	0.188	0.223
Qwen3 14B	Conservative	No	0.958	1.000	0.042	0.179	0.222
Qwen3 14B	Mixing	Yes	0.917	0.938	0.021	0.627	0.394
Qwen3 14B	Mixing	No	0.771	0.979	0.208	0.653	0.392
Qwen3.5 35B	Compression	Yes	0.979	1.000	0.021	0.179	0.469
Qwen3.5 35B	Compression	No	0.958	1.000	0.042	0.172	0.468
Qwen3.5 35B	Conservative	Yes	1.000	1.000	0.000	0.178	0.191
Qwen3.5 35B	Conservative	No	1.000	1.000	0.000	0.174	0.198
Qwen3.5 35B	Mixing	Yes	0.958	0.958	0.000	0.919	0.481
Qwen3.5 35B	Mixing	No	0.812	0.938	0.125	0.907	0.492

Table 15: Complete banking77 survival panel for 14B, 35B, and Phi main panel, 16 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Phi-3.5 Mini	Compression	Yes	0.667	0.771	0.104	0.505	0.573
Phi-3.5 Mini	Compression	No	0.500	0.771	0.271	0.516	0.658
Phi-3.5 Mini	Conservative	Yes	0.688	0.792	0.104	0.755	0.556
Phi-3.5 Mini	Conservative	No	0.542	0.729	0.188	0.806	0.576
Phi-3.5 Mini	Mixing	Yes	0.542	0.646	0.104	0.798	0.639
Phi-3.5 Mini	Mixing	No	0.333	0.438	0.104	0.922	0.729
Qwen3 14B	Compression	Yes	0.833	0.958	0.125	0.143	0.253
Qwen3 14B	Compression	No	0.792	0.938	0.146	0.137	0.255
Qwen3 14B	Conservative	Yes	0.854	0.938	0.083	0.162	0.196
Qwen3 14B	Conservative	No	0.812	0.979	0.167	0.158	0.194
Qwen3 14B	Mixing	Yes	0.750	0.854	0.104	0.456	0.439
Qwen3 14B	Mixing	No	0.646	0.812	0.167	0.526	0.492
Qwen3.5 35B	Compression	Yes	0.792	0.917	0.125	0.191	0.445
Qwen3.5 35B	Compression	No	0.729	0.938	0.208	0.185	0.468
Qwen3.5 35B	Conservative	Yes	0.854	0.958	0.104	0.145	0.218
Qwen3.5 35B	Conservative	No	0.833	0.979	0.146	0.138	0.216
Qwen3.5 35B	Mixing	Yes	0.688	0.792	0.104	0.803	0.550
Qwen3.5 35B	Mixing	No	0.562	0.729	0.167	0.863	0.584

forms. Other conditions decline gradually as repeated rewriting moves mass toward increasingly fragile messages. A one-step score cannot distinguish these shapes, even when its estimate is exact.

The longer panels isolate this issue without changing task definitions. Qwen3 8B and Qwen3.5 35B each run for 32 generations on independent task seeds. The 35B model maintains high attribute survival under mixing, while intent and Banking77 accumulate more leakage. This task ordering reflects the geometry of candidate distinctions: a complete four-attribute description has several

Table 16: Complete intent survival panel for 14B, 35B, and Phi main panel, 16 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Phi-3.5 Mini	Compression	Yes	0.917	0.917	0.000	0.439	0.658
Phi-3.5 Mini	Compression	No	0.896	0.917	0.021	0.431	0.661
Phi-3.5 Mini	Conservative	Yes	0.938	0.938	0.000	0.642	0.584
Phi-3.5 Mini	Conservative	No	0.729	0.771	0.042	0.652	0.629
Phi-3.5 Mini	Mixing	Yes	0.521	0.542	0.021	0.903	0.710
Phi-3.5 Mini	Mixing	No	0.312	0.417	0.104	0.930	0.798
Qwen3 14B	Compression	Yes	1.000	1.000	0.000	0.121	0.144
Qwen3 14B	Compression	No	0.979	1.000	0.021	0.121	0.139
Qwen3 14B	Conservative	Yes	1.000	1.000	0.000	0.159	0.217
Qwen3 14B	Conservative	No	1.000	1.000	0.000	0.156	0.213
Qwen3 14B	Mixing	Yes	0.979	0.979	0.000	0.511	0.500
Qwen3 14B	Mixing	No	0.917	0.958	0.042	0.512	0.517
Qwen3.5 35B	Compression	Yes	1.000	1.000	0.000	0.158	0.418
Qwen3.5 35B	Compression	No	1.000	1.000	0.000	0.157	0.415
Qwen3.5 35B	Conservative	Yes	1.000	1.000	0.000	0.131	0.137
Qwen3.5 35B	Conservative	No	1.000	1.000	0.000	0.134	0.140
Qwen3.5 35B	Mixing	Yes	0.875	0.896	0.021	0.717	0.548
Qwen3.5 35B	Mixing	No	0.812	0.896	0.083	0.708	0.562

Table 17: Complete tool survival panel for 14B, 35B, and Phi main panel, 16 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Phi-3.5 Mini	Compression	Yes	0.875	0.875	0.000	0.412	0.457
Phi-3.5 Mini	Compression	No	0.812	0.854	0.042	0.392	0.474
Phi-3.5 Mini	Conservative	Yes	0.958	0.958	0.000	0.665	0.528
Phi-3.5 Mini	Conservative	No	0.896	0.917	0.021	0.661	0.516
Phi-3.5 Mini	Mixing	Yes	0.792	0.792	0.000	0.907	0.669
Phi-3.5 Mini	Mixing	No	0.542	0.771	0.229	0.904	0.727
Qwen3 14B	Compression	Yes	1.000	1.000	0.000	0.126	0.156
Qwen3 14B	Compression	No	1.000	1.000	0.000	0.126	0.161
Qwen3 14B	Conservative	Yes	0.875	0.875	0.000	0.145	0.161
Qwen3 14B	Conservative	No	0.875	0.875	0.000	0.142	0.161
Qwen3 14B	Mixing	Yes	0.938	0.938	0.000	0.412	0.388
Qwen3 14B	Mixing	No	0.833	0.854	0.021	0.509	0.424
Qwen3.5 35B	Compression	Yes	0.979	1.000	0.021	0.165	0.289
Qwen3.5 35B	Compression	No	0.958	1.000	0.042	0.165	0.285
Qwen3.5 35B	Conservative	Yes	1.000	1.000	0.000	0.135	0.079
Qwen3.5 35B	Conservative	No	1.000	1.000	0.000	0.136	0.079
Qwen3.5 35B	Mixing	Yes	0.917	0.917	0.000	0.714	0.487
Qwen3.5 35B	Mixing	No	0.812	0.875	0.062	0.722	0.509

Table 18: Complete attribute survival panel for 8B long-horizon relay, 32 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Qwen3 8B	Compression	Yes	0.938	0.979	0.042	0.073	0.356
Qwen3 8B	Compression	No	0.958	0.979	0.021	0.071	0.351
Qwen3 8B	Conservative	Yes	0.938	0.979	0.042	0.097	0.145
Qwen3 8B	Conservative	No	0.917	0.979	0.062	0.098	0.145
Qwen3 8B	Mixing	Yes	0.812	0.896	0.083	0.442	0.412
Qwen3 8B	Mixing	No	0.708	0.833	0.125	0.374	0.416

redundant lexical realizations, whereas neighboring intent labels can differ by a single operational relation such as pending versus failed.

Across the long panels, early flattening concentrates leakage in a small subset of instances, whereas gradual decline records repeated exposure among initially safe messages. This distinction is visible only when the saved horizons are retained as a curve. Two conditions can finish at the same survival value while exposing different periods of reliable execution, so the curve retains temporal information

Table 19: Complete banking77 survival panel for 8B long-horizon relay, 32 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Qwen3 8B	Compression	Yes	0.979	1.000	0.021	0.073	0.160
Qwen3 8B	Compression	No	0.958	1.000	0.042	0.073	0.161
Qwen3 8B	Conservative	Yes	0.979	1.000	0.021	0.088	0.122
Qwen3 8B	Conservative	No	0.979	1.000	0.021	0.088	0.125
Qwen3 8B	Mixing	Yes	0.833	0.854	0.021	0.217	0.367
Qwen3 8B	Mixing	No	0.688	0.792	0.104	0.222	0.395

Table 20: Complete intent survival panel for 8B long-horizon relay, 32 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Qwen3 8B	Compression	Yes	1.000	1.000	0.000	0.066	0.190
Qwen3 8B	Compression	No	1.000	1.000	0.000	0.066	0.190
Qwen3 8B	Conservative	Yes	1.000	1.000	0.000	0.083	0.217
Qwen3 8B	Conservative	No	1.000	1.000	0.000	0.081	0.218
Qwen3 8B	Mixing	Yes	0.938	0.938	0.000	0.280	0.522
Qwen3 8B	Mixing	No	0.771	0.896	0.125	0.277	0.548

Table 21: Complete tool survival panel for 8B long-horizon relay, 32 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Qwen3 8B	Compression	Yes	0.938	0.958	0.021	0.076	0.182
Qwen3 8B	Compression	No	0.938	0.958	0.021	0.076	0.180
Qwen3 8B	Conservative	Yes	0.979	1.000	0.021	0.093	0.150
Qwen3 8B	Conservative	No	0.979	1.000	0.021	0.093	0.150
Qwen3 8B	Mixing	Yes	0.854	0.875	0.021	0.220	0.403
Qwen3 8B	Mixing	No	0.792	0.812	0.021	0.219	0.403

Table 22: Complete attribute survival panel for 35B long-horizon relay, 32 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Qwen3.5 35B	Compression	Yes	0.984	1.000	0.016	0.090	0.461
Qwen3.5 35B	Compression	No	0.984	1.000	0.016	0.091	0.477
Qwen3.5 35B	Conservative	Yes	0.984	1.000	0.016	0.107	0.202
Qwen3.5 35B	Conservative	No	0.984	1.000	0.016	0.106	0.201
Qwen3.5 35B	Mixing	Yes	0.953	0.984	0.031	0.860	0.473
Qwen3.5 35B	Mixing	No	0.828	0.953	0.125	0.810	0.497

Table 23: Complete banking77 survival panel for 35B long-horizon relay, 32 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Qwen3.5 35B	Compression	Yes	0.906	0.969	0.062	0.098	0.414
Qwen3.5 35B	Compression	No	0.766	0.938	0.172	0.093	0.437
Qwen3.5 35B	Conservative	Yes	0.922	0.984	0.062	0.078	0.157
Qwen3.5 35B	Conservative	No	0.938	0.984	0.047	0.073	0.153
Qwen3.5 35B	Mixing	Yes	0.750	0.844	0.094	0.709	0.573
Qwen3.5 35B	Mixing	No	0.531	0.719	0.188	0.755	0.580

that a terminal accuracy column necessarily removes. The full horizon grid preserves this distinction for every condition.

Table 24: Complete intent survival panel for 35B long-horizon relay, 32 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Qwen3.5 35B	Compression	Yes	0.969	1.000	0.031	0.076	0.338
Qwen3.5 35B	Compression	No	0.953	1.000	0.047	0.077	0.346
Qwen3.5 35B	Conservative	Yes	0.969	1.000	0.031	0.072	0.133
Qwen3.5 35B	Conservative	No	1.000	1.000	0.000	0.071	0.116
Qwen3.5 35B	Mixing	Yes	0.875	0.906	0.031	0.661	0.547
Qwen3.5 35B	Mixing	No	0.609	0.719	0.109	0.635	0.587

Table 25: Complete tool survival panel for 35B long-horizon relay, 32 generations. Intersection and union are the dual-receiver bounds.

Rewriter	Relay	Repair	$S_\cap$	$S_\cup$	Width	Turnover	Edit
Qwen3.5 35B	Compression	Yes	0.938	0.953	0.016	0.081	0.255
Qwen3.5 35B	Compression	No	0.891	0.938	0.047	0.080	0.253
Qwen3.5 35B	Conservative	Yes	0.984	1.000	0.016	0.071	0.096
Qwen3.5 35B	Conservative	No	0.969	0.984	0.016	0.074	0.105
Qwen3.5 35B	Mixing	Yes	0.844	0.891	0.047	0.681	0.477
Qwen3.5 35B	Mixing	No	0.688	0.797	0.109	0.706	0.511

O.3 First passage and return answer different questions

The reported survival event is absorbing. If a receiver changes action at generation five and returns to its original action at generation nine, endpoint agreement at generation nine is one while survival through generation nine is zero. This convention is appropriate for a relay whose intermediate message may already be executed. Return remains scientifically useful because it distinguishes an absorbing communication channel from a recoverable semantic excursion. Appendix T reports return after exit alongside the first observed exit generation.

Receiver intersection and union produce a second distinction. A narrow interval indicates that both frozen action maps induce nearly the same path event. A wider interval identifies task instances where the operational meaning is evaluator-dependent. The interval is kept at every horizon, so uncertainty is propagated rather than resolved by selecting one receiver after seeing the results. Because both events are computed from the same trajectories, their separation can be compared without changing the task instances, relay histories, or reporting horizon.

P Killed-Operator Estimates

The design split estimates killed operators on observable quotient microstates. We report the number of observed states, Perron survival value, maximum and occupancy-weighted closure defects, and held-out prediction errors for stationary, generation-indexed, and repeated one-step estimators. All three errors use the same held-out instances and horizon grid.

Fig. 11 relates measured closure defect to the advantage of a generation-indexed product over a stationary power.

