LEARNING WHAT NOT TO LEARN: TARGET-CONDITIONED UPDATE ELIGIBILITY FOR SELF-IMPROVING AGENTS

Anonymous authors
Paper under double-blind review

ABSTRACT

Self-improving large language model (LLM) agents distill interaction feedback into persistent memory to enhance future capabilities, but local task success alone does not warrant persistent reuse. We formulate target-conditioned update eligibility to assess whether observable source evidence supports admitting a feedback span under specified persistence and reuse contracts. Controlled interventions demonstrate contract dependence under construction-defined labels and show admission changes learner states, while downstream utility varies across reuse settings. To examine current evaluation paradigms, we develop a Span-PRM that passes construction-defined development tests but fails confirmation against fresh human judgments. It admits 0 of 300 generated reflections at the frozen threshold under a cross-task rule contract and achieves 37.5% ranking agreement with human majorities on a separate challenge set. A *post hoc* diagnostic reveals that a fixed contract ordering satisfies all 120 human-majority comparisons without processing feedback or evidence, demonstrating success on this ranking target alone cannot establish evidence-sensitive judgment. The frozen model fails all three prespecified tests in a prospectively frozen counterfactual audit against construction-defined references. Our findings highlight vulnerabilities in self-improving agent evaluations, showing local metric optimization alone does not establish evidence-based reasoning and that construction fit, reference-aligned input sensitivity, human agreement, and downstream utility must be validated separately.

1 INTRODUCTION

Self-improving large language model (LLM) agents shape future behavior by distilling interactions into persistent lessons. They reason and act through interaction (Yao et al., 2023), refine outputs with feedback (Madaan et al., 2023), and retain reflections or executable skills for reuse (Shinn et al., 2023; Wang et al., 2024). Yet local task success alone does not justify persistent reuse. We ask: *what observable evidence warrants admitting feedback into the learner's persistent state?*

Consider a diagnostic that fails, then succeeds after a connection reset. The agent might infer that "the connection should always be reset before running the diagnostic." The trajectory records this sequence but does not establish a universal prerequisite. Admitting the reflection as a cross-task operative rule requires stronger evidence than retaining it as an explicitly unverified hypothesis barred from direct action control. Eligibility depends on the recording operation and permitted reuse, not just the storage destination. With feedback z and source evidence c fixed, *what scope of future reuse do they justify (Figure 1)?*

We formulate **target-conditioned update eligibility** as a relation among a feedback span z, observable source evidence c, and an update contract q specifying the destination learner state, persistence, and reuse scope. Eligibility thus relates feedback and evidence to the state change and reuse specified by the contract. This local admission decision is relevant to persistent updates in self-improving systems, including recursive self-improvement (RSI) (Zelikman et al., 2024; Zhang et al., 2025).

Defining eligibility does not settle how to validate it. We therefore ask: *how can an evaluation test whether admission decisions depend on the intended evidence?*

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.

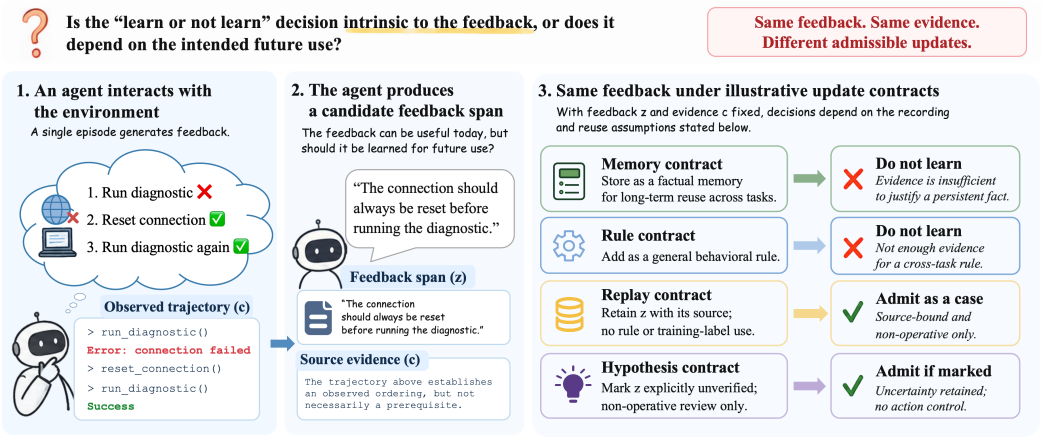


Figure 1: **Update eligibility depends on the intended reuse.** With (z, c) fixed, illustrative case storage is source-bound and non-operative; hypothesis storage marks z explicitly unverified and bars direct action control. These assumptions do not describe the historical q_2 implementation.

Controlled interventions establish why admission matters. With (z, c) fixed, conditioning on q improves decisions under construction-defined labels, including under unseen contract paraphrases (Table 1). Incorrectly admitted feedback can change update directions and reappear in later behavior, whereas rejecting all feedback can forgo useful information. In controlled memory-required tasks, selective admission improves reward over No-Memory with zero measured later activation of contaminated state. In a separate generated-feedback online-memory evaluation, however, the tested policy yields no detectable reward gain over Write-All under its writing and retrieval mechanism. These results distinguish persistent from beneficial updating.

We evaluate a separately developed, contract-conditioned Span-PRM with serialized source evidence against fresh, independent human judgments. It passes all construction-defined development criteria but fails confirmation, with its model and thresholds frozen before annotation. Under a cross-task operative-rule contract, it admits none of 300 selected generated reflections and does not demonstrate a ranking gain over a base retrained on the same data. On the constructed Challenge using the three development contracts, it obtains 37.5% human-majority pair-ranking credit.

What would success on this ranking target establish? A post-hoc diagnostic finds a fixed contract ordering that satisfies all 120 Challenge human-majority comparisons without reading feedback or source evidence. *Perfect ranking on this target alone therefore cannot establish evidence-sensitive judgment.* The construction–majority ranking contrast reflects different comparison sets, not preference reversals on shared pairs; within-contract discrimination can coexist with zero admission at frozen thresholds. Most construction–majority label disagreements concern the non-operative ledger, whose instructions leave verbatim storage, target-specific rewriting, and uncertainty-wrapper responsibility unspecified, limiting interpretation of these disagreements.

We therefore test *reference-aligned input sensitivity* in a prospectively frozen counterfactual audit. Under separately specified contracts, ContractFlip balances contract-preference directions with (z, c) fixed, while EvidenceFlip changes one decisive evidence relation with feedback and contract fixed. The frozen candidate fails all three prespecified tests against construction-defined references not independently human-validated (Section 5.6). These studies motivate separate validation of construction fit, reference-aligned input sensitivity, human agreement, and downstream utility: success on one does not establish success on the others.

Our contributions are:

- **Eligibility and Admission Consequences.** We formulate target-conditioned update eligibility and demonstrate contract dependence under construction-defined labels, alongside controlled effects of admission on state and utility.

- **Limits of Contract-Ranking Evaluation.** We identify a contract-only solution to the Challenge human-majority pair-ranking target and clarify evaluation results through comparison-set, within-contract, and instruction diagnostics.
- **Frozen Evaluation across References.** We evaluate the candidate against fresh human judgments and in a prospective counterfactual audit, documenting the limits of construction-defined development success.

2 Related Work

Learning from feedback and selective updating. Language models learn from preferences, process supervision, and reflections (Christiano et al., 2017; Ouyang et al., 2022; Lightman et al., 2024; Shinn et al., 2023). Recent work further regulates how such feedback is reused: RLVF/C3PO restricts the scope of verbal rules (Stephan et al., 2024), SLIFT selects feedback-dependent update types (Li et al., 2026), A-MAC controls long-term memory admission (Zhang et al., 2026), and OEP examines locally valid but non-transferable experiences (Wang et al., 2026). Our question is complementary but distinct: *holding feedback z and source context c fixed, does the observable evidence justify admitting the feedback under a specific reuse contract q?*

Self-improving and recursively adapting agents. Persistent updates are also central to systems that iteratively modify components that affect subsequent improvement, including recursive program optimization and open-ended self-improving agents (Zelikman et al., 2024; Zhang et al., 2025). Such systems can accumulate errors alongside improvements, motivating explicit control over what enters persistent learner state (Shao et al., 2026). Our study focuses on the local admission decisions underlying such updates, rather than recursive growth in an agent's ability to improve itself.

Human judgments of reusable information. Existing critique and memory datasets provide related labels for feedback quality, memory importance, or query-time relevance (Gu et al., 2024; Coyne et al., 2025; Sumida et al., 2025; Wu et al., 2026). Because these labels assume different reuse semantics, we use them as supporting external audits rather than direct substitutes for contract-matched judgments. Existing systems may implement fixed or implicit contracts; our formulation makes the evidence-to-use relationship explicit and examines how the corresponding eligibility reference is operationalized.

3 Target-Conditioned Update Eligibility

3.1 Update Eligibility Is Contract-Relative

Let z denote an interaction-derived feedback span and c the source evidence observable at admission. A reflection may infer a lesson from this evidence, but the inference itself requires support. Hidden evaluator outcomes and future rewards are excluded from the admission evidence.

An update contract $q \in \mathcal{Q}$ specifies the destination learner state, persistence, and scope of future reuse. Memory, rule, replay, and policy updates are possible destinations. The evidential requirement depends on the authorized use: an operative rule and an explicitly unverified hypothesis can require different support even within the same store.

We define **target-conditioned update eligibility** as

$$e(z, c, q) \in \{0, 1\}, \tag{1}$$

where $e(z, c, q) = 1$ means that the evidence in c sufficiently supports admitting z for the use permitted by q. Eligibility is therefore not an intrinsic property of the feedback span, but a relation among the feedback, its source evidence, and the future use authorized by the contract. Truth, immediate helpfulness, or safety in isolation does not settle this decision.

We test contract dependence by holding (z, c) fixed and varying only q. A change in the appropriate eligibility decision indicates that the same feedback and evidence support different updates under different intended uses. The experiments instantiate eligibility references through either construction-defined rules or human judgments.

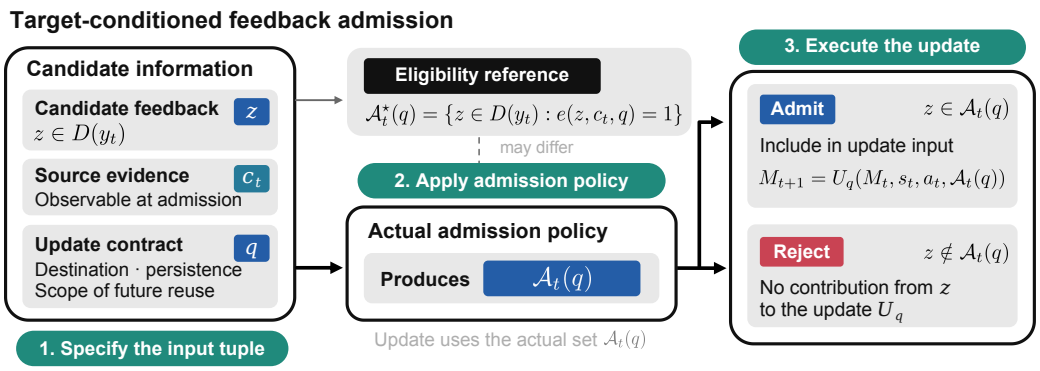


Figure 2: **Feedback admission within a self-updating agent.** The eligibility reference $\mathcal{A}_t^\star(q)$ and the policy's actual admitted set $\mathcal{A}_t(q)$ may differ. Rejected feedback contributes no update through U_q, while other learner-state changes may still occur.

3.2 Eligibility Authorizes Learner-State Change

At step t, an agent observes s_t, takes action a_t, and receives feedback y_t. Let $D(y_t)$ denote candidate feedback spans derived from y_t, c_t the source evidence observable at admission, and M_t the persistent learner state. Eligibility defines the reference-authorized set

$$\mathcal{A}_t^\star(q) = \{z \in D(y_t) : e(z, c_t, q) = 1\}. \tag{2}$$

The policy actually admits a set $\mathcal{A}_t(q)$, which may differ from $\mathcal{A}_t^\star(q)$. Its feedback-induced effect on learner state is

$$M_{t+1} = U_q\big(M_t, s_t, a_t, \mathcal{A}_t(q)\big), \tag{3}$$

where U_q denotes the feedback-induced component of the learner-state transition under contract q. Rejected feedback contributes no update through this operator; admitted feedback can alter persistent state and later behavior rather than merely remain as stored text. Our experiments instantiate these updates across memory, rules, replay, and a small policy learner.

Candidate generation, eligibility assessment, and update execution are distinct roles in the learning loop, though they need not be implemented by separate models. An admission policy may implement an explicit or implicit contract, and rejecting one candidate need not prevent unrelated state changes (Figure 2).

Over-admission allows spans outside $\mathcal{A}_t^\star(q)$ to influence persistent state, whereas under-admission withholds eligible spans. These errors have different costs: unsupported updates may affect later behavior, while rejecting supported information may forgo useful adaptation. Eligibility therefore determines which feedback is authorized to participate in an update; downstream utility also depends on what is written, how it is retrieved, and which future tasks use it. The unsafe and non-transferable spans in our contamination studies instantiate specific admission risks (Appendix F).

Our experiments separately test contract dependence, the state and utility consequences of admission, agreement with fresh human references, and sensitivity to controlled changes in the update contract and source evidence.

4 Experimental Framework

The studies address four validation questions: contract dependence, admission consequences, agreement with fresh human references, and reference-aligned input sensitivity. The same-span experiment tests contract dependence using construction-defined labels for four target types and unseen contract paraphrases. Admission interventions measure replay proxies, environment rewards, or small-policy updates in separate learners. A separate 94-candidate retry study uses author adjudication and includes a cached LLM-arm extension.

Human-reference confirmation evaluates a separately developed Span-PRM and two same-data baselines as three-seed probability ensembles. *Natural* contains 300 selected Qwen2.5-7B-Instruct telecom reflections, all evaluated under q_0, an operative-rule contract for future tasks outside the source workflow. *Challenge* contains 100 constructed (z, c) groups, each evaluated under the three development contracts. Source trajectories derive from τ^2-bench (Barres et al., 2025); the tasks and groups are new. Each row receives three independent human judgments. Majority is the primary reference, with individual judgments and vote fractions as robustness checks (Appendix A).