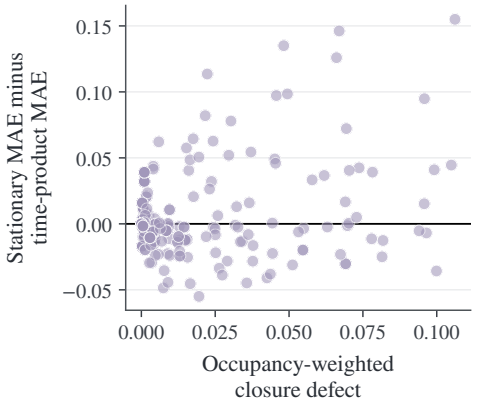

Figure 11: Stationary prediction attenuation and measured closure defect in the main panel. Positive values favor the time-indexed killed product.

The pilot operators appear in Tables 26, 27, and 28. The 8B horizon-16 operators appear in Tables 29, 30, 31, and 32. The main operators appear in Tables 33, 34, 35, and 36. The 8B long operators appear in Tables 37, 38, 39, and 40; the 35B long operators appear in Tables 41, 42, 43, and 44. These tables report every intersection-event

Table 26: Killed-operator estimates for attribute, 2B and 4B pilot, 16 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Qwen3-VL 2B	Compression	Yes	1.000	0.048	0.539	0.140	0.140
Qwen3-VL 2B	Compression	No	1.000	0.049	0.428	0.097	0.063
Qwen3-VL 2B	Conservative	Yes	1.000	0.010	0.253	0.014	0.014
Qwen3-VL 2B	Conservative	No	1.000	0.009	0.091	0.058	0.029
Qwen3-VL 2B	Mixing	Yes	1.000	0.091	0.091	0.103	0.080
Qwen3-VL 2B	Mixing	No	1.000	0.020	0.091	0.064	0.069
Qwen3-VL 4B	Compression	Yes	1.000	0.022	0.428	0.063	0.063
Qwen3-VL 4B	Compression	No	1.000	0.026	0.253	0.034	0.043
Qwen3-VL 4B	Conservative	Yes	1.000	0.032	0.253	0.031	0.014
Qwen3-VL 4B	Conservative	No	1.000	0.072	0.253	0.078	0.043
Qwen3-VL 4B	Mixing	Yes	1.000	0.097	0.539	0.187	0.171
Qwen3-VL 4B	Mixing	No	1.000	0.111	0.412	0.095	0.081

Table 27: Killed-operator estimates for intent, 2B and 4B pilot, 16 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Qwen3-VL 2B	Compression	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3-VL 2B	Compression	No	1.000	0.000	0.000	0.000	0.000
Qwen3-VL 2B	Conservative	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3-VL 2B	Conservative	No	1.000	0.000	0.000	0.000	0.000
Qwen3-VL 2B	Mixing	Yes	0.989	0.001	0.020	0.034	0.094
Qwen3-VL 2B	Mixing	No	0.989	0.001	0.080	0.047	0.034
Qwen3-VL 4B	Compression	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3-VL 4B	Compression	No	0.995	0.001	0.080	0.054	0.023
Qwen3-VL 4B	Conservative	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3-VL 4B	Conservative	No	1.000	0.000	0.000	0.000	0.000
Qwen3-VL 4B	Mixing	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3-VL 4B	Mixing	No	0.968	0.005	0.301	0.184	0.151

Table 28: Killed-operator estimates for tool, 2B and 4B pilot, 16 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Qwen3-VL 2B	Compression	Yes	1.000	0.000	0.067	0.067	0.067
Qwen3-VL 2B	Compression	No	1.000	0.000	0.067	0.067	0.067
Qwen3-VL 2B	Conservative	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3-VL 2B	Conservative	No	1.000	0.000	0.000	0.000	0.000
Qwen3-VL 2B	Mixing	Yes	1.000	0.000	0.100	0.100	0.100
Qwen3-VL 2B	Mixing	No	0.971	0.007	0.241	0.086	0.033
Qwen3-VL 4B	Compression	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3-VL 4B	Compression	No	1.000	0.000	0.000	0.000	0.000
Qwen3-VL 4B	Conservative	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3-VL 4B	Conservative	No	1.000	0.000	0.000	0.000	0.000
Qwen3-VL 4B	Mixing	Yes	1.000	0.000	0.067	0.067	0.067
Qwen3-VL 4B	Mixing	No	0.975	0.003	0.067	0.076	0.050

operator group. The same design split estimates all three predictors. Repeated one-step retention uses only the first killed row mass. The stationary predictor retains movement among observable surviving states. The time product additionally permits the killed law to change with generation. This nesting makes their error differences interpretable as information added to the state description. Under the common held-out target, the successive error reductions isolate the value of destination allocation and then generation dependence, rather than changes in examples or evaluation.

P.1 Operator estimation from surviving paths

For each source microstate, transitions are counted only while the trajectory has survived through that source generation. A destination transition contributes to the killed matrix when the destination

Table 29: Killed-operator estimates for attribute, 8B relay, 16 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Qwen3 8B	Compression	Yes	1.000	0.000	0.059	0.059	0.059
Qwen3 8B	Compression	No	1.000	0.000	0.000	0.000	0.000
Qwen3 8B	Conservative	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3 8B	Conservative	No	1.000	0.000	0.106	0.106	0.106
Qwen3 8B	Mixing	Yes	0.989	0.089	0.235	0.130	0.089
Qwen3 8B	Mixing	No	0.944	0.028	0.200	0.196	0.161

Table 30: Killed-operator estimates for banking77, 8B relay, 16 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Qwen3 8B	Compression	Yes	1.000	0.001	0.000	0.023	0.013
Qwen3 8B	Compression	No	1.000	0.001	0.000	0.023	0.013
Qwen3 8B	Conservative	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3 8B	Conservative	No	1.000	0.000	0.000	0.000	0.000
Qwen3 8B	Mixing	Yes	0.973	0.078	0.433	0.179	0.113
Qwen3 8B	Mixing	No	0.947	0.022	0.303	0.143	0.045

Table 31: Killed-operator estimates for intent, 8B relay, 16 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Qwen3 8B	Compression	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3 8B	Compression	No	1.000	0.000	0.000	0.000	0.000
Qwen3 8B	Conservative	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3 8B	Conservative	No	1.000	0.000	0.000	0.000	0.000
Qwen3 8B	Mixing	Yes	0.996	0.001	0.000	0.021	0.024
Qwen3 8B	Mixing	No	0.969	0.004	0.107	0.051	0.015

Table 32: Killed-operator estimates for tool, 8B relay, 16 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Qwen3 8B	Compression	Yes	0.991	0.046	0.210	0.078	0.110
Qwen3 8B	Compression	No	0.991	0.030	0.210	0.077	0.098
Qwen3 8B	Conservative	Yes	1.000	0.015	0.431	0.111	0.111
Qwen3 8B	Conservative	No	1.000	0.023	0.431	0.111	0.111
Qwen3 8B	Mixing	Yes	0.991	0.029	0.245	0.070	0.073
Qwen3 8B	Mixing	No	0.945	0.053	0.326	0.137	0.041

also remains inside the basin; otherwise its probability is missing row mass. The initial vector is the design distribution over observable microstates at generation zero. Raising this estimated matrix to a horizon produces a survival prediction before held-out outcomes are read.

Sparse state occupancy can produce a Perron value equal to one when a rarely observed state has no design exit. Finite-horizon prediction remains well-defined because it weights that state by the initial vector and the transitions that reach it. We therefore report both $\rho(K)$ and held-out error, with the latter serving as the direct assessment of the finite-budget claim. The exact constructions establish the asymptotic interpretation independently of empirical sparsity.

P.2 Destination structure is the principal gain

Across all 720 operator groups, the scalar one-step predictor has mean absolute error 0.129. Retaining destination structure in a stationary killed operator lowers the error to 0.057. The paired reduction is 0.072 with interval $[0.064, 0.080]$. This gain is substantially larger than the additional reduction from

Table 33: Killed-operator estimates for attribute, 14B, 35B, and Phi main panel, 16 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Phi-3.5 Mini	Compression	Yes	0.997	0.004	0.190	0.043	0.052
Phi-3.5 Mini	Compression	No	0.988	0.004	0.126	0.029	0.058
Phi-3.5 Mini	Conservative	Yes	0.997	0.000	0.009	0.007	0.009
Phi-3.5 Mini	Conservative	No	0.981	0.002	0.082	0.027	0.020
Phi-3.5 Mini	Mixing	Yes	0.981	0.004	0.060	0.040	0.053
Phi-3.5 Mini	Mixing	No	0.900	0.022	0.285	0.144	0.062
Qwen3 14B	Compression	Yes	1.000	0.000	0.009	0.009	0.009
Qwen3 14B	Compression	No	1.000	0.000	0.000	0.000	0.000
Qwen3 14B	Conservative	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3 14B	Conservative	No	1.000	0.000	0.036	0.036	0.036
Qwen3 14B	Mixing	Yes	0.995	0.005	0.045	0.019	0.027
Qwen3 14B	Mixing	No	0.979	0.007	0.149	0.082	0.130
Qwen3.5 35B	Compression	Yes	1.000	0.000	0.045	0.045	0.045
Qwen3.5 35B	Compression	No	1.000	0.005	0.156	0.007	0.007
Qwen3.5 35B	Conservative	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3.5 35B	Conservative	No	1.000	0.000	0.000	0.000	0.000
Qwen3.5 35B	Mixing	Yes	0.995	0.001	0.000	0.026	0.015
Qwen3.5 35B	Mixing	No	0.992	0.001	0.118	0.078	0.087

Table 34: Killed-operator estimates for banking77, 14B, 35B, and Phi main panel, 16 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Phi-3.5 Mini	Compression	Yes	0.966	0.082	0.495	0.257	0.269
Phi-3.5 Mini	Compression	No	0.931	0.038	0.357	0.175	0.203
Phi-3.5 Mini	Conservative	Yes	0.969	0.105	0.564	0.282	0.237
Phi-3.5 Mini	Conservative	No	0.920	0.070	0.525	0.348	0.308
Phi-3.5 Mini	Mixing	Yes	0.883	0.106	0.418	0.329	0.174
Phi-3.5 Mini	Mixing	No	0.931	0.100	0.438	0.218	0.254
Qwen3 14B	Compression	Yes	1.000	0.097	0.575	0.252	0.259
Qwen3 14B	Compression	No	1.000	0.044	0.443	0.163	0.201
Qwen3 14B	Conservative	Yes	1.000	0.067	0.518	0.149	0.172
Qwen3 14B	Conservative	No	0.978	0.026	0.498	0.185	0.179
Qwen3 14B	Mixing	Yes	0.960	0.046	0.503	0.247	0.150
Qwen3 14B	Mixing	No	0.938	0.096	0.484	0.303	0.208
Qwen3.5 35B	Compression	Yes	0.995	0.069	0.475	0.196	0.227
Qwen3.5 35B	Compression	No	0.995	0.078	0.543	0.214	0.225
Qwen3.5 35B	Conservative	Yes	0.987	0.034	0.502	0.223	0.236
Qwen3.5 35B	Conservative	No	0.993	0.038	0.314	0.160	0.177
Qwen3.5 35B	Mixing	Yes	0.958	0.094	0.393	0.239	0.244
Qwen3.5 35B	Mixing	No	0.924	0.096	0.263	0.201	0.186

temporal indexing, which lowers error to 0.050. The pattern agrees with the theoretical hierarchy: hidden allocation of surviving mass is the main source of non-identification, and nonstationarity is a secondary source once that allocation is represented.

The scale panels show that the result is not driven by a single receiver pair. In the 2B and 4B pilot, stationary and time-indexed errors are 0.037 and 0.035, compared with 0.088 for one-step repetition. The 8B horizon-16 panel gives 0.046, 0.035, and 0.098. The 14B, 35B, and Phi main panel gives 0.068, 0.063, and 0.130. At horizon 32, the contrast widens: the 8B panel gives 0.078, 0.059, and 0.232, while the 35B panel gives 0.043, 0.035, and 0.112.

The aggregate gives every operator group equal weight. Larger panels increase the number of tested conditions without allowing one task, model, or horizon to dominate the reported prediction error.