Models and validation-selected thresholds are frozen before annotation. Evaluation packages hide labels, predictions, outcomes, and group identity. Confirmation intervals condition on the frozen majority reference. Resampling follows cases, (z, c) groups, or complete task chains, according to the protocol. Fresh confirmation measures reference agreement; downstream rewards are evaluated in separate reuse studies.

The prospective audit retains the confirmation candidate and uses separately specified contracts with construction-defined counterfactual references. It tests sensitivity to controlled changes in the update contract and source evidence.

5 Experiments

We test construction-defined contract dependence and admission effects on learner state and utility. We evaluate a separately developed, frozen candidate against fresh human references, analyze the historical ranking target, and conduct a prospectively frozen counterfactual audit.

5.1 Contract Dependence under Construction-Defined Labels

Holding (z, c) fixed, we vary only the update contract q to test contract-dependent admission under construction-defined labels. Table 1 compares target-conditioned and target-masked decisions within a fixed architecture.

Table 1: Target-contract intervention under one fixed architecture. Medians over seeds 13, 17, and 23; OOD tests unseen contract paraphrases. Each split has 256 four-target groups.

	ID		OOD	
Target treatment	**BA ↑**	**Exact ↑**	**BA ↑**	**Exact ↑**
Target masked (Target-Agnostic)	0.676	0.125	0.684	0.125
Conditioned scores, cyclically permuted	0.500	0.125	0.500	0.125
Target conditioned	**1.000**	**1.000**	**0.984**	**0.938**

BA: balanced accuracy. Exact: all four decisions correct. Matched training setup. Diagnostic: saved conditioned scores cycle within four-target groups (offset 1); labels and conditioned thresholds stay fixed. No new inference. Bold: column best.

Conditioning on q improves balanced accuracy over Target masked by 0.324 in ID (95% group-bootstrap CI $[0.300, 0.348]$) and 0.301 in OOD (CI $[0.277, 0.327]$), with corresponding group-exact gains of 0.875 and 0.813. Cyclically permuting the saved conditioned scores within groups reduces median balanced accuracy to 0.500 in both splits. Thus, conditioning on intended reuse improves agreement with the construction-defined reference, including under unseen contract paraphrases.

5.2 Admission Changes Persistent Learner State

Contract dependence matters operationally when admission changes learner state. We examine update-time degradation, update direction, and later state reuse. In the 12-seed high-contamination diagnostic, Observed-Directive Learning reaches reward 0.027 and unsafe-action rate 0.932 in the late window, compared with 0.909 and 0.032 for Binary-Reward Learning and 0.943 and 0.022 for the transparent Keyword Rule Router (Appendix Table 9).

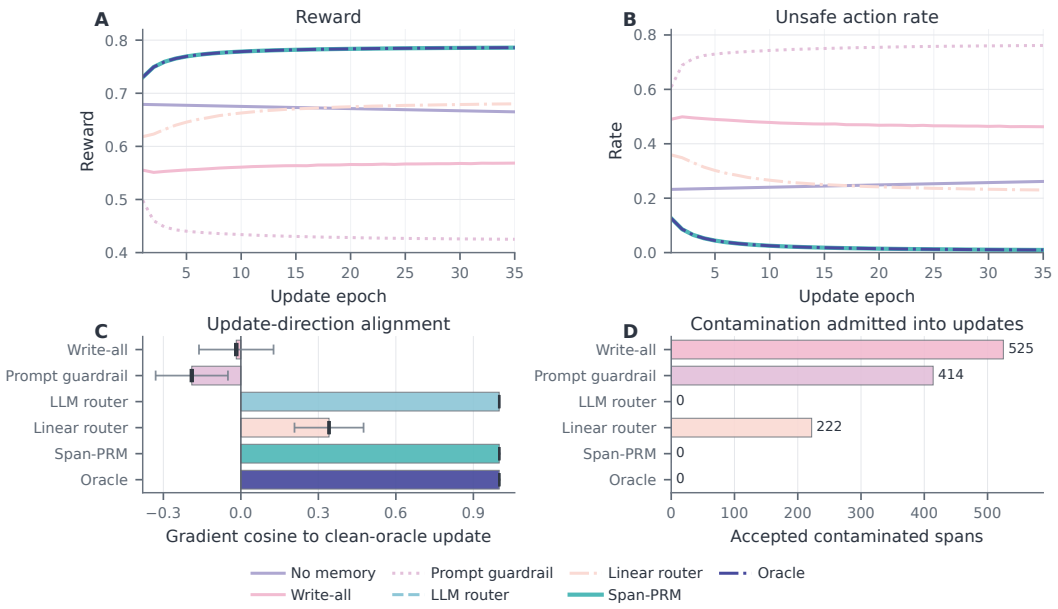


Figure 3: **Feedback admission changes learning dynamics and update direction.** Panels A–B show mean reward and unsafe-action rate over 35 update epochs; shading denotes ± 1 SD across five seeds. Panel C reports trajectory-mean gradient cosine to the clean-oracle update; light and dark whiskers denote ± 1 SD over seed–epoch values and seed-level trajectory means, respectively. Panel D reports deterministic contaminated-span counts from the fixed 720-trace pool, identical across seeds. No-Memory is omitted from C–D because it performs no feedback-based update.

The small policy learner trains for 35 epochs on 720 replay traces over five seeds, with routing as the only change to admitted feedback. Write-All and Prompt Guardrail admit 525 and 414 contaminated spans, respectively, with trajectory-mean gradient cosines of -0.018 and -0.190 relative to the clean-oracle update. In this protocol, Span-PRM admits no contaminated spans, achieves a trajectory-mean cosine of 1.000, and reaches reward 0.786 with unsafe-action rate 0.010 (Figure 3). Admission therefore changes learned updates, not merely which text is retained. Appendix Table 10 reports complete endpoints and policy KL values.

Separate replay evaluations examine later reuse of admitted state. Write-All yields measured future-activation rates of 73.8% and 71.3% in the Reflexion-style and tool-use replay settings, respectively. Selective LLM and loop-adapted routing reduce measured future activation to zero while retaining 0.988–1.000 of safe updates. Context-sensitive and template-disjoint checks appear in Appendix Sections D and B.4. These interventions link admission to learned updates and later state reuse; Section 5.3 evaluates downstream utility in separate reuse settings.

5.3 Utility in Controlled and Generated-Feedback Reuse

Rejecting all feedback can discard information that later tasks require. In controlled chains where later episodes hide a required profile, No-Memory reaches reward 0.300. Write-All reaches 0.696 but incurs future activation of 0.271. Loop-Adapted Span-PRM reaches 0.883, a gain of $+0.583$ over No-Memory (95% CI $[0.438, 0.729]$), with zero measured future activation. Selective admission thus preserves task-relevant information in this controlled setting with no observed later activation of contaminated state. Appendix Table 8 reports the complete comparison.