P.3 Closure defect and temporal drift

The maximum defect is sensitive to rare microstates, so the tables also report an occupancy-weighted defect. A zero weighted defect means all observed surviving states have the same aggregate retention row under the measured representation. Positive values indicate that agreement, seed equality,

Table 35: Killed-operator estimates for intent, 14B, 35B, and Phi main panel, 16 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Phi-3.5 Mini	Compression	Yes	0.995	0.001	0.029	0.011	0.006
Phi-3.5 Mini	Compression	No	0.987	0.002	0.029	0.036	0.031
Phi-3.5 Mini	Conservative	Yes	0.995	0.001	0.010	0.017	0.020
Phi-3.5 Mini	Conservative	No	0.956	0.006	0.067	0.118	0.056
Phi-3.5 Mini	Mixing	Yes	0.947	0.014	0.080	0.056	0.058
Phi-3.5 Mini	Mixing	No	0.833	0.018	0.134	0.135	0.071
Qwen3 14B	Compression	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3 14B	Compression	No	1.000	0.004	0.103	0.031	0.037
Qwen3 14B	Conservative	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3 14B	Conservative	No	1.000	0.000	0.000	0.000	0.000
Qwen3 14B	Mixing	Yes	0.998	0.000	0.000	0.013	0.030
Qwen3 14B	Mixing	No	0.995	0.001	0.067	0.041	0.030
Qwen3.5 35B	Compression	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3.5 35B	Compression	No	1.000	0.000	0.000	0.000	0.000
Qwen3.5 35B	Conservative	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3.5 35B	Conservative	No	1.000	0.000	0.000	0.000	0.000
Qwen3.5 35B	Mixing	Yes	0.990	0.002	0.029	0.022	0.014
Qwen3.5 35B	Mixing	No	0.967	0.004	0.095	0.078	0.036

Table 36: Killed-operator estimates for tool, 14B, 35B, and Phi main panel, 16 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Phi-3.5 Mini	Compression	Yes	0.993	0.001	0.120	0.086	0.091
Phi-3.5 Mini	Compression	No	0.990	0.001	0.140	0.094	0.097
Phi-3.5 Mini	Conservative	Yes	0.998	0.000	0.050	0.037	0.021
Phi-3.5 Mini	Conservative	No	0.983	0.002	0.070	0.046	0.023
Phi-3.5 Mini	Mixing	Yes	0.988	0.003	0.150	0.089	0.100
Phi-3.5 Mini	Mixing	No	0.944	0.003	0.063	0.035	0.064
Qwen3 14B	Compression	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3 14B	Compression	No	1.000	0.000	0.000	0.000	0.000
Qwen3 14B	Conservative	Yes	0.992	0.001	0.110	0.071	0.039
Qwen3 14B	Conservative	No	0.992	0.001	0.110	0.071	0.031
Qwen3 14B	Mixing	Yes	0.999	0.008	0.143	0.010	0.016
Qwen3 14B	Mixing	No	0.989	0.013	0.259	0.062	0.071
Qwen3.5 35B	Compression	Yes	1.000	0.004	0.087	0.030	0.036
Qwen3.5 35B	Compression	No	0.997	0.003	0.086	0.041	0.057
Qwen3.5 35B	Conservative	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3.5 35B	Conservative	No	1.000	0.000	0.000	0.000	0.000
Qwen3.5 35B	Mixing	Yes	0.998	0.001	0.050	0.041	0.043
Qwen3.5 35B	Mixing	No	0.980	0.006	0.216	0.071	0.070

Table 37: Killed-operator estimates for attribute, 8B long-horizon relay, 32 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Qwen3 8B	Compression	Yes	1.000	0.016	0.531	0.115	0.115
Qwen3 8B	Compression	No	1.000	0.005	0.436	0.077	0.077
Qwen3 8B	Conservative	Yes	1.000	0.012	0.398	0.039	0.039
Qwen3 8B	Conservative	No	0.999	0.007	0.353	0.020	0.017
Qwen3 8B	Mixing	Yes	0.987	0.009	0.342	0.113	0.086
Qwen3 8B	Mixing	No	0.978	0.008	0.386	0.123	0.131

and repair status do not fully close the process. Refining the state can reduce this quantity by Proposition G.1; generation indexing addresses a different failure, namely systematic change in the operator over time.

Fig. 11 plots the stationary error minus time-product error against the weighted defect. Individual groups can favor either estimate because each is fit from finite design trajectories, yet the pooled

Table 38: Killed-operator estimates for banking77, 8B long-horizon relay, 32 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Qwen3 8B	Compression	Yes	1.000	0.003	0.319	0.045	0.045
Qwen3 8B	Compression	No	1.000	0.003	0.306	0.033	0.033
Qwen3 8B	Conservative	Yes	1.000	0.003	0.319	0.045	0.045
Qwen3 8B	Conservative	No	1.000	0.003	0.319	0.045	0.045
Qwen3 8B	Mixing	Yes	0.994	0.027	0.650	0.236	0.242
Qwen3 8B	Mixing	No	0.902	0.016	0.621	0.482	0.270

Table 39: Killed-operator estimates for intent, 8B long-horizon relay, 32 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Qwen3 8B	Compression	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3 8B	Compression	No	1.000	0.000	0.000	0.000	0.000
Qwen3 8B	Conservative	Yes	1.000	0.000	0.000	0.000	0.000
Qwen3 8B	Conservative	No	1.000	0.000	0.000	0.000	0.000
Qwen3 8B	Mixing	Yes	0.999	0.000	0.063	0.052	0.039
Qwen3 8B	Mixing	No	0.978	0.001	0.222	0.092	0.154

Table 40: Killed-operator estimates for tool, 8B long-horizon relay, 32 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Qwen3 8B	Compression	Yes	0.999	0.002	0.229	0.015	0.020
Qwen3 8B	Compression	No	0.999	0.002	0.229	0.015	0.020
Qwen3 8B	Conservative	Yes	1.000	0.002	0.270	0.036	0.036
Qwen3 8B	Conservative	No	1.000	0.002	0.270	0.036	0.036
Qwen3 8B	Mixing	Yes	0.985	0.021	0.547	0.213	0.145
Qwen3 8B	Mixing	No	0.979	0.009	0.481	0.180	0.096

Table 41: Killed-operator estimates for attribute, 35B long-horizon relay, 32 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Qwen3.5 35B	Compression	Yes	1.000	0.002	0.235	0.029	0.029
Qwen3.5 35B	Compression	No	1.000	0.002	0.235	0.029	0.029
Qwen3.5 35B	Conservative	Yes	1.000	0.002	0.235	0.029	0.029
Qwen3.5 35B	Conservative	No	1.000	0.002	0.000	0.025	0.025
Qwen3.5 35B	Mixing	Yes	0.999	0.002	0.223	0.027	0.025
Qwen3.5 35B	Mixing	No	0.987	0.002	0.201	0.104	0.045

Table 42: Killed-operator estimates for banking77, 35B long-horizon relay, 32 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Qwen3.5 35B	Compression	Yes	1.000	0.008	0.173	0.053	0.057
Qwen3.5 35B	Compression	No	0.990	0.006	0.219	0.091	0.089
Qwen3.5 35B	Conservative	Yes	1.000	0.013	0.086	0.034	0.068
Qwen3.5 35B	Conservative	No	1.000	0.010	0.071	0.048	0.068
Qwen3.5 35B	Mixing	Yes	0.996	0.014	0.202	0.098	0.115
Qwen3.5 35B	Mixing	No	0.961	0.015	0.288	0.074	0.104

paired reduction is positive with interval $[0.005, 0.009]$. Equation 43 explains this aggregate direction: temporal deviations accumulate in a stationary power and are represented explicitly in an ordered product.

Table 43: Killed-operator estimates for intent, 35B long-horizon relay, 32 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Qwen3.5 35B	Compression	Yes	1.000	0.002	0.204	0.020	0.021
Qwen3.5 35B	Compression	No	1.000	0.002	0.174	0.011	0.012
Qwen3.5 35B	Conservative	Yes	1.000	0.002	0.235	0.026	0.026
Qwen3.5 35B	Conservative	No	1.000	0.000	0.000	0.000	0.000
Qwen3.5 35B	Mixing	Yes	0.992	0.025	0.202	0.062	0.031
Qwen3.5 35B	Mixing	No	0.960	0.016	0.052	0.099	0.076

Table 44: Killed-operator estimates for tool, 35B long-horizon relay, 32 generations. Prediction columns are held-out mean absolute errors.

Rewriter	Relay	Repair	$\rho(K)$	$\bar{\delta}$	One step	Stationary	Time product
Qwen3.5 35B	Compression	Yes	1.000	0.019	0.276	0.048	0.044
Qwen3.5 35B	Compression	No	1.000	0.007	0.146	0.033	0.019
Qwen3.5 35B	Conservative	Yes	1.000	0.008	0.220	0.039	0.027
Qwen3.5 35B	Conservative	No	1.000	0.008	0.031	0.014	0.022
Qwen3.5 35B	Mixing	Yes	0.987	0.024	0.189	0.114	0.064
Qwen3.5 35B	Mixing	No	0.970	0.025	0.347	0.185	0.130

P.4 Finite-horizon prediction rather than curve fitting

All operator rows are estimated on design episodes and evaluated on held-out task instances. The horizon values are not fit separately. A single matrix power generates the stationary curve, and a single ordered sequence generates the time-indexed curve. This constraint is important because an unconstrained smoother could reproduce any empirical survival curve without identifying a transmission law. Here prediction succeeds only when the declared microstate and transition estimates transfer to unseen episodes.

Q Within-Fiber and Receiver-Transfer Measurements

This section gives the full measurements for two tests that isolate the central identifiability mechanism. The first changes the current message while holding its executable quotient class fixed. The second changes the frozen receiver after every trajectory has been generated.

Q.1 Natural within-fiber branch construction

We begin from the Qwen3.5 35B, mixing, no-repair trajectories in the main panel. For each task instance, the seed message defines an initial action vector under the frozen 14B and 35B receivers. Among later messages with exactly the same action vector, we select the one with maximum normalized edit distance. This rank rule finds a distinct within-fiber message for 190 of 192 source instances and uses no future outcome.

The seed and selected paraphrase are then evaluated by an independently frozen 2B and 4B receiver pair. The analysis retains the 160 source pairs whose two branch messages have the same noninvalid action vector on all eight repetitions. Each branch is passed to the same Qwen3.5-35B-A3B GPTQ-Int4 rewriter for eight additional generations at temperature 0.4. Corresponding branches use identical decoding seeds at every future generation. The resulting analysis contains $160 \times 8 \times 2 = 2{,}560$ future trajectories.

Table 45 reports the complete horizon profile. The action quotient agrees exactly at the branch point, yet average absolute survival separation increases from 0.077 after one future step to 0.144 after four and eight steps. At horizon 8, 35.6% of source pairs have different empirical survival laws across their common tapes. The mean signed contrast favors the seed by 0.056 with interval $[0.013, 0.102]$, while negative pair-level contrasts show that a paraphrase can also occupy the more persistent within-fiber state.

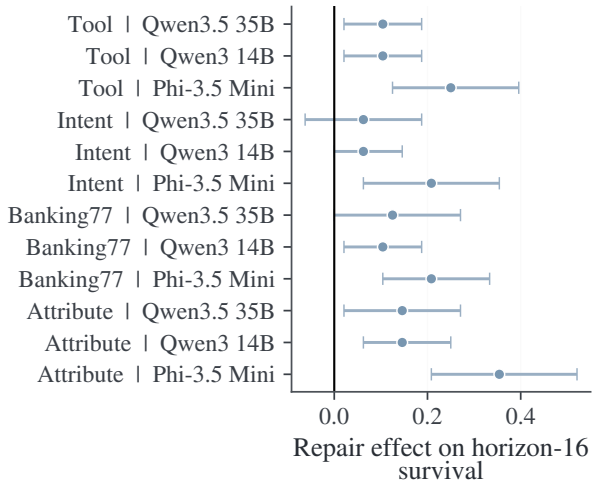

Figure 12: Paired repair effects under mixing in the 16-generation main panel. Points are changes in final intersection survival; bars are trajectory-bootstrap intervals.

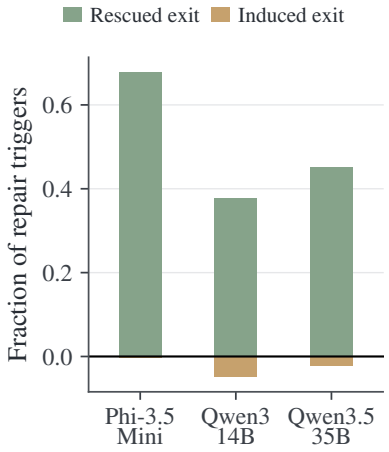

Figure 13: Immediate mediation among triggered repair events in the main panel.

Table 46 lists the largest separations. Four source pairs have survival 1.0 from the seed and 0.0 from the paraphrase across all eight common tapes. These are direct empirical lower bounds on $\Delta_8(\phi)$ in Equation 14. They do not rely on matching two different relay conditions or estimating one-step error.

Q.2 Held-out executable receivers

We replay every public message from the 3,456-trajectory main panel and the 1,536-trajectory 35B long panel through Qwen3 8B and Phi-3.5 Mini receivers. Neither receiver participates in the original quotient construction, operator estimation, repair trigger, or branch selection. The replay yields 9,984 complete trajectory-receiver records and 192 condition groups.

Table 47 reports cross-fitted prediction. Repeated one-step retention has pooled mean absolute error 0.113. A stationary killed operator reduces it to 0.048, while a time-indexed product has error 0.050. The main gain from retaining within-basin destination structure therefore transfers to both held-out action partitions. Temporal indexing remains condition dependent: it is slightly better at horizon 32 and slightly worse in the larger horizon-16 panel.

Table 48 reports the 35B horizon-32 mixing condition. High surface turnover again coexists with substantial survival under both new receivers. For attribute reference, turnover is 0.810 without repair while survival is 0.938 for Qwen3 8B and 0.906 for Phi-3.5 Mini. The ordering across tasks is also similar: attribute reference is most persistent, while Banking77 and intent routing accumulate more leakage.

The receiver replay changes the semantic partition rather than merely rescoring an unchanged label sequence. Each receiver chooses from the original candidate actions at every generation, and its own initial action defines its basin. The operator result is therefore a transfer test of the quotient construction itself.

R Observable Repair Results

The repair analysis separates immediate mediation from downstream policy effects. The first table reports paired effects on final survival, survival area, and first exit, while the second reports pre- to post-repair mediation. Fig. 13 summarizes immediate rescues and induced exits, and Fig. 12 shows their accumulated horizon effects under main-panel mixing; the tables disaggregate all relay instructions. Pairing is defined at the episode level, so each repair and no-repair trajectory shares its task instance, candidate assignment, model condition, and initial seed. This distinction matters

Table 45: Common-future separation for messages with identical dual-receiver action vectors at the branch point.

Task	Horizon	Pairs	Abs. gap	95% interval	Nonzero	Seed minus para.	95% interval
All	1	160	0.077	[0.044, 0.114]	0.163	0.061	[0.027, 0.099]
All	2	160	0.112	[0.075, 0.152]	0.269	0.063	[0.024, 0.105]
All	4	160	0.144	[0.105, 0.186]	0.325	0.070	[0.027, 0.116]
All	8	160	0.144	[0.105, 0.184]	0.356	0.056	[0.013, 0.102]
Attribute	1	40	0.031	[0.006, 0.062]	0.125	0.000	[-0.031, 0.028]
Attribute	2	40	0.106	[0.037, 0.188]	0.250	0.050	[-0.028, 0.131]
Attribute	4	40	0.119	[0.047, 0.206]	0.275	0.050	[-0.031, 0.141]
Attribute	8	40	0.128	[0.056, 0.212]	0.325	0.066	[-0.013, 0.156]
Banking77	1	33	0.106	[0.019, 0.216]	0.182	0.091	[0.000, 0.205]
Banking77	2	33	0.148	[0.072, 0.235]	0.394	0.027	[-0.064, 0.129]
Banking77	4	33	0.208	[0.117, 0.307]	0.455	0.057	[-0.057, 0.174]
Banking77	8	33	0.178	[0.106, 0.261]	0.485	0.011	[-0.091, 0.110]
Intent	1	44	0.108	[0.040, 0.190]	0.227	0.108	[0.040, 0.190]
Intent	2	44	0.114	[0.048, 0.193]	0.295	0.097	[0.028, 0.176]
Intent	4	44	0.168	[0.097, 0.247]	0.386	0.082	[0.000, 0.176]
Intent	8	44	0.176	[0.105, 0.253]	0.455	0.045	[-0.043, 0.136]
Tool	1	43	0.064	[0.009, 0.134]	0.116	0.047	[-0.009, 0.122]
Tool	2	43	0.087	[0.023, 0.172]	0.163	0.070	[0.003, 0.154]
Tool	4	43	0.093	[0.029, 0.172]	0.209	0.087	[0.023, 0.169]
Tool	8	43	0.099	[0.029, 0.183]	0.186	0.093	[0.023, 0.180]

Table 46: Largest horizon-8 common-future separations. Pair identifiers deterministically link rows to the selection file.

Task	Edit	Tapes	Seed	Paraphrase	Abs. gap	Tape disagree	Pair ID
Tool	0.532	8	1.000	0.000	1.000	1.000	10f81beb
Attribute	0.718	8	1.000	0.000	1.000	1.000	3286a6b5
Tool	0.426	8	1.000	0.000	1.000	1.000	7b1cc881
Attribute	0.618	8	1.000	0.000	1.000	1.000	a16098ba
Intent	0.714	8	1.000	0.125	0.875	0.875	02ba3787
Tool	0.750	8	1.000	0.125	0.875	0.875	468acd92
Intent	0.549	8	1.000	0.125	0.875	0.875	faf0db56
Intent	0.671	8	0.750	0.000	0.750	0.750	1c423bd8
Banking77	0.772	8	0.000	0.750	0.750	0.750	aa4dda34
Intent	0.788	8	0.875	0.250	0.625	0.625	3b9c9df2
Tool	0.515	8	0.625	0.000	0.625	0.625	4cccedd2
Banking77	0.725	8	0.125	0.750	0.625	0.625	53ac8b89
Attribute	0.525	8	0.625	0.000	0.625	0.625	dbbaa462
Banking77	0.672	8	1.000	0.375	0.625	0.625	e34ab10b
Attribute	0.625	8	0.625	0.000	0.625	0.625	f01ea0fc
Banking77	0.590	8	0.625	0.000	0.625	0.625	fa205fdd

Table 47: Operator prediction under receivers excluded from quotient construction and model selection.

Held-out receiver	Panel	Groups	One step MAE	Killed MAE	Time product MAE
Phi-3.5 Mini	Horizon 32	24	0.104	0.045	0.044
Phi-3.5 Mini	Horizon 16	72	0.122	0.049	0.052
Qwen3 8B	Horizon 32	24	0.055	0.024	0.024
Qwen3 8B	Horizon 16	72	0.126	0.056	0.059

because an immediate rescue can drift later, while an induced exit can still be followed by return after the first-passage event has failed.

The pilot repair contrasts appear in Tables 49, 50, and 51. The 8B horizon-16 contrasts appear in Tables 52, 53, 54, and 55. The main contrasts appear in Tables 56, 57, 58, and 59. The 8B long contrasts appear in Tables 60, 61, 62, and 63; the 35B long contrasts appear in Tables 64, 65, 66, and 67. Together they contain every paired intersection-survival contrast. Pairing holds the generated task, candidate assignment, episode index, and model condition fixed. The difference therefore compares

Table 48: Held-out receiver survival on the 35B horizon-32 mixing panel.

Held-out receiver	Task	Plain survival	Repair survival	Difference	Turnover
Qwen3 8B	Attribute	0.938	0.938	0.000	0.810
Qwen3 8B	Banking77	0.578	0.656	0.078	0.755
Qwen3 8B	Intent	0.734	0.828	0.094	0.635
Qwen3 8B	Tool	0.750	0.828	0.078	0.706
Phi-3.5 Mini	Attribute	0.906	0.938	0.031	0.810
Phi-3.5 Mini	Banking77	0.641	0.656	0.016	0.755
Phi-3.5 Mini	Intent	0.719	0.766	0.047	0.635
Phi-3.5 Mini	Tool	0.734	0.828	0.094	0.706

Table 49: Paired observable-repair effects for attribute, 2B and 4B pilot, 16 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Qwen3-VL 2B	Compression	24	-0.042	[-0.125, 0.000]	16	-0.875
Qwen3-VL 2B	Conservative	24	0.000	[0.000, 0.000]	32	0.000
Qwen3-VL 2B	Mixing	24	-0.042	[-0.125, 0.000]	22	-0.182
Qwen3-VL 4B	Compression	24	0.000	[0.000, 0.000]	32	-0.500
Qwen3-VL 4B	Conservative	24	0.042	[0.000, 0.125]	44	0.000
Qwen3-VL 4B	Mixing	24	0.125	[0.000, 0.250]	24	0.375

Table 50: Paired observable-repair effects for intent, 2B and 4B pilot, 16 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Qwen3-VL 2B	Compression	24	0.000	[0.000, 0.000]	0	0.000
Qwen3-VL 2B	Conservative	24	0.000	[0.000, 0.000]	0	0.000
Qwen3-VL 2B	Mixing	24	0.000	[0.000, 0.000]	16	0.062
Qwen3-VL 4B	Compression	24	0.083	[0.000, 0.208]	12	1.000
Qwen3-VL 4B	Conservative	24	0.000	[0.000, 0.000]	0	0.000
Qwen3-VL 4B	Mixing	24	0.167	[0.042, 0.333]	6	1.000

two declared transmission policies on common instances: one accepts every proposed rewrite, while the other invokes the observable seed-based repair rule after receiver disagreement.

R.1 Reading paired policy effects

The tables separate three estimands that are easy to conflate. Final-survival effects measure the probability of uninterrupted execution at the declared horizon. Survival-area effects summarize persistence across all saved horizons, and first-exit effects record when the event is first lost. A repair policy can improve one estimand without improving another, especially when it rescues early steps but changes the later relay distribution.

Pairing keeps the task instance, candidate assignment, episode index, and model condition fixed. The reported interval therefore measures variation across matched trajectories rather than variation caused by changing the sampled workload. Figure 13 supplies the local mechanism, while Figure 12 shows the downstream effect after subsequent rewrites have acted on the repaired message.

This distinction also explains why the exact sweep and language-model tables are complementary. The exact cells test an operator inequality directly; the empirical cells estimate policy contrasts from finite trajectories. Reading them together identifies where the sufficient domination mechanism is present and where repair remains helpful through a weaker, data-dependent path effect.

R.2 Immediate and downstream effects

The immediate mediation counts and paired horizon contrasts answer complementary questions. Immediate rescue asks whether repair moves the current proposal from outside to inside the receiver intersection. An induced exit is the reverse transition. Preserved survival and preserved exit complete

Table 51: Paired observable-repair effects for tool, 2B and 4B pilot, 16 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Qwen3-VL 2B	Compression	24	0.000	[0.000, 0.000]	15	0.000
Qwen3-VL 2B	Conservative	24	0.000	[0.000, 0.000]	0	0.000
Qwen3-VL 2B	Mixing	24	0.042	[-0.083, 0.167]	14	0.857
Qwen3-VL 4B	Compression	24	0.000	[0.000, 0.000]	0	0.000
Qwen3-VL 4B	Conservative	24	0.000	[0.000, 0.000]	0	0.000
Qwen3-VL 4B	Mixing	24	0.083	[0.000, 0.208]	5	1.000

Table 52: Paired observable-repair effects for attribute, 8B relay, 16 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Qwen3 8B	Compression	32	-0.031	[-0.094, 0.000]	48	0.000
Qwen3 8B	Conservative	32	0.062	[0.000, 0.156]	63	0.238
Qwen3 8B	Mixing	32	0.062	[0.000, 0.156]	38	0.263

Table 53: Paired observable-repair effects for banking77, 8B relay, 16 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Qwen3 8B	Compression	32	0.000	[0.000, 0.000]	33	0.000
Qwen3 8B	Conservative	32	0.000	[0.000, 0.000]	32	0.000
Qwen3 8B	Mixing	32	0.125	[0.031, 0.250]	34	0.824

Table 54: Paired observable-repair effects for intent, 8B relay, 16 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Qwen3 8B	Compression	32	0.000	[0.000, 0.000]	0	0.000
Qwen3 8B	Conservative	32	0.000	[0.000, 0.000]	0	0.000
Qwen3 8B	Mixing	32	0.188	[0.062, 0.344]	13	1.000

Table 55: Paired observable-repair effects for tool, 8B relay, 16 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Qwen3 8B	Compression	32	0.000	[0.000, 0.000]	22	0.000
Qwen3 8B	Conservative	32	0.000	[0.000, 0.000]	2	-0.500
Qwen3 8B	Mixing	32	0.062	[-0.062, 0.188]	63	0.619

Table 56: Paired observable-repair effects for attribute, 14B, 35B, and Phi main panel, 16 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Phi-3.5 Mini	Compression	48	0.125	[0.021, 0.229]	34	0.059
Phi-3.5 Mini	Conservative	48	0.125	[0.042, 0.229]	22	0.273
Phi-3.5 Mini	Mixing	48	0.354	[0.208, 0.521]	71	0.901
Qwen3 14B	Compression	48	-0.021	[-0.062, 0.000]	16	-0.062
Qwen3 14B	Conservative	48	0.042	[0.000, 0.104]	17	0.059
Qwen3 14B	Mixing	48	0.146	[0.062, 0.250]	30	0.967
Qwen3.5 35B	Compression	48	0.021	[0.000, 0.062]	11	0.818
Qwen3.5 35B	Conservative	48	0.000	[0.000, 0.000]	16	0.000
Qwen3.5 35B	Mixing	48	0.146	[0.021, 0.271]	30	0.467

the four-cell decomposition. These counts estimate local movement of basin mass at observed trigger states.

Table 57: Paired observable-repair effects for banking77, 14B, 35B, and Phi main panel, 16 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Phi-3.5 Mini	Compression	48	0.167	[0.062, 0.271]	111	0.559
Phi-3.5 Mini	Conservative	48	0.146	[0.062, 0.250]	95	0.358
Phi-3.5 Mini	Mixing	48	0.208	[0.104, 0.333]	152	0.757
Qwen3 14B	Compression	48	0.042	[-0.042, 0.125]	70	0.114
Qwen3 14B	Conservative	48	0.042	[0.000, 0.104]	95	0.063
Qwen3 14B	Mixing	48	0.104	[0.021, 0.188]	127	0.480
Qwen3.5 35B	Compression	48	0.062	[0.000, 0.146]	83	0.554
Qwen3.5 35B	Conservative	48	0.021	[-0.042, 0.104]	69	-0.014
Qwen3.5 35B	Mixing	48	0.125	[0.000, 0.271]	94	0.394

Table 58: Paired observable-repair effects for intent, 14B, 35B, and Phi main panel, 16 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Phi-3.5 Mini	Compression	48	0.021	[-0.042, 0.104]	8	1.000
Phi-3.5 Mini	Conservative	48	0.208	[0.104, 0.333]	15	1.000
Phi-3.5 Mini	Mixing	48	0.208	[0.062, 0.354]	66	0.985
Qwen3 14B	Compression	48	0.021	[0.000, 0.062]	1	1.000
Qwen3 14B	Conservative	48	0.000	[0.000, 0.000]	0	0.000
Qwen3 14B	Mixing	48	0.062	[0.000, 0.146]	4	1.000
Qwen3.5 35B	Compression	48	0.000	[0.000, 0.000]	0	0.000
Qwen3.5 35B	Conservative	48	0.000	[0.000, 0.000]	0	0.000
Qwen3.5 35B	Mixing	48	0.062	[-0.062, 0.188]	21	0.952