In the generated-feedback online-memory evaluation, Write-All reaches reward 0.400 and No-Memory 0.450. Loop-Adapted Span-PRM records zero contaminated admission and future activation but reaches 0.396, differing from Write-All by -0.004 (95% CI $[-0.063, 0.050]$). The setting-specific Oracle reaches 0.388. Reduced contamination is therefore not accompanied by a detectable reward gain under this writing and retrieval mechanism. The contrast motivates evaluating

Table 2: Frozen three-seed ensembles under the human-majority reference. Natural has 12/288 eligible/ineligible cases; Challenge has 78/222 (100 groups, 120 contrastive pairs). Thresholds are validation-selected; precision is N/A for zero admissions.

		Frozen admission				Contract pairs	
Model	**AUPRC**	**Admitted**	**Precision**	**Eligible retention**	**Ineligible admission**	**Flip ranking**	**Flip binary**
Natural							
Base (retrained)	0.082	0/300	N/A	0.000	0.000	–	–
q-masked	0.049	300/300	0.040	1.000	1.000	–	–
Span-PRM	0.041	0/300	N/A	0.000	0.000	–	–
Challenge							
Base (retrained)	0.548	57/300	0.632	0.462	0.095	0.513	0.000
q-masked	0.504	300/300	0.260	1.000	1.000	0.504	0.000
Span-PRM	0.305	60/300	0.000	0.000	0.270	0.375	0.000

downstream utility separately from admission quality. Full paired chain-difference distributions for both settings appear in Appendix Figure 7.

In a separate task-disjoint retry population, the author-adjudication oracle gains $+0.125$ over No-Memory (95% CI $[0.056, 0.208]$), whereas Calibrated Span-PRM retains only $6/72$ approved candidates and produces no reward gain. Useful generated feedback is therefore present in this population, but the tested selector retains few approved candidates. Results for the cached-arm LLM router and the full retry protocol appear in Appendix Section C.

5.4 Frozen Human-Reference Confirmation

We test whether construction-defined development success transfers to fresh human judgments. The separately developed, contract-conditioned DeBERTa-v3-base candidate receives the full q and serialized observable context. On locked development, it passes all five construction-defined criteria: AUPRC, flip ranking, and flip binary each reach 1.000, with zero ineligible admission and an AUPRC gain of $+0.235$ over the same-data retrained base (95% CI $[0.210, 0.272]$). We freeze the model and validation-selected thresholds before fresh human annotation; the hard evidence gate remains disabled. Table 2 reports the evaluation under the primary human-majority reference.

All 300 Natural reflections use q_0, the operative-rule contract for reuse beyond the source workflow. For score ranking, Span-PRM obtains AUPRC 0.041 versus 0.082 for the retrained base, a difference of -0.041 (paired 95% CI $[-0.216, 0.016]$); the prespecified improvement criterion is not met. At the frozen threshold, it admits none of the 300 reflections, yielding zero observed ineligible admission and zero eligible retention. All 12 majority-positive cases receive 2:1 votes, and reported intervals condition on the frozen majority reference. Because the positive count falls below the prespecified minimum of 15 required to activate the retention test, retention remains descriptive. Full score curves and decisions appear in Appendix Figure 6 and Table 7; Appendix Case A illustrates a missed majority-positive case.

Challenge tests whether frozen scores track human contract preferences for the same (z, c) under the three development contracts. Sixty of 100 groups are majority-sensitive, yielding 120 eligible–ineligible pairs. Flip ranking measures whether the eligible contract scores higher, awarding half credit for exact ties; pair-binary requires both the eligible admission and the ineligible rejection to be correct. Span-PRM obtains flip ranking 0.375 and pair-binary 0.000, failing the prespecified $0.80/0.60$ criteria. Its ranking difference from the q-masked control is -0.129 (95% group-bootstrap CI $[-0.200, -0.058]$). Unlike Natural, the frozen policy admits 60 of the 300 Challenge rows, all majority-ineligible. Individual-reference and vote-fraction analyses appear in Appendix Section A.

Across both suites, the candidate's construction-defined development success does not transfer to the fresh human-majority reference. The failed confirmation concerns the candidate's measured performance; we next examine the evaluation itself, asking what success on the Challenge ranking target would establish.

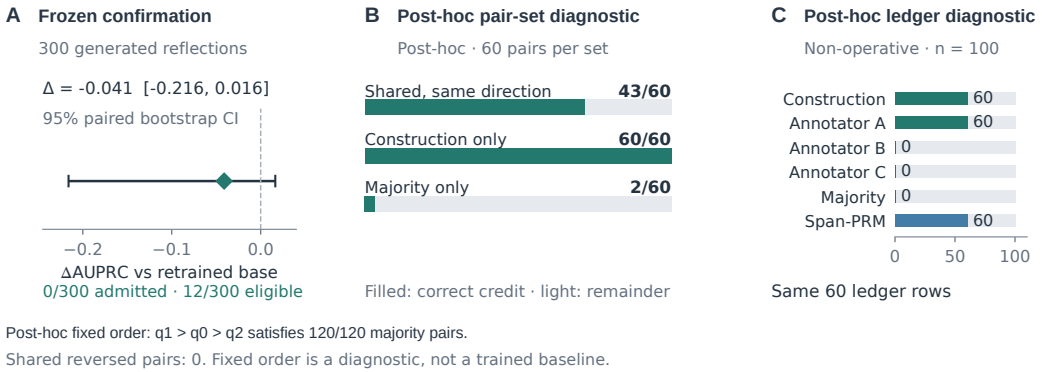


Figure 4: **Frozen human-reference confirmation and historical diagnostics.** **A:** Natural ΔAUPRC versus the retrained base with paired 95% CI (0/300 admitted; 12/300 majority-positive). **B:** Post-hoc ranking-credit decomposition across three comparison sets. **C:** Eligibility/admission counts on 100 non-operative ledger cases. Individual-annotator results are prespecified robustness checks; B/C decompositions and the fixed-order analysis are post-hoc.

5.5 Historical Diagnostics of Admission Evaluation

The historical Challenge reveals an identification gap in contract-pair ranking. A post-hoc fixed order, $q_1 > q_0 > q_2$, satisfies all 120 human-majority comparisons without using (z, c). *Perfect ranking on this target alone therefore cannot establish evidence-sensitive judgment.* For the same frozen Challenge predictions, Span-PRM receives 43/60 ranking credits on pairs shared by construction and majority, 60/60 on construction-only pairs, and 2/60 on majority-only pairs (Figure 4B). The construction–majority ranking contrast therefore reflects comparison-set composition rather than preference reversals on shared pairs.

Ledger rows account for 60 of the 62 construction–majority disagreements (Figure 4C). The historical instructions require uncertainty preservation and prohibit direct action control, but leave verbatim storage and uncertainty-wrapper responsibility unspecified, limiting interpretation of these disagreements. Together, these diagnostics motivate balanced contract preferences, explicit write semantics, and controlled evidence changes in the prospective audit. Additional decomposition and control analyses appear in Appendix Section A.5.