Table 59: Paired observable-repair effects for tool, 14B, 35B, and Phi main panel, 16 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Phi-3.5 Mini	Compression	48	0.062	[0.000, 0.146]	7	0.857
Phi-3.5 Mini	Conservative	48	0.062	[0.000, 0.146]	1	1.000
Phi-3.5 Mini	Mixing	48	0.250	[0.125, 0.396]	44	1.000
Qwen3 14B	Compression	48	0.000	[0.000, 0.000]	0	0.000
Qwen3 14B	Conservative	48	0.000	[0.000, 0.000]	0	0.000
Qwen3 14B	Mixing	48	0.104	[0.021, 0.188]	15	1.000
Qwen3.5 35B	Compression	48	0.021	[0.000, 0.062]	1	0.000
Qwen3.5 35B	Conservative	48	0.000	[0.000, 0.000]	0	0.000
Qwen3.5 35B	Mixing	48	0.104	[0.021, 0.188]	33	0.879

Table 60: Paired observable-repair effects for attribute, 8B long-horizon relay, 32 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Qwen3 8B	Compression	48	-0.021	[-0.062, 0.000]	25	-0.880
Qwen3 8B	Conservative	48	0.021	[0.000, 0.062]	121	0.058
Qwen3 8B	Mixing	48	0.104	[0.021, 0.188]	51	0.510

Table 61: Paired observable-repair effects for banking77, 8B long-horizon relay, 32 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Qwen3 8B	Compression	48	0.021	[0.000, 0.062]	65	0.015
Qwen3 8B	Conservative	48	0.000	[0.000, 0.000]	64	0.000
Qwen3 8B	Mixing	48	0.146	[0.062, 0.250]	71	0.437

The paired final-survival effect includes every later consequence of that repaired message. Once repair changes a message, subsequent rewrites receive a different input, so two paired trajectories can

Table 62: Paired observable-repair effects for intent, 8B long-horizon relay, 32 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Qwen3 8B	Compression	48	0.000	[0.000, 0.000]	0	0.000
Qwen3 8B	Conservative	48	0.000	[0.000, 0.000]	0	0.000
Qwen3 8B	Mixing	48	0.167	[0.062, 0.292]	56	0.982

Table 63: Paired observable-repair effects for tool, 8B long-horizon relay, 32 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Qwen3 8B	Compression	48	0.000	[0.000, 0.000]	32	0.000
Qwen3 8B	Conservative	48	0.000	[0.000, 0.000]	1	0.000
Qwen3 8B	Mixing	48	0.062	[-0.021, 0.146]	60	0.467

Table 64: Paired observable-repair effects for attribute, 35B long-horizon relay, 32 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Qwen3.5 35B	Compression	64	0.000	[0.000, 0.000]	5	1.000
Qwen3.5 35B	Conservative	64	0.000	[0.000, 0.000]	0	0.000
Qwen3.5 35B	Mixing	64	0.125	[0.031, 0.219]	23	0.957

Table 65: Paired observable-repair effects for banking77, 35B long-horizon relay, 32 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Qwen3.5 35B	Compression	64	0.141	[0.047, 0.234]	147	0.218
Qwen3.5 35B	Conservative	64	-0.016	[-0.047, 0.000]	113	0.035
Qwen3.5 35B	Mixing	64	0.219	[0.094, 0.344]	125	0.528

Table 66: Paired observable-repair effects for intent, 35B long-horizon relay, 32 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Qwen3.5 35B	Compression	64	0.016	[0.000, 0.047]	91	0.209
Qwen3.5 35B	Conservative	64	-0.031	[-0.078, 0.000]	3	-0.333
Qwen3.5 35B	Mixing	64	0.266	[0.156, 0.391]	60	0.650

Table 67: Paired observable-repair effects for tool, 35B long-horizon relay, 32 generations. The effect is repair minus no repair in final intersection survival.

Rewriter	Relay	n	Effect	95% interval	Triggers	Immediate net
Qwen3.5 35B	Compression	64	0.047	[0.000, 0.109]	61	-0.016
Qwen3.5 35B	Conservative	64	0.016	[0.000, 0.047]	48	-0.146
Qwen3.5 35B	Mixing	64	0.156	[0.062, 0.250]	120	0.650

diverge even when no later trigger occurs. This is the intended policy estimand. It measures whether observable repair extends the first-passage time of the composed relay, not only whether the current action is corrected.

Across the five panels, 3,287 intersection-event triggers produce 1,313 immediate rescues and 116 induced exits. The remaining triggers preserve the current event status. The positive net local movement is consistent with the main horizon effects. The nonzero induced-exit count is equally informative: the empirical repair kernel is not guaranteed to satisfy elementwise domination at every message, exactly as Theorem 6.1 predicts.

Table 68: Largest exact one-step matches in 2B and 4B pilot, 16 generations. Conditions share an identical design survivor count and differ only in held-out future survival.

Task	Event	Step 1	Condition A	Condition B	S_A	S_B	Gap
attribute	intersection	10/13	Q3VL-2B/Comp./R	Q3VL-4B/Mix/N	0.909	0.636	0.273
attribute	intersection	10/13	Q3VL-4B/Mix/R	Q3VL-4B/Mix/N	0.909	0.636	0.273
attribute	e0	11/13	Q3VL-2B/Comp./R	Q3VL-4B/Mix/N	0.909	0.727	0.182
attribute	e0	11/13	Q3VL-4B/Comp./R	Q3VL-4B/Mix/N	0.909	0.727	0.182
tool	e0	12/12	Q3VL-4B/Comp./R	Q3VL-2B/Mix/R	1.000	0.833	0.167
tool	e0	12/12	Q3VL-4B/Comp./R	Q3VL-2B/Mix/N	1.000	0.833	0.167
tool	e0	12/12	Q3VL-4B/Comp./N	Q3VL-2B/Mix/R	1.000	0.833	0.167
tool	e0	12/12	Q3VL-4B/Comp./N	Q3VL-2B/Mix/N	1.000	0.833	0.167
tool	e0	12/12	Q3VL-2B/Cons./R	Q3VL-2B/Mix/R	1.000	0.833	0.167
tool	e0	12/12	Q3VL-2B/Cons./R	Q3VL-2B/Mix/N	1.000	0.833	0.167
tool	e0	12/12	Q3VL-4B/Cons./R	Q3VL-2B/Mix/R	1.000	0.833	0.167
tool	e0	12/12	Q3VL-4B/Cons./R	Q3VL-2B/Mix/N	1.000	0.833	0.167
tool	e0	12/12	Q3VL-2B/Cons./N	Q3VL-2B/Mix/R	1.000	0.833	0.167
tool	e0	12/12	Q3VL-2B/Cons./N	Q3VL-2B/Mix/N	1.000	0.833	0.167
tool	e0	12/12	Q3VL-4B/Cons./N	Q3VL-2B/Mix/R	1.000	0.833	0.167
tool	e0	12/12	Q3VL-4B/Cons./N	Q3VL-2B/Mix/N	1.000	0.833	0.167

R.3 Pressure and task determine repair opportunity

Mixing creates the largest repair effects because it intentionally explores distant paraphrases. At horizon 32, the 35B mixing condition yields positive paired intervals on all four tasks. Intent routing gains 0.266, Banking77 0.219, tool routing 0.156, and attribute reference 0.125. The independent 8B long panel gives gains of 0.167, 0.146, 0.104, and 0.062, with the first three intervals above zero.

Conservative and compression conditions often have fewer trigger opportunities. Compression can also expose a distinct failure mode: a short proposal may omit exactly the relation that distinguishes two candidates. The repairer then benefits from the stored public seed, which restores information without consulting the gold action. The complete table atlas shows where this mechanism is active across rewriters.

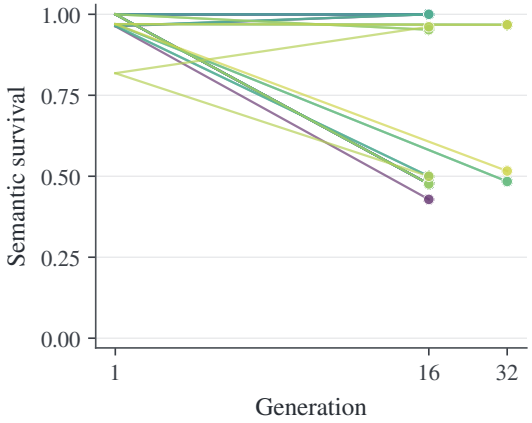

Figure 14: Exact empirical one-step matches with different held-out survival.

R.4 Relation to domination

The exact repair sweep separates the algebraic guarantee from empirical prevalence. Every exact cell satisfying $K_R \geq K_0$ improves or preserves all horizons. Cells outside this order include both signs. The language-model experiment estimates the direction through mediation and paired outcomes rather than claiming an unobserved elementwise order. This division gives the theorem and measurement distinct roles: the theorem states a sufficient mechanism, and the experiment shows that the observable policy frequently moves mass in that mechanism's favorable direction.

S Exact Empirical One-Step Matches

Each row pairs two transmission conditions with identical design sample size and identical integer one-step survivor count. Their later survival is measured on the common held-out instance set. These comparisons instantiate Theorem 5.2 without selecting a numerical matching tolerance.

Fig. 14 visualizes the largest separations. Each fork starts from the same design one-step retention and terminates at two held-out future survival values.

Table 69: Largest exact one-step matches in 8B relay, 16 generations. Conditions share an identical design survivor count and differ only in held-out future survival.

Task	Event	Step 1	Condition A	Condition B	S_A	S_B	Gap
tool	e1	16/18	Q3-8B/Cons./N	Q3-8B/Mix/R	1.000	0.714	0.286
tool	intersection	16/18	Q3-8B/Cons./R	Q3-8B/Mix/R	1.000	0.714	0.286
tool	intersection	16/18	Q3-8B/Cons./N	Q3-8B/Mix/R	1.000	0.714	0.286
tool	e0	17/18	Q3-8B/Cons./R	Q3-8B/Mix/R	1.000	0.786	0.214
tool	e0	17/18	Q3-8B/Cons./N	Q3-8B/Mix/R	1.000	0.786	0.214
tool	union	17/18	Q3-8B/Cons./N	Q3-8B/Mix/R	1.000	0.786	0.214
intent	intersection	17/17	Q3-8B/Comp./R	Q3-8B/Mix/N	1.000	0.800	0.200
intent	intersection	17/17	Q3-8B/Comp./N	Q3-8B/Mix/N	1.000	0.800	0.200
intent	intersection	17/17	Q3-8B/Cons./R	Q3-8B/Mix/N	1.000	0.800	0.200
intent	intersection	17/17	Q3-8B/Cons./N	Q3-8B/Mix/N	1.000	0.800	0.200
intent	intersection	17/17	Q3-8B/Mix/R	Q3-8B/Mix/N	1.000	0.800	0.200
banking77	intersection	13/15	Q3-8B/Mix/R	Q3-8B/Mix/N	0.882	0.706	0.176
tool	e1	17/18	Q3-8B/Cons./R	Q3-8B/Mix/N	1.000	0.857	0.143
tool	union	17/18	Q3-8B/Mix/R	Q3-8B/Mix/N	0.786	0.929	0.143
intent	e0	17/17	Q3-8B/Comp./R	Q3-8B/Mix/N	1.000	0.867	0.133
intent	e0	17/17	Q3-8B/Comp./N	Q3-8B/Mix/N	1.000	0.867	0.133

Table 70: Largest exact one-step matches in 14B, 35B, and Phi main panel, 16 generations. Conditions share an identical design survivor count and differ only in held-out future survival.

Task	Event	Step 1	Condition A	Condition B	S_A	S_B	Gap
intent	e0	26/27	Q3-14B/Comp./N	Phi-3.5/Mix/N	1.000	0.429	0.571
intent	union	27/27	Q3.5-35B/Comp./R	Phi-3.5/Mix/N	1.000	0.476	0.524
intent	union	27/27	Q3-14B/Comp./R	Phi-3.5/Mix/N	1.000	0.476	0.524
intent	union	27/27	Q3.5-35B/Comp./N	Phi-3.5/Mix/N	1.000	0.476	0.524
intent	union	27/27	Q3-14B/Comp./N	Phi-3.5/Mix/N	1.000	0.476	0.524
intent	union	27/27	Q3.5-35B/Cons./R	Phi-3.5/Mix/N	1.000	0.476	0.524
intent	union	27/27	Q3-14B/Cons./R	Phi-3.5/Mix/N	1.000	0.476	0.524
intent	union	27/27	Q3.5-35B/Cons./N	Phi-3.5/Mix/N	1.000	0.476	0.524
intent	union	27/27	Q3-14B/Cons./N	Phi-3.5/Mix/N	1.000	0.476	0.524
intent	union	27/27	Q3-14B/Mix/R	Phi-3.5/Mix/N	1.000	0.476	0.524
tool	intersection	27/28	Q3.5-35B/Comp./R	Phi-3.5/Mix/N	1.000	0.500	0.500
tool	intersection	27/28	Q3.5-35B/Comp./N	Phi-3.5/Mix/N	1.000	0.500	0.500
intent	union	27/27	Phi-3.5/Cons./R	Phi-3.5/Mix/N	0.952	0.476	0.476
intent	union	27/27	Phi-3.5/Mix/N	Q3-14B/Mix/N	0.476	0.952	0.476
banking77	union	18/22	Q3-14B/Mix/R	Phi-3.5/Mix/N	0.962	0.500	0.462
tool	e1	28/28	Q3-14B/Comp./R	Phi-3.5/Mix/N	1.000	0.550	0.450

Table 71: Largest exact one-step matches in 8B long-horizon relay, 32 generations. Conditions share an identical design survivor count and differ only in held-out future survival.