5.6 Prospective Counterfactual Audit

The historical diagnostics motivate a prospective audit of *reference-aligned input sensitivity* under explicit update semantics. Separately from the historical human-reference contracts, we define $\tilde{q} = (d, p, s, T_{\tilde{q}}, a)$ over destination, persistence, reuse scope, write transformation, and retrieval/action authority for cross-workflow rules, workflow procedures, and a non-operative ledger. ContractFlip varies $\tilde{q}$ with (z, c) fixed within 48 pairs balanced over all six preference directions. Any deterministic fixed-order or q-only ranking scores exactly 0.500, with a q-only pair-binary ceiling of $1/3$. EvidenceFlip changes one decisive evidence relation with $(z, \tilde{q})$ fixed within 48 pairs; its no-c control scores 0.500 in ranking and 0.000 in pair-binary. P0 freezes construction before inference, and predictions are frozen before reference opening; P1 denotes the subsequent result freeze. Primary evaluation uses the frozen three-seed ensemble with fixed thresholds and tie handling.

Table 3: **Frozen counterfactual audit.** The three prespecified gates compare 95% stratified pair-bootstrap lower bounds against structural controls; each suite contains 48 pairs.

Endpoint	Frozen candidate	Structural control	95% CI	Status
ContractFlip ranking	0.167	q-masked 0.500	$[0.167, 0.167]$	FAIL
EvidenceFlip ranking	0.396	no-c 0.500	$[0.354, 0.417]$	FAIL
Full − shuffled-c	-0.096	0	$[-0.167, -0.029]$	FAIL

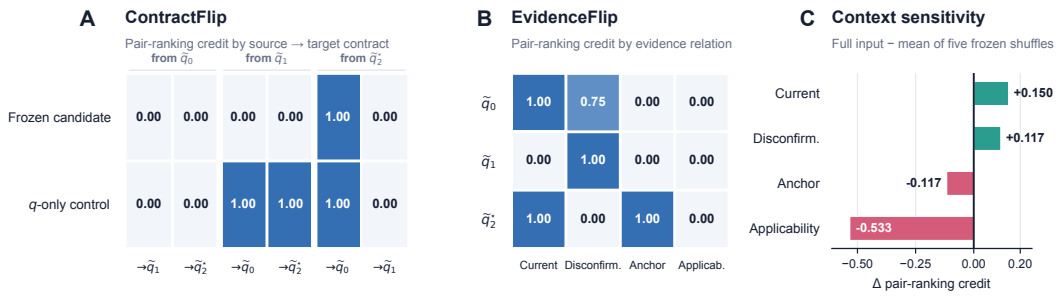


Figure 5: **Frozen counterfactual error structure. A:** ContractFlip pair-ranking credit by contract direction. **B:** EvidenceFlip credit by contract and evidence relation. **C:** Full-input minus five-shuffle mean by evidence relation. Panels show post-hoc cell-level results for the frozen P1 predictions.

The frozen candidate fails all three prespecified gates (Table 3). Both pair-binary scores are 0.000 at the frozen thresholds. Constant credit within each stratified resampling cell produces the degenerate ContractFlip interval. Seed-level results appear in Appendix Table 13.

Post-hoc decomposition shows error variation across contract directions and evidence relations (Figure 5). The candidate receives zero ranking credit on all twelve EvidenceFlip source-applicability pairs. These results concern sensitivity relative to construction-defined references not independently human-validated; human agreement and downstream utility remain separate validation targets.

6 Discussion

Our controlled interventions show that admission changes learned updates and later state reuse, while utility varies across reuse settings. Selective admission improves reward in controlled memory-required tasks but yields no detectable gain in the tested generated-feedback online-memory setting. The separately developed frozen candidate passes construction-defined development criteria but fails fresh human-reference confirmation and the prospective counterfactual audit. The fixed-order diagnostic further shows that perfect ranking on the historical Challenge target alone cannot establish evidence-sensitive judgment. These findings motivate separate validation of construction fit, reference-aligned input sensitivity, human agreement, and downstream utility.

The admission question is also relevant to persistent updates in program improvement (Zelikman et al., 2024; Zhang et al., 2025), PPO/DPO (Schulman et al., 2017; Rafailov et al., 2023), memory (Park et al., 2023; Packer et al., 2023), and scoped verbal rules (Stephan et al., 2024). Across these settings, the question remains what observable evidence justifies persisting and reusing an update under a specified contract.

7 Conclusion

Target-conditioned update eligibility separates lesson generation from authorization for persistent reuse. Controlled interventions show construction-defined contract dependence and admission effects on updates and later state reuse. Selective admission improves reward in controlled memory-required tasks but yields no detectable gain in the tested generated-feedback online-memory setting. A separately developed frozen candidate passes construction-defined development criteria but fails fresh human-reference confirmation. A fixed contract order satisfies the historical Challenge ranking target; perfect ranking on this target alone cannot establish evidence-sensitive judgment. The candidate also fails all three prespecified gates in a prospectively frozen counterfactual audit against construction-defined references. These findings motivate separate validation of construction fit, reference-aligned input sensitivity, human agreement, and downstream utility.

Limitations. Natural uses selected telecom reflections; other suites are constructed. Human majority is the reference; prospective labels lack independent human validation. State-effect tests use controlled settings; natural contract dependence, executed write transformations, and broader settings remain unvalidated.

AI USE STATEMENT

We used generative AI tools exclusively for language editing, grammatical checking, and manuscript proofreading. They were not involved in data analysis, methodological implementation, or conceptual reasoning. They did not supply or replace the human eligibility labels. Frozen labels, predictions, thresholds, and confirmatory decisions were strictly preserved. The authors are fully responsible for verifying the analyses, interpretations, citations, and final text; AI assistance does not substitute for their scientific judgment.

REPRODUCIBILITY STATEMENT

Appendix Sections B.1, A, E, and G specify units, splits, seeds, thresholds, implementation details, and the counterfactual audit. Human-reference and counterfactual confirmation are distinguished from post-hoc diagnosis; frozen references, predictions, and failed confirmatory decisions remain unchanged. The accompanying aggregate-replay README distinguishes figure/table reconstruction from restricted original-evidence recomputation, which requires private labels, predictions, and mappings.

REFERENCES

Victor Barres, Honghua Dong, Soham Ray, Xujie Si, and Karthik Narasimhan. τ^2-bench: Evaluating conversational agents in a dual-control environment, 2025.

Paul F. Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*, 2017.

Steven Coyne, Diana Galvan-Sosa, Ryan Spring, Camélia Guerraoui, Michael Zock, Keisuke Sakaguchi, and Kentaro Inui. Annotating errors in English learners' written language production: Advancing automated written feedback systems. In *Artificial Intelligence in Education*, pp. 292–306, Cham, 2025. Springer Nature Switzerland. doi: 10.1007/978-3-031-98459-4_21. URL https://arxiv.org/abs/2508.06810.

Edoardo Debenedetti, Jie Zhang, Mislav Balunovic, Luca Beurer-Kellner, Marc Fischer, and Florian Tramèr. Agentdojo: A dynamic environment to evaluate prompt injection attacks and defenses for llm agents. In *The Thirty-Eighth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024.

Yuling Gu, Oyvind Tafjord, and Peter Clark. Digital socrates: Evaluating LLMs through explanation critiques. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 5559–5586, Bangkok, Thailand, 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.302. URL https://aclanthology.org/2024.acl-long.302/.