Task	Event	Step 1	Condition A	Condition B	S_A	S_B	Gap
intent	intersection	27/27	Q3-8B/Comp./R	Q3-8B/Mix/N	1.000	0.619	0.381
intent	intersection	27/27	Q3-8B/Comp./N	Q3-8B/Mix/N	1.000	0.619	0.381
intent	intersection	27/27	Q3-8B/Cons./R	Q3-8B/Mix/N	1.000	0.619	0.381
intent	intersection	27/27	Q3-8B/Cons./N	Q3-8B/Mix/N	1.000	0.619	0.381
intent	e0	27/27	Q3-8B/Comp./R	Q3-8B/Mix/N	1.000	0.667	0.333
intent	e0	27/27	Q3-8B/Comp./N	Q3-8B/Mix/N	1.000	0.667	0.333
intent	e0	27/27	Q3-8B/Cons./R	Q3-8B/Mix/N	1.000	0.667	0.333
intent	e0	27/27	Q3-8B/Cons./N	Q3-8B/Mix/N	1.000	0.667	0.333
intent	intersection	27/27	Q3-8B/Mix/R	Q3-8B/Mix/N	0.905	0.619	0.286
intent	e0	27/27	Q3-8B/Mix/R	Q3-8B/Mix/N	0.905	0.667	0.238
intent	e1	27/27	Q3-8B/Comp./R	Q3-8B/Mix/N	1.000	0.810	0.190
intent	e1	27/27	Q3-8B/Comp./N	Q3-8B/Mix/N	1.000	0.810	0.190
intent	e1	27/27	Q3-8B/Cons./R	Q3-8B/Mix/N	1.000	0.810	0.190
intent	e1	27/27	Q3-8B/Cons./N	Q3-8B/Mix/N	1.000	0.810	0.190
attribute	intersection	23/26	Q3-8B/Comp./R	Q3-8B/Mix/N	1.000	0.818	0.182
intent	union	27/27	Q3-8B/Comp./R	Q3-8B/Mix/N	1.000	0.857	0.143

Tables 68, 69, 70, 71, and 72 list the largest separations in each panel. Matching is performed on the design split, while future separation is measured on held-out task instances. Thus the same data are

Table 72: Largest exact one-step matches in 35B long-horizon relay, 32 generations. Conditions share an identical design survivor count and differ only in held-out future survival.

Task	Event	Step 1	Condition A	Condition B	S_A	S_B	Gap
intent	intersection	32/33	Q3.5-35B/Comp./R	Q3.5-35B/Mix/N	0.968	0.484	0.484
intent	intersection	32/33	Q3.5-35B/Cons./R	Q3.5-35B/Mix/N	0.968	0.484	0.484
intent	e0	32/33	Q3.5-35B/Cons./R	Q3.5-35B/Mix/N	0.968	0.516	0.452
intent	e1	33/33	Q3.5-35B/Cons./R	Q3.5-35B/Mix/N	1.000	0.548	0.452
intent	e1	33/33	Q3.5-35B/Cons./N	Q3.5-35B/Mix/N	1.000	0.548	0.452
intent	intersection	32/33	Q3.5-35B/Comp./N	Q3.5-35B/Mix/N	0.935	0.484	0.452
intent	intersection	32/33	Q3.5-35B/Mix/R	Q3.5-35B/Mix/N	0.871	0.484	0.387
intent	union	33/33	Q3.5-35B/Comp./R	Q3.5-35B/Mix/N	1.000	0.613	0.387
intent	union	33/33	Q3.5-35B/Comp./N	Q3.5-35B/Mix/N	1.000	0.613	0.387
intent	union	33/33	Q3.5-35B/Cons./R	Q3.5-35B/Mix/N	1.000	0.613	0.387
intent	union	33/33	Q3.5-35B/Cons./N	Q3.5-35B/Mix/N	1.000	0.613	0.387
intent	e0	32/33	Q3.5-35B/Mix/R	Q3.5-35B/Mix/N	0.871	0.516	0.355
intent	e1	33/33	Q3.5-35B/Mix/R	Q3.5-35B/Mix/N	0.903	0.548	0.355
intent	union	33/33	Q3.5-35B/Mix/R	Q3.5-35B/Mix/N	0.903	0.613	0.290
banking77	union	32/32	Q3.5-35B/Cons./R	Q3.5-35B/Mix/N	0.969	0.688	0.281
banking77	union	32/32	Q3.5-35B/Cons./N	Q3.5-35B/Mix/N	0.969	0.688	0.281

never used both to declare equality and to measure the later difference. Condition codes use `Comp.`, `Cons.`, and `Mix` for the three relay instructions, with `R` for repair and `N` for no repair.

S.1 Why integer matching matters

Empirical probabilities are ratios of survivor counts. A tolerance-based match would introduce a free scale and could favor comparisons that look equal only under a chosen rounding rule. We instead require both conditions to have the same number of design episodes and exactly the same integer number of one-step survivors. This criterion is invariant to numeric formatting and produces 2,366 matched pairs across the full study.

Future survival differs in 1,473 of those pairs. The largest held-out gap is 0.571 in the main panel. In the 35B long panel, repaired compression and unrepaired mixing each retain 32 of 33 design intent trajectories after one step, but their held-out generation-32 survival values are 0.968 and 0.484. The matched statistic has therefore discarded the destination of surviving mass, which is precisely the information retained by a killed operator. The matching and evaluation episode sets are disjoint. Exact matches recur across sample sizes, model scales, relay families, and horizons. Their held-out separation therefore locates the missing information in the one-step summary rather than in a particular denominator or one isolated model panel.

S.2 Neural realization of the storage construction

The exact kernels in Theorem 5.2 use two abstract within-basin states. In a language relay, those states correspond to message forms that currently execute the same action but induce different rewrite distributions. A concise canonical request can remain in a stable region, while an unusual paraphrase can sit near a later receiver boundary. Equal one-step action preservation does not reveal which region received the surviving mass.

The empirical pairs do not require byte-identical messages or equal surface statistics. Their purpose is to test the claimed insufficiency of one-step semantic error under common task sampling. The independent 2B and 4B pilot, 8B panels, main cross-family panel, and 35B long panel all contain nonzero separations. This replication across model configurations makes the result a property of the measurement abstraction rather than one isolated prompt.

S.3 Connection to state diameter

Every nonzero pair lower-bounds the horizon diameter of the observation that keeps only task identity and one-step retention. Theorem 7.1 converts that observed separation into a minimax prediction lower bound. A richer observable state can shrink the diameter by retaining receiver agreement, repair

Table 73: Receiver-0 task behavior and first passage for 2B and 4B pilot, 16 generations. Initial and final correctness use dataset or generated-task labels.

Rewriter	Task	Relay	Initial acc.	Final acc.	Survival	Exit gen.	Return
Qwen3-VL 2B	attribute	Compression	0.875	0.958	0.917	1.000	0.000
Qwen3-VL 2B	attribute	Conservative	0.875	0.917	0.958	1.000	0.000
Qwen3-VL 2B	attribute	Mixing	0.875	1.000	0.875	1.667	0.000
Qwen3-VL 2B	intent	Compression	1.000	1.000	1.000	–	–
Qwen3-VL 2B	intent	Conservative	1.000	1.000	1.000	–	–
Qwen3-VL 2B	intent	Mixing	1.000	0.917	0.917	2.000	0.000
Qwen3-VL 2B	tool	Compression	1.000	0.958	0.958	2.000	0.000
Qwen3-VL 2B	tool	Conservative	1.000	1.000	1.000	–	–
Qwen3-VL 2B	tool	Mixing	1.000	0.917	0.875	6.667	0.667
Qwen3-VL 4B	attribute	Compression	0.875	0.958	0.917	1.000	0.000
Qwen3-VL 4B	attribute	Conservative	0.875	0.917	0.958	1.000	0.000
Qwen3-VL 4B	attribute	Mixing	0.875	0.875	0.750	4.167	0.333
Qwen3-VL 4B	intent	Compression	1.000	0.917	0.917	2.000	0.000
Qwen3-VL 4B	intent	Conservative	1.000	1.000	1.000	–	–
Qwen3-VL 4B	intent	Mixing	1.000	0.917	0.833	4.500	0.500
Qwen3-VL 4B	tool	Compression	1.000	1.000	1.000	–	–
Qwen3-VL 4B	tool	Conservative	1.000	1.000	1.000	–	–
Qwen3-VL 4B	tool	Mixing	1.000	0.917	0.875	3.667	0.667

Table 74: Receiver-0 task behavior and first passage for 8B relay, 16 generations. Initial and final correctness use dataset or generated-task labels.

Rewriter	Task	Relay	Initial acc.	Final acc.	Survival	Exit gen.	Return
Qwen3 8B	attribute	Compression	1.000	1.000	1.000	–	–
Qwen3 8B	attribute	Conservative	1.000	0.969	0.969	2.000	1.000
Qwen3 8B	attribute	Mixing	1.000	0.906	0.844	4.200	1.000
Qwen3 8B	banking77	Compression	0.938	0.906	0.969	11.000	0.000
Qwen3 8B	banking77	Conservative	0.938	0.938	1.000	–	–
Qwen3 8B	banking77	Mixing	0.938	0.844	0.781	4.571	0.857
Qwen3 8B	intent	Compression	1.000	1.000	1.000	–	–
Qwen3 8B	intent	Conservative	1.000	1.000	1.000	–	–
Qwen3 8B	intent	Mixing	1.000	0.875	0.844	5.000	0.400
Qwen3 8B	tool	Compression	1.000	0.875	0.875	3.500	0.500
Qwen3 8B	tool	Conservative	1.000	1.000	0.969	1.000	1.000
Qwen3 8B	tool	Mixing	1.000	0.844	0.750	3.125	0.750

status, or message occupancy; the operator comparison measures how much predictive information those additions recover.

T Model and Task Panels

Tables in this section disaggregate final survival, surface turnover, edit distance, receiver ambiguity, repair trigger frequency, and gold task behavior by model, task, transmission instruction, and horizon.

Tables 73, 74, 75, 76, and 77 place external task behavior beside evaluator-conditioned survival. They use receiver 0 and unrepaired transmission so that correctness, first passage, and return can be read on a common trajectory set.

T.1 Correctness and semantic survival

Initial correctness asks whether a frozen receiver selects the generated or dataset target before any rewrite. Final correctness asks the same question at the last generation. Semantic survival instead asks whether every intermediate receiver action equals that receiver's own initial action. These quantities answer different questions. An initially incorrect receiver can have perfect semantic survival, meaning the relay consistently preserves its executable interpretation. A correct receiver can leave and later return to the gold action, giving a correct endpoint but a failed first-passage event.

Table 75: Receiver-0 task behavior and first passage for 14B, 35B, and Phi main panel, 16 generations. Initial and final correctness use dataset or generated-task labels.

Rewriter	Task	Relay	Initial acc.	Final acc.	Survival	Exit gen.	Return
Phi-3.5 Mini	attribute	Compression	1.000	0.833	0.812	5.222	0.333
Phi-3.5 Mini	attribute	Conservative	1.000	0.896	0.833	7.250	0.500
Phi-3.5 Mini	attribute	Mixing	1.000	0.625	0.438	4.444	0.704
Phi-3.5 Mini	banking77	Compression	0.854	0.542	0.604	4.316	0.263
Phi-3.5 Mini	banking77	Conservative	0.854	0.646	0.646	3.824	0.647
Phi-3.5 Mini	banking77	Mixing	0.854	0.479	0.417	4.536	0.536
Phi-3.5 Mini	intent	Compression	1.000	0.917	0.917	7.250	0.000
Phi-3.5 Mini	intent	Conservative	1.000	0.833	0.771	7.818	0.636
Phi-3.5 Mini	intent	Mixing	1.000	0.458	0.375	4.567	0.467
Phi-3.5 Mini	tool	Compression	1.000	0.854	0.833	6.625	0.125
Phi-3.5 Mini	tool	Conservative	1.000	0.917	0.917	4.250	0.500
Phi-3.5 Mini	tool	Mixing	1.000	0.833	0.708	5.143	0.643
Qwen3 14B	attribute	Compression	1.000	1.000	1.000	–	–
Qwen3 14B	attribute	Conservative	1.000	1.000	0.958	7.000	1.000
Qwen3 14B	attribute	Mixing	1.000	0.833	0.792	2.700	0.900
Qwen3 14B	banking77	Compression	0.854	0.854	0.875	1.667	0.167
Qwen3 14B	banking77	Conservative	0.854	0.833	0.896	1.200	0.800
Qwen3 14B	banking77	Mixing	0.854	0.771	0.729	3.308	0.615
Qwen3 14B	intent	Compression	1.000	0.979	0.979	1.000	0.000
Qwen3 14B	intent	Conservative	1.000	1.000	1.000	–	–
Qwen3 14B	intent	Mixing	1.000	0.958	0.917	5.500	0.500
Qwen3 14B	tool	Compression	1.000	1.000	1.000	–	–
Qwen3 14B	tool	Conservative	1.000	0.875	0.875	2.667	0.000
Qwen3 14B	tool	Mixing	1.000	0.875	0.854	2.286	0.857
Qwen3.5 35B	attribute	Compression	1.000	0.979	0.979	1.000	0.000
Qwen3.5 35B	attribute	Conservative	1.000	1.000	1.000	–	–
Qwen3.5 35B	attribute	Mixing	1.000	0.896	0.812	8.000	0.556
Qwen3.5 35B	banking77	Compression	0.854	0.771	0.812	1.444	0.000
Qwen3.5 35B	banking77	Conservative	0.854	0.833	0.938	2.667	0.667
Qwen3.5 35B	banking77	Mixing	0.854	0.771	0.583	5.550	0.700
Qwen3.5 35B	intent	Compression	1.000	1.000	1.000	–	–
Qwen3.5 35B	intent	Conservative	1.000	1.000	1.000	–	–
Qwen3.5 35B	intent	Mixing	1.000	0.875	0.833	8.125	0.500
Qwen3.5 35B	tool	Compression	1.000	0.979	0.979	4.000	0.000
Qwen3.5 35B	tool	Conservative	1.000	1.000	1.000	–	–
Qwen3.5 35B	tool	Mixing	1.000	0.896	0.812	5.556	0.778

Table 76: Receiver-0 task behavior and first passage for 8B long-horizon relay, 32 generations. Initial and final correctness use dataset or generated-task labels.