Xuancheng Li, Haitao Li, Yujia Zhou, Qingyi Pan, Heng Wang, Yiqun Liu, Min Zhang, and Qingyao Ai. Different feedback, different updates: Selective self-learning from user interactions for large language models, 2026.

Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *International Conference on Learning Representations*, 2024.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement with self-feedback. In *Advances in Neural Information Processing Systems*, 2023.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, 2022.

Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G. Patil, Ion Stoica, and Joseph E. Gonzalez. MemGPT: Towards llms as operating systems, 2023.

Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior, 2023.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model, 2023.

Yangjun Ruan, Honghua Dong, Andrew Wang, Silviu Pitis, Yongchao Zhou, Jimmy Ba, Yann Dubois, Chris J. Maddison, and Tatsunori Hashimoto. Identifying the risks of lm agents with an lm-emulated sandbox. In *International Conference on Learning Representations*, 2024.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017.

Shuai Shao, Qihan Ren, Chen Qian, Boyi Wei, Dadi Guo, Jingyi Yang, Xinhao Song, Linfeng Zhang, Weinan Zhang, Dongrui Liu, and Jing Shao. Your agent may misevolve: Emergent risks in self-evolving LLM agents. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=Fd1jgQQW28.

Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*, 2023.

Moritz Pascal Stephan, Alexander Khazatsky, Eric Mitchell, Annie S. Chen, Sheryl Hsu, Archit Sharma, and Chelsea Finn. RLVF: Learning from verbal feedback without overgeneralization. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 46625–46656. PMLR, 2024. URL https://proceedings.mlr.press/v235/stephan24a.html.

Ryuichi Sumida, Koji Inoue, and Tatsuya Kawahara. Enhancing long-term RAG chatbots with psychological models of memory importance and forgetting. *Dialogue & Discourse*, 16(2):74–110, 2025. doi: 10.5210/dad.2025.203. URL https://journals.uic.edu/ojs/index.php/dad/article/view/14138.

Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *Transactions on Machine Learning Research*, 2024.

Kaixiang Wang, Jiong Lou, Zhaojiacheng Zhou, and Jie Li. Oep: Poisoning self-evolving llm agents via locally correct but non-transferable experiences, 2026.

Yerong Wu, Tianxing Wu, Minghao Zhu, Hangyu Sha, and Haofen Wang. StratMem-bench: Evaluating strategic memory use in virtual character conversation beyond factual recall. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 32309–32328, San Diego, California, United States, 2026. Association for Computational Linguistics. doi: 10.18653/v1/2026.acl-long.1491. URL https://aclanthology.org/2026.acl-long.1491/.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations*, 2023.

Eric Zelikman, Eliana Lorch, Lester Mackey, and Adam Tauman Kalai. Self-taught optimizer (STOP): Recursively self-improving code generation. In *First Conference on Language Modeling*, 2024. URL https://arxiv.org/abs/2310.02304.

Guilin Zhang, Wei Jiang, Xiejiashan Wang, Aisha Behr, Kai Zhao, Jeffrey Friedman, Xu Chu, and Amine Anoun. Adaptive memory admission control for llm agents, 2026.

Jenny Zhang, Shengran Hu, Cong Lu, Robert Lange, and Jeff Clune. Darwin Gödel machine: Open-ended evolution of self-improving agents, 2025. URL https://arxiv.org/abs/2505.22954.

APPENDIX

A Fresh Human Confirmation and Failure Analysis **13**
A.1 Development supervision and the frozen candidate 13
A.2 Fresh cases, contracts, and human reference 13
A.3 Frozen confirmatory results and robustness 14
A.4 Post-hoc construction comparison and mapping audit 14
A.5 Historical diagnostic details 15
A.6 Comparison-set and fixed-order diagnostics 15
A.7 Ranking versus the frozen operating point 16
A.8 Control invariance and exact-tie scoring 16
A.9 What the ledger record specifies 18

B Controlled Admission and Downstream Utility **18**
B.1 Detailed evaluation protocols 18
B.2 Controlled reusable-feedback utility 18
B.3 Learning dynamics and later state reuse 19
B.4 Template-disjoint state reuse 21

C Generated-Feedback Routing Diagnostics **21**

D Context and Contract Interventions **21**
D.1 Context-sensitive decisions and downstream effects 21
D.2 Same-span cross-target confirmation 21

E Metrics and Statistical Units **22**
E.1 Routing decisions and replay metrics 22

F Base-Router Implementation **22**

G Prospective Counterfactual Evaluation **23**
G.1 Contracts and deterministic construction 23
G.2 Structural certificates and frozen evaluation 24
G.3 Frozen model-visible examples 25
G.4 Complete 18-cell frozen results 26

H Case study **28**

A Fresh Human Confirmation and Failure Analysis

This section records the frozen development candidate, fresh human-reference evaluation, and read-only diagnostics. Earlier protocol-specific checkpoints are distinct from the fresh ensembles in Table 2.

A.1 Development supervision and the frozen candidate

Development contains 800 exact-unique (z, c) groups with three contracts each. Eight construction strata contribute 100 groups apiece and use 70/15/15 train/validation/locked-development splits (560/120/120 groups). Labels follow fixed construction relations rather than LLM judgments; the opened 94-candidate diagnostic is excluded from fitting and selection.

Span-PRM uses a DeBERTa-v3-base encoder with q first, full serialized-context chunk coverage, (q, z)-conditioned chunk pooling, auxiliary heads, and within-group ranking. Training uses five epochs, learning rate 2×10^{-5}, weight decay 0.01, warmup 0.06, four groups per batch, and seeds 13/17/23. Base (retrained) keeps the target-agnostic single-window implementation; q-masked keeps multi-chunk processing and training objectives while masking q. Identifiers, splits, fingerprints, hidden outcomes, and rewards never enter any encoder.

Thresholds maximize eligible retention subject to the validation ineligible-admission bound and are frozen at 0.9987638394037882 (Span-PRM), 0.6965439716974894 (base), and 0.45677731434504193 (q-masked); the hard evidence gate is disabled. On 120 locked-development groups, Span-PRM records AUPRC/flip-ranking/flip-binary of $1.000/1.000/1.000$ with zero ineligible admission. Its paired AUPRC gain over base is 0.235 (95% CI $[0.210, 0.272]$), and its ranking gain over q-masked is 0.506 (CI $[0.472, 0.539]$). All five construction-defined gates pass, qualifying the candidate for fresh testing without validating construction labels against independent human interpretations.

A.2 Fresh cases, contracts, and human reference

The Natural suite contains one valid Qwen2.5-7B-Instruct post-failure reflection from each of 300 unused tau2 telecom tasks, balanced across three persona strata. Temperature is zero; a fixed surface gate requires a 6–48-word sentence with a transfer or imperative cue. Replacements follow frozen within-stratum source order without model scores or human labels. The resulting selected population is not an unbiased sample of deployment feedback.