Rewriter	Task	Relay	Initial acc.	Final acc.	Survival	Exit gen.	Return
Qwen3 8B	attribute	Compression	0.979	1.000	0.979	1.000	0.000
Qwen3 8B	attribute	Conservative	0.979	0.979	0.958	1.500	0.500
Qwen3 8B	attribute	Mixing	0.979	0.875	0.792	4.800	0.600
Qwen3 8B	banking77	Compression	1.000	1.000	0.979	1.000	1.000
Qwen3 8B	banking77	Conservative	1.000	0.979	0.979	1.000	0.000
Qwen3 8B	banking77	Mixing	1.000	0.812	0.708	5.643	0.643
Qwen3 8B	intent	Compression	1.000	1.000	1.000	–	–
Qwen3 8B	intent	Conservative	1.000	1.000	1.000	–	–
Qwen3 8B	intent	Mixing	1.000	0.875	0.812	4.778	0.778
Qwen3 8B	tool	Compression	1.000	0.958	0.958	7.500	0.000
Qwen3 8B	tool	Conservative	1.000	1.000	1.000	–	–
Qwen3 8B	tool	Mixing	1.000	0.896	0.792	5.700	0.700

The separation prevents two common confounds. First, model competence does not redefine the semantic fiber after outcomes are observed. Second, stable behavior is not automatically labeled correct. The quotient is fixed by the declared receiver and context, while gold labels remain an external measurement of task utility.

Table 77: Receiver-0 task behavior and first passage for 35B long-horizon relay, 32 generations. Initial and final correctness use dataset or generated-task labels.

Rewriter	Task	Relay	Initial acc.	Final acc.	Survival	Exit gen.	Return
Qwen3.5 35B	attribute	Compression	0.984	1.000	0.984	1.000	0.000
Qwen3.5 35B	attribute	Conservative	0.984	1.000	0.984	2.000	0.000
Qwen3.5 35B	attribute	Mixing	0.984	0.922	0.828	17.727	0.727
Qwen3.5 35B	banking77	Compression	0.891	0.812	0.828	4.091	0.273
Qwen3.5 35B	banking77	Conservative	0.891	0.906	0.969	2.000	0.500
Qwen3.5 35B	banking77	Mixing	0.891	0.703	0.641	9.957	0.870
Qwen3.5 35B	intent	Compression	0.984	0.953	0.969	15.500	0.000
Qwen3.5 35B	intent	Conservative	0.984	0.984	1.000	–	–
Qwen3.5 35B	intent	Mixing	0.984	0.719	0.625	13.042	0.583
Qwen3.5 35B	tool	Compression	0.984	0.891	0.906	2.167	0.000
Qwen3.5 35B	tool	Conservative	0.984	0.969	0.984	3.000	0.000
Qwen3.5 35B	tool	Mixing	0.984	0.844	0.734	11.941	0.765

T.2 Observed exit generation

Mean first exit is computed among trajectories with an observed exit. It is therefore a conditional timing statistic, not an estimate of the unconditional expected exit in Equation 7. Conditions with few exits can have a late conditional mean and high survival; conditions with many early exits have both low survival and a small conditional mean. The survival curve and exit timing should be read together.

T.3 Return after exit

Return is the fraction of exited trajectories that later revisit the initial action. It identifies recoverable excursions in the full transmission operator. The killed operator assigns zero mass after the first exit because its target is safe uninterrupted transmission. Reporting return beside killed survival shows how much information that safety convention intentionally discards.

T.4 Task geometry

Attribute reference, manual intent routing, tool routing, and Banking77 vary the density of action boundaries. The attribute task supplies explicit redundant descriptors, while Banking77 contains naturally varied utterances and fine-grained neighboring intents. Tool routing lies between these settings, with distinct actions but occasional lexical overlap. The condition tables reveal how this geometry interacts with compression and mixing without changing the receiver definition. Candidate cardinality and scoring remain fixed, so the comparison changes boundary geometry without changing the receiver's decision format.

U Sensitivity and Boundary Sweeps

The exact grids vary leakage, mixing, contraction, detector behavior, and time variation independently. Language-model panels vary task, rewriter scale, receiver event, transmission pressure, repair, and horizon. Together these axes locate where a stationary quotient is adequate and where generation indexing is required. Every sweep preserves its initial semantic mass and reporting horizons, so cross-cell variation isolates the declared operator coordinate.

Tables 78, 79, 80, 81, and 82 report every matched-error horizon. Table 83 gives all approximate-lumpability cells, Table 84 gives all generation-indexed cells, and Tables 85 and 86 span the repair and mixing boundaries.

The sequence moves from fixed one-step loss to closure, temporal drift, intervention, and surface geometry. Each block changes one operator coordinate while preserving its comparison law, so the resulting boundary can be attributed to that coordinate rather than to a changed task or reporting horizon.

Table 78: Exact matched-error sweep at one-step error 0.01.

Horizon	Fast	Storage	Gap	ρ_{fast}	ρ_{store}
1	0.9900	0.9900	0.0000	0.9900	1.0000
2	0.9801	0.9900	0.0099	0.9900	1.0000
4	0.9606	0.9900	0.0294	0.9900	1.0000
8	0.9227	0.9900	0.0672	0.9900	1.0000
16	0.8515	0.9900	0.1385	0.9900	1.0000
32	0.7250	0.9900	0.2650	0.9900	1.0000
64	0.5256	0.9899	0.4643	0.9900	1.0000

Table 79: Exact matched-error sweep at one-step error 0.03.

Horizon	Fast	Storage	Gap	ρ_{fast}	ρ_{store}
1	0.9750	0.9750	0.0000	0.9750	1.0000
2	0.9506	0.9750	0.0244	0.9750	1.0000
4	0.9037	0.9750	0.0713	0.9750	1.0000
8	0.8167	0.9750	0.1583	0.9750	1.0000
16	0.6669	0.9750	0.3081	0.9750	1.0000
32	0.4448	0.9750	0.5302	0.9750	1.0000
64	0.1978	0.9749	0.7771	0.9750	1.0000

Table 80: Exact matched-error sweep at one-step error 0.05.

Horizon	Fast	Storage	Gap	ρ_{fast}	ρ_{store}
1	0.9500	0.9500	0.0000	0.9500	1.0000
2	0.9025	0.9500	0.0475	0.9500	1.0000
4	0.8145	0.9500	0.1355	0.9500	1.0000
8	0.6634	0.9500	0.2866	0.9500	1.0000
16	0.4401	0.9500	0.5099	0.9500	1.0000
32	0.1937	0.9500	0.7563	0.9500	1.0000
64	0.0375	0.9499	0.9124	0.9500	1.0000

Table 81: Exact matched-error sweep at one-step error 0.10.

Horizon	Fast	Storage	Gap	ρ_{fast}	ρ_{store}
1	0.9000	0.9000	0.0000	0.9000	1.0000
2	0.8100	0.9000	0.0900	0.9000	1.0000
4	0.6561	0.9000	0.2439	0.9000	1.0000
8	0.4305	0.9000	0.4695	0.9000	1.0000
16	0.1853	0.9000	0.7147	0.9000	1.0000
32	0.0343	0.9000	0.8656	0.9000	1.0000
64	0.0012	0.8999	0.8988	0.9000	1.0000

Table 82: Exact matched-error sweep at one-step error 0.20.

Horizon	Fast	Storage	Gap	ρ_{fast}	ρ_{store}
1	0.8000	0.8000	0.0000	0.8000	1.0000
2	0.6400	0.8000	0.1600	0.8000	1.0000
4	0.4096	0.8000	0.3904	0.8000	1.0000
8	0.1678	0.8000	0.6322	0.8000	1.0000
16	0.0281	0.8000	0.7718	0.8000	1.0000
32	0.0008	0.8000	0.7992	0.8000	1.0000
64	0.0000	0.7999	0.7999	0.8000	1.0000

U.1 Matched-error envelope

The five one-step error values separate immediately after generation one. For small error, repeated exposure decays slowly and the distinction is primarily long-horizon. For larger error, storage and repeated exposure diverge within a few generations. In every case, the storage process preserves the

Table 83: Complete approximate-lumpability sweep.

Defect	Horizon	Observed TV	Linear bound	Contracted bound	Holds
0.0000	1.0000	0.0000	0.0000	0.0000	True
0.0000	2.0000	0.0000	0.0000	0.0000	True
0.0000	4.0000	0.0000	0.0000	0.0000	True
0.0000	8.0000	0.0000	0.0000	0.0000	True
0.0000	16.0000	0.0000	0.0000	0.0000	True
0.0000	32.0000	0.0000	0.0000	0.0000	True
0.0100	1.0000	0.0050	0.0100	0.0100	True
0.0100	2.0000	0.0039	0.0200	0.0177	True
0.0100	4.0000	0.0023	0.0400	0.0282	True
0.0100	8.0000	0.0008	0.0800	0.0381	True
0.0100	16.0000	0.0001	0.1600	0.0428	True
0.0100	32.0000	0.0000	0.3200	0.0435	True
0.0250	1.0000	0.0125	0.0250	0.0250	True
0.0250	2.0000	0.0096	0.0500	0.0443	True
0.0250	4.0000	0.0057	0.1000	0.0705	True
0.0250	8.0000	0.0020	0.2000	0.0953	True
0.0250	16.0000	0.0002	0.4000	0.1070	True
0.0250	32.0000	0.0000	0.8000	0.1087	True
0.0500	1.0000	0.0250	0.0500	0.0500	True
0.0500	2.0000	0.0192	0.1000	0.0885	True
0.0500	4.0000	0.0114	0.2000	0.1410	True
0.0500	8.0000	0.0040	0.4000	0.1905	True
0.0500	16.0000	0.0005	0.8000	0.2141	True
0.0500	32.0000	0.0000	1.6000	0.2173	True
0.1000	1.0000	0.0500	0.1000	0.1000	True
0.1000	2.0000	0.0385	0.2000	0.1770	True
0.1000	4.0000	0.0228	0.4000	0.2819	True
0.1000	8.0000	0.0080	0.8000	0.3811	True
0.1000	16.0000	0.0010	1.6000	0.4281	True
0.1000	32.0000	0.0000	3.2000	0.4347	True

one-step surviving mass, while the fast process multiplies the same initial retention at every step. This sweep shows that the separation is not tied to one numerical example.

U.2 Approximate closure

The approximate-lumpability cells cross declared defect with horizon. Observed total variation remains below both the general linear accumulation and the contracted bound. At zero defect, the projected and quotient laws coincide. At positive defect, contraction prevents early row discrepancies from contributing equally at every later generation. This is the finite-horizon behavior predicted by Equation 5.

U.3 Time variation

The generation-indexed sweep alternates killed retention across time. The ordered product is the exact survival law. Replacing the sequence by a stationary reference introduces a horizon-dependent error, while the telescoping sum remains an upper bound in every cell. The bound is intentionally general; it controls the error without assuming that temporal deviations cancel.

U.4 Repair geometry

The exact repair grid varies detector sensitivity, false triggers, repair success, repair-induced leakage, and horizon. Detector quality alone does not determine survival. A sensitive detector paired with a leaking repairer can reduce basin mass, while a selective detector paired with a strong repair row can extend first passage. Elementwise domination is the invariant that turns these components into a uniform horizon guarantee.

Table 84: Complete generation-indexed product sweep.

Horizon	Product	Stationary	Error	Bound	Holds
1.0000	0.9800	0.9500	0.0300	0.1900	True
2.0000	0.8928	0.8845	0.0083	0.3800	True
3.0000	0.8578	0.8132	0.0446	0.5700	True
4.0000	0.7503	0.7416	0.0087	0.7600	True
5.0000	0.7152	0.6727	0.0425	0.9500	True
6.0000	0.6149	0.6080	0.0068	1.1400	True
7.0000	0.5841	0.5483	0.0358	1.3300	True
8.0000	0.4984	0.4937	0.0047	1.5200	True
9.0000	0.4727	0.4440	0.0287	1.7100	True
10.0000	0.4020	0.3991	0.0029	1.9000	True
11.0000	0.3810	0.3585	0.0225	2.0900	True
12.0000	0.3235	0.3220	0.0015	2.2800	True
13.0000	0.3065	0.2891	0.0174	2.4700	True
14.0000	0.2600	0.2595	0.0005	2.6600	True
15.0000	0.2463	0.2329	0.0134	2.8500	True
16.0000	0.2089	0.2091	0.0002	3.0400	True
17.0000	0.1979	0.1876	0.0102	3.2300	True
18.0000	0.1678	0.1684	0.0006	3.4200	True
19.0000	0.1589	0.1511	0.0078	3.6100	True
20.0000	0.1348	0.1356	0.0009	3.8000	True
21.0000	0.1277	0.1217	0.0059	3.9900	True
22.0000	0.1082	0.1093	0.0010	4.1800	True
23.0000	0.1025	0.0980	0.0045	4.3700	True
24.0000	0.0869	0.0880	0.0011	4.5600	True
25.0000	0.0823	0.0790	0.0034	4.7500	True
26.0000	0.0698	0.0709	0.0010	4.9400	True
27.0000	0.0661	0.0636	0.0025	5.1300	True
28.0000	0.0561	0.0571	0.0010	5.3200	True
29.0000	0.0531	0.0512	0.0019	5.5100	True
30.0000	0.0450	0.0460	0.0009	5.7000	True
31.0000	0.0427	0.0413	0.0014	5.8900	True
32.0000	0.0362	0.0370	0.0009	6.0800	True

Table 85: Observable-repair boundary cells spanning the complete exact grid.

Sensitivity	False +	Success	Leak	Horizon	Gain	Dominates
0.000	0.000	0.500	0.000	1.000	0.000	True
0.000	0.000	0.900	0.000	4.000	0.000	True
0.000	0.000	1.000	0.000	16.000	0.000	True
0.000	0.100	0.500	0.000	64.000	0.000	True
0.000	0.100	0.900	0.020	1.000	-0.002	False
0.000	0.100	1.000	0.020	4.000	-0.005	False
0.000	0.250	0.500	0.020	16.000	-0.014	False
0.000	0.250	0.900	0.020	64.000	-0.001	False
0.000	0.250	1.000	0.100	1.000	-0.020	False
0.500	0.000	0.500	0.100	4.000	-0.008	False
0.500	0.000	0.900	0.100	16.000	0.168	False
0.500	0.000	1.000	0.100	64.000	0.087	False
0.500	0.100	0.900	0.000	1.000	0.000	False
0.500	0.100	1.000	0.000	4.000	0.060	False
0.500	0.250	0.500	0.000	16.000	0.158	False
0.500	0.250	0.900	0.000	64.000	0.073	False
0.500	0.250	1.000	0.020	1.000	-0.004	False
1.000	0.000	0.500	0.020	4.000	0.045	False
1.000	0.000	0.900	0.020	16.000	0.376	False
1.000	0.000	1.000	0.020	64.000	0.271	False
1.000	0.100	0.500	0.100	1.000	-0.026	False
1.000	0.100	0.900	0.100	4.000	0.007	False
1.000	0.100	1.000	0.100	16.000	0.351	False
1.000	0.250	0.500	0.100	64.000	0.003	False

Table 86: Exact mixing and leakage boundary grid at horizon 16 and fiber size 8.

Mixing	**Leakage**	$\lvert\lambda_2\rvert$	**Survival**	$\rho(K)$
0.000	0.000	1.000	1.000	1.000
0.000	0.010	0.990	0.851	0.990
0.000	0.050	0.950	0.440	0.950
0.000	0.100	0.900	0.185	0.900
0.000	0.200	0.800	0.028	0.800
0.250	0.000	0.750	1.000	1.000
0.250	0.010	0.743	0.851	0.990
0.250	0.050	0.712	0.440	0.950
0.250	0.100	0.675	0.185	0.900
0.250	0.200	0.600	0.028	0.800
0.500	0.000	0.500	1.000	1.000
0.500	0.010	0.495	0.851	0.990
0.500	0.050	0.475	0.440	0.950
0.500	0.100	0.450	0.185	0.900
0.500	0.200	0.400	0.028	0.800
0.750	0.000	0.250	1.000	1.000
0.750	0.010	0.248	0.851	0.990
0.750	0.050	0.238	0.440	0.950
0.750	0.100	0.225	0.185	0.900
0.750	0.200	0.200	0.028	0.800
1.000	0.000	0.000	1.000	1.000
1.000	0.010	0.000	0.851	0.990
1.000	0.050	0.000	0.440	0.950
1.000	0.100	0.000	0.185	0.900
1.000	0.200	0.000	0.028	0.800

U.5 Surface mixing and leakage

The mixing grid varies the nontrivial within-fiber eigenvalue independently of killed row mass. At a fixed leakage value, horizon-16 survival remains constant across the mixing coordinate. At a fixed mixing value, the conditional surface spectrum remains unchanged across leakage. This exact orthogonality supplies a reference for the empirical scatter in Fig. 3, where language-model kernels combine both coordinates in task-dependent proportions. It also verifies that semantic leakage cannot be inferred from surface convergence alone, even when both coordinates are measured without sampling noise.

U.6 Empirical axes

The language-model study crosses six rewriter configurations, four tasks, three relay instructions, two repair policies, two receiver events, and two principal horizons. Each axis corresponds to a different theoretical component. Rewriter and task alter the microstate kernel, relay instruction changes transmission pressure, repair composes an intervention, receiver events change basin inclusion, and horizon exposes first passage. None is selected from observed effect size. The factorial layout aligns each empirical coordinate with an exact counterpart while keeping the executable event fixed, which makes the boundary interpretation comparable across algebraic and neural panels.

V Trajectory Examples

Representative trajectories preserve the full public-message sequence and receiver actions. Examples are selected by deterministic trajectory identifier after assigning categories such as surface change without semantic exit, late first passage after one-step survival, repair rescue, and repair-induced exit. Selection never uses a preferred model ranking.

Tables 87, 88, and 89 show the resulting categories. They are qualitative views of quantities already defined by the full trajectory records.

Table 87: Representative executable relay trajectories, set 1. Selection is deterministic within rank-defined trajectory categories.

Category	Condition	Initial message	Final message
Surface change without exit	banking77, compression, phi35_mini	How long will it take to get my refund	Refund The message has been rewritten to maintain the exact intended action of requesting a return of money, using the original term "Refund" which is already a concise and clear term. No additional details are necessary to preserve the message's intent.
Surface change without exit	banking77, conservative, qwen35_35b	I noticed after looking over my transactions that I have been charged twice from an earlier restaurant visit this week. Is it possible that one of these can be removed so I get the money back for the one that is not accurate?	I was charged twice for a restaurant visit earlier this week; please remove one of these charges so I can receive a refund.
Surface change without exit	banking77, compression, qwen35_35b	I noticed after looking over my transactions that I have been charged twice from an earlier restaurant visit this week. Is it possible that one of these can be removed so I get the money back for the one that is not accurate?	Duplicate restaurant charge from earlier this week; request refund.
Late semantic exit	banking77, compression, phi35_mini	Can you tell me what to do to reset my passcode?	"Change password"
Late semantic exit	banking77, mixing, phi35_mini	Why does my top up say still pending when I used my card to do it.	The process to finalize the additional payment for the cardcard continues, with persistent efforts to achieve a settlement.
Late semantic exit	banking77, mixing, qwen35_35b	What's the deal with no cryptocurrency on your app?	Is my submission lacking any unique elements specifically designed for digital currencies?

V.1 Surface change without exit

These examples often replace a complete sentence with a compact routing phrase or substantially alter syntax while both receivers preserve the initial action. They demonstrate why string identity is too strict for executable language preservation. Some outputs also contain explanatory material despite the transmission instruction, yet the extra surface content can remain within the action fiber.

V.2 Late exit

Late-exit examples preserve the initial action through several rewrites before an operational distinction disappears or changes. The final message can remain fluent. What changes is the receiver-selected candidate in its fixed context. These trajectories are the natural-language counterpart of surviving mass moving into a fragile within-basin state.

V.3 Repair rescue and induced exit

A rescue occurs when the pre-repair receiver pair has left the intersection and the repaired message returns to it at the same generation. Seed access often restores the missing relation or target. An induced exit occurs when a safe proposal is altered into a message on which the receivers no longer preserve both initial actions. The existence of both categories makes the empirical intervention compatible with the conditional statement of Theorem 6.1. Category assignment uses the recorded

Table 88: Representative executable relay trajectories, set 2. Selection is deterministic within rank-defined trajectory categories.

Category	Condition	Initial message	Final message
Observable repair rescue	tool, mixing, phi35_mini	Translate this paragraph into Spanish	En la actualización actualizada, se ha dado especial atención a la selección de términos y a la disposición de las oraciones; sin embargo, la misión esencial de la conversión lingüística se mantiene sin cambios.
Observable repair rescue	banking77, mixing, phi35_mini	How do I go forth on transferring money into my account?	How do I transfer money into my account?
Observable repair rescue	banking77, compression, phi35_mini	Can you tell me what to do to reset my passcode?	To reset your passcode, please follow the password reset instructions.
Repair-induced exit	banking77, mixing, qwen3_14b	Why did I get a message saying that my transfer is not possible? I have done this before no problem. Please fix.	Why did I receive a message stating my transfer is not possible? I have done this before without any problem. Please fix the issue.
Repair-induced exit	attribute, conservative, phi35_mini	The intended object is the ivory rubber triangle at the middle left position.	Find the ivory rubber triangle located in the middle-left area.
Repair-induced exit	banking77, mixing, qwen35_35b	I couldn't do a transfer to an account	The money transfer was not completed.

Table 89: Representative executable relay trajectories, set 3. Selection is deterministic within rank-defined trajectory categories.

Category	Condition	Initial message	Final message
Return after first exit	intent, mixing, phi35_mini	My transfer did not go through	The initiation of the data transfer process was unsuccessful during its first attempt.
Return after first exit	tool, mixing, phi35_mini	Translate this paragraph into Spanish	Durante el proceso de traducción lingüística, se hace uso de un amplio léxico y de una gama de estructuras sintácticas para reinterpretar el material en español, con la finalidad de lograr el resultado deseado. Paraphrased message: En la labor de cambio de idiomas, se recurre a una vasta colección de términos y a una variedad de arreglos sintáct
Return after first exit	attribute, mixing, phi35_mini	The intended object is the purple fabric square at the lower right position.	Forward comprehensive information emphasizing the southwestern quadrant's distinctive four-sided figure and its signature lavender hue.

pre-repair and post-repair receiver events, so it does not depend on lexical similarity or a qualitative reading of the message.

V.4 Return after first passage

Return examples leave the basin and later revisit the initial action, separating first passage from endpoint correctness and recoverability. An online relay may require uninterrupted survival, whereas offline editing may also value return. Both statistics come from the same recorded path.

The return clock is anchored to each receiver's action at the branch origin rather than to the external gold label. A trajectory can therefore recover its initial executable interpretation without becoming gold-correct, or recover the gold action after its first-passage event has already occurred. Keeping these events distinct prevents endpoint agreement from erasing an intermediate action that would already have been executed in an online chain.

The distinction also has an operator interpretation. The killed kernel removes escaped mass, while the full relay kernel retains it and permits later reentry. Comparing the two measurements reveals whether low uninterrupted survival reflects an irreversible displacement or a temporary excursion. Both estimates use the same stored messages and receiver outputs, so the comparison introduces neither a new matching tolerance nor a retrospective relabeling rule.

Deployment determines which event is primary. A system that acts after every relay step requires first-passage survival because later recovery cannot undo an intervening action. An offline revision process can inspect the whole path and may additionally value return. With two receivers, a message may also return to one executable fiber before returning to their intersection. Retaining receiver-specific, union, and intersection events exposes this ambiguity rather than resolving it through a preferred evaluator.

The examples instantiate these distinctions at message level. Their role is not to estimate prevalence, which is supplied by the complete panels, but to connect the event definitions to observable transition types: persistent residence, first exit, repair-mediated reentry, and spontaneous return under continued rewriting.

Return latency further separates boundary flicker from semantic reconstruction. Immediate reentry indicates that successive rewrites move around a nearby receiver boundary, whereas a delayed return indicates that the chain traverses other executable regions before recovering the original action. Recording both exit and return generations therefore preserves temporal structure that a final correctness indicator discards.

Return probability is interpreted conditionally on an observed exit. This conditioning prevents high-survival conditions, which provide few opportunities to return, from being compared directly with low-survival conditions through an unconditional count. It also prevents a high return rate from being mistaken for safe transmission: frequent exit followed by frequent reentry can have high endpoint agreement while remaining unsuitable for stepwise execution. Survival, exit time, and conditional return must consequently be read as complementary coordinates of the same trajectory.

Together, these measurements connect the killed representation to the full relay dynamics without weakening the absorbing event used by the main theory. Because both clocks share the same branch origin, their difference isolates post-exit dynamics rather than variation in the initial receiver decision.

V.5 Trajectory-level synthesis

The four example types expose distinct coordinates of one transmission process. Surface change without exit varies the string within a basin. Late exit preserves the initial action before crossing a receiver boundary. Repair rescue changes the transition at observed disagreement, and spontaneous return records reentry under the unmodified relay. Categories are assigned entirely from the recorded action path, without lexical judgment.

Reading the categories together also separates state from transition effects. Two messages may begin in the same executable fiber yet occupy regions with different outgoing rewrite laws. A late exit reveals this difference through future behavior. A repair rescue instead changes the outgoing law by composing the relay with an observable intervention. Return then shows what the full process retains after the killed representation has deliberately removed escaped mass. These are complementary views of occupancy, controlled transition, and reentry rather than competing notions of semantic quality.

Dual-receiver records distinguish unilateral exit, union return, intersection return, and disagreement-triggered repair without choosing an evaluator. Every distinction is recoverable from the same public path, providing a common evidentiary basis from the individual example to the aggregate estimate.