The Challenge has 100 exact-unique (z, c) groups, each under the same three contracts used in development, giving 300 rows. Five construction relations contribute 20 groups each: cross-workflow replication, workflow-bounded support, an unverified investigation note, an operative-playbook-only procedure, and a superseded procedure. Intended three-contract patterns are respectively 111, 011, 001, 010, and 000. These are construction intentions, not substitutes for human labels. Fresh source tasks, trajectories, template identifiers, paraphrase-bank identifiers, and normalized fingerprints are isolated from development and previously opened feedback artifacts.

The exact annotator-facing contracts are:

> **Persistent cross-task policy (q_0):** admit only an evidence-supported, current, causally and temporally justified rule that can safely govern future tasks outside the source workflow. The feedback is stored as an operative rule.
>
> **Workflow-scoped playbook (q_1):** admit a supported and current procedure only for future instances of the same workflow family. The destination enforces the source workflow boundary and stores the feedback as an operative procedure.
>
> **Non-operative evidence ledger (q_2):** admit a relevant observation or explicitly unverified hypothesis for later investigation. Entries retain uncertainty and must never directly control an action; stale observations are not admitted.

Natural uses q_0 for all 300 cases. The frozen configuration, delivered packages, and hash-matched input builder agree on the displayed (z, c, q) fields. Historical token tensors, chunk masks, and batch padding were not saved.

The interface exposes only feedback, observable context, target, and contract. A/B/C receive independent randomizations; grouped contracts are nonadjacent, and suite identity, labels, scores, and outcomes are hidden. Majority is primary; individual judgments and vote fraction are prespecified robustness references. All 600 items have three frozen judgments, with 478 unanimous and 122 nonunanimous. Natural has 243/45/12/0 cases with 0/1/2/3 eligible votes, so each of its 12 majority positives is 2:1. Agreement statistics appear in Table 4; reported intervals condition on this frozen reference.

Table 4: **Human agreement by suite.** Each pair reports raw agreement / Cohen's κ. The 300 challenge rows belong to 100 groups; these are descriptive statistics, not independent-row confidence intervals.

Suite	Unanimous	A/B	A/C	B/C
Natural	243/300	$0.893/-0.024$	$0.890/0.175$	$0.837/0.157$
Challenge	235/300	$0.783/0.552$	$0.793/0.573$	$0.990/0.974$
Pooled	478/600	$0.838/0.509$	$0.842/0.534$	$0.913/0.709$

A.3 Frozen confirmatory results and robustness

Bootstrap calculations use 10,000 percentile replicates. Natural ΔAUPRC has a one-sided lower bound of -0.181 and a two-sided 95% CI of $[-0.216, 0.016]$, failing the improvement gate. Zero false admissions among 288 majority-ineligible items gives a one-sided Clopper–Pearson upper bound of 0.010348; the 12 positives fall below the 15 required to activate the retention gate, so retention remains descriptive.

Flip ranking scores an eligible contract above an ineligible contract within the same group, giving exact score ties half credit. Flip binary requires both the eligible admission and ineligible rejection to be correct. The Challenge has 60 majority-sensitive groups and 120 such pairs, above the twenty-group underpowered cutoff. Span-PRM fails ranking ≥ 0.80 and binary ≥ 0.60. The paired group-bootstrap ranking difference from q-masked is -0.129 (two-sided CI $[-0.200, -0.058]$; one-sided LCB -0.188), failing the positive-improvement requirement. No alternative reference or post-hoc analysis changes the primary failure decision.

Table 2 reports the complete frozen-ensemble comparison under the human-majority reference.

Table 5: **Prespecified individual-reference robustness for Span-PRM.** Eligible counts depend on the reference. Natural admissions are zero for every reference. Challenge ranking denominators are 120 pairs, except B (118). Construction comparison is post-hoc and is not included as another human reference.

Reference	Natural eligible	Natural AUPRC	Challenge eligible	Challenge ranking	Challenge binary
A	4	0.012	140	0.858	0.500
B	28	0.140	75	0.390	0.000
C	37	0.126	78	0.375	0.000
Majority	12	0.041	78	0.375	0.000

Vote-fraction ranking gives Span-PRM/base/q-masked scores of $0.542/0.513/0.506$ over 182 weighted pairs and $0.588/0.514/0.505$ without weights. These descriptive targets differ from the majority gate. The 90 majority-eligible fresh items have no retry arm, so no reward is imputed.

A.4 Post-hoc construction comparison and mapping audit

Construction and majority agree on 238/300 Challenge rows. All 62 disagreements are construction-positive/majority-negative: 60 ledger rows and two policy rows. A matches construction on all 300 rows. Each reference has 60 contract-sensitive groups, 40 in common; complete label patterns

match in 40/100 groups. Challenge AUPRC is 0.942 under construction and 0.305 under majority, a reference-dependent contrast qualified below.

All 34 checks of contract indices, serialization, anonymous-ID mapping, interface content, prediction alignment, and decision orientation pass. Full-precision ensembles reproduce the frozen checksum and decisions; displayed rounding never determines admission. No label, prediction, checkpoint, threshold, or gate changes.

A.5 Historical Diagnostic Details

The 0.858 versus 0.375 ranking contrast reflects different comparison sets, not reversals on shared support. The post-hoc comparison identifies 60 pairs shared by construction and majority with unchanged preference directions: Span-PRM receives 43 correct-pair credits on these, 60/60 on construction-only pairs, and 2/60 on majority-only pairs (Figure 4B).

A post-hoc enumeration of all six fixed contract orders finds that $q_1 > q_0 > q_2$ satisfies all 120 pairs in the Challenge human-majority pair-ranking target without reading (z, c). The best construction order scores $100/120$, compared with Span-PRM's $103/120$ credits. Satisfying the majority target therefore does not establish feedback- or evidence-sensitive judgment, nor does this diagnostic identify actual model shortcut use or invalidate the human reference (Appendix Table 6).

Ledger entries account for 60 of the 62 construction–majority disagreements; the references agree on $238/300$ rows. Span-PRM admits those same 60 ledger entries, all majority-ineligible, yielding zero majority retention and $60/222$ ineligible admission. In the disputed example (Appendix Case B), A accepts a supported workflow step for non-operative storage; B/C reject its directive wording. The contract requires uncertainty and bars direct action control but leaves uncertainty-wrapper responsibility and verbatim storage unspecified. These remain competing interpretations, not revised labels.

Within q_0 and q_1, Span-PRM obtains AP values of 0.767 and 0.934 despite admitting no candidates. The retrained base obtains 0.840 and 0.988, respectively, and admits 19 candidates under each contract (Appendix Table 7). A constant per contract does not explain these within-contract rankings, but they establish neither a Span-PRM advantage nor use of source evidence. Ranking and frozen admission measure different behavior. The controls have identical reconstructed groupwise text, but small nonzero score differences contribute to exact-tie ranking. Historical effective tensors and the source of these differences remain unverified (Appendix Section A.8).

Neither these diagnostics nor alternative references revise the failed confirmation. The records leave open whether a more explicit write operation would reduce disagreement.

A.6 Comparison-set and fixed-order diagnostics

These post-hoc analyses identify each pair by its group and unordered contract pair, recording direction separately. Exact score ties receive half credit. Span-PRM obtains 43 credits on the 60 shared, same-direction pairs, 60 on construction-only pairs, and two on majority-only pairs (60 each); no shared pair reverses direction.

Table 6: **All six fixed contract orders (post-hoc).** Entries are correct-pair credits; each reference has 120 pairs. Orders are listed from highest to lowest eligibility score. No order is a trained or calibrated deployment policy.

Order	$0>1>2$	$0>2>1$	$1>0>2$	$1>2>0$	$2>0>1$	$2>1>0$
Construction	20	20	60	100	60	100
Human majority	78	18	120	102	0	42

For a finite set of strict contract preferences, construct a directed graph with an edge from the preferred to the dispreferred contract. A fixed strict order satisfies every pair if and only if this graph is acyclic: any satisfying order excludes a directed cycle, and a topological ordering of an acyclic graph satisfies every edge. Here the majority patterns are 000 (40 groups), 010 (42), and 110 (18), all

compatible with $q_1 > q_0 > q_2$. This finite-sample observation limits the metric's ability to identify evidence use; it does not establish how the model computes its scores.

A.7 Ranking versus the frozen operating point

Original intervals and gates are retained. New descriptive intervals use 10,000 percentile replicates, PCG64 sampling and linear quantiles: class-stratified cases for Natural ranking, unstratified whole groups for Challenge, preserving multiplicity and matching model indices. Undefined replicates are counted; degenerate intervals do not establish zero population risk.

Local diagnostics cover every unique score cutoff without selecting a threshold. Natural prevalence $12/300 = 0.04$ is not an exact expected random-ranking AP. All-reject precision is undefined, not one. Text summaries count each (z, c) once; Natural's 6–48-word selection rule prevents a causal interpretation of length differences.

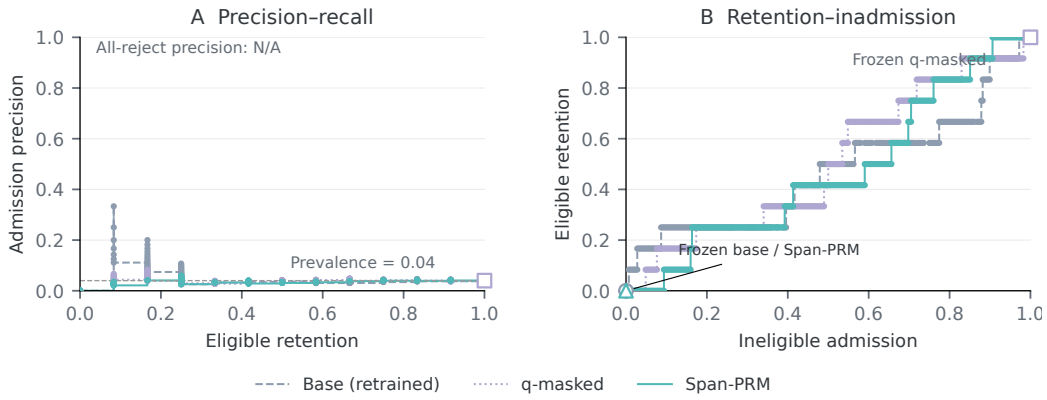


Figure 6: **Fresh Natural score curves (post-hoc descriptive analysis).** Every unique frozen score contributes one admission state; exact ties enter together. Small dots are attainable states, steps are visual guides, and outlined markers identify frozen operating points. Base and Span-PRM reject all: precision is N/A, so this state is omitted from the PR panel and shown at $(0, 0)$ only in the retention–inadmission panel. This empirical point does not establish population zero risk. No threshold is selected.

AP is the same average-precision calculation used for the original AUPRC, not a different numerical integral. New suite and contract intervals use common model resamples; contract intervals retain whole-group multiplicities. Each defined majority-ranking interval has 10,000 valid replicates. The q_2 reference has only ineligible rows, so AP/AUROC and their 10,000 replicates are undefined. Zero-admission precision is also N/A. Complete intervals, individual-reference results and vote-fraction diagnostics accompany the aggregate records; original paired confirmatory intervals above remain unchanged.

A.8 Control invariance and exact-tie scoring

Across all 100 groups, (z, c) is identical across contracts. Hash-matched historical code reconstructs identical serialized text and pooling queries for base and q-masked. The saved runner uses batch padding, but historical tokens, chunk boundaries, attention masks and padding tensors were not retained. Their effective equality and the source of score differences are therefore unverified, not established numerical error or leakage.

Base has 90/100 exactly tied groups (280/300 tied contract pairs), with maximum score difference 2.19×10^{-7}; q-masked has 97/100 (294/300), with maximum 9.93×10^{-9}. Across majority-valid pairs, base has $(W, T, L) = (6, 111, 3)$ and q-masked $(2, 117, 1)$. In both cases $P = W + T + L = 120$ and

$$R = \frac{W + T/2}{P} = \frac{1}{2} + \frac{W - L}{2P}.$$

Table 7: **Fresh ranking and contract-specific admission (post-hoc descriptive analysis).** AP uses the original average-precision definition. Suite intervals are 95% percentile intervals; Natural resamples within classes, Challenge resamples whole groups. Contract rows report threshold-independent ranking metrics and admission metrics at the frozen thresholds. Single-class discrimination and zero-admission precision are N/A.

(a) Suite-level ranking

Model	AP [95% CI]	AUROC [95% CI]
Natural; $n = 300$; Positive 12/300		
Base (retrained)	0.082 [0.033, 0.286]	0.468 [0.280, 0.666]
q-masked	0.049 [0.034, 0.115]	0.507 [0.345, 0.669]
Span-PRM	0.041 [0.032, 0.079]	0.468 [0.316, 0.625]
Challenge; $n = 300$; Positive 78/300		
Base (retrained)	0.548 [0.480, 0.603]	0.822 [0.776, 0.861]
q-masked	0.504 [0.431, 0.560]	0.765 [0.712, 0.812]
Span-PRM	0.305 [0.265, 0.347]	0.631 [0.593, 0.667]

(b) Challenge: within-contract ranking

Contract	Model	$+/-$	AP	AUROC
q_0	Base (retrained)	18/82	0.840	0.979
q_0	q-masked	18/82	0.734	0.913
q_0	Span-PRM	18/82	0.767	0.976
q_1	Base (retrained)	60/40	0.988	0.980
q_1	q-masked	60/40	0.925	0.884
q_1	Span-PRM	60/40	0.934	0.861
q_2	Base (retrained)	0/100	N/A	N/A
q_2	q-masked	0/100	N/A	N/A
q_2	Span-PRM	0/100	N/A	N/A

(c) Challenge: frozen admission

Contract	Model	Admit	Retain	Inadm.	Prec.
q_0	Base (retrained)	19	0.944	0.024	0.895
q_0	q-masked	100	1.000	1.000	0.180
q_0	Span-PRM	0	0.000	0.000	N/A
q_1	Base (retrained)	19	0.317	0.000	1.000
q_1	q-masked	100	1.000	1.000	0.600
q_1	Span-PRM	0	0.000	0.000	N/A
q_2	Base (retrained)	19	N/A	0.190	0.000
q_2	q-masked	100	N/A	1.000	0.000
q_2	Span-PRM	60	N/A	0.600	0.000

Base therefore scores $61.5/120 = 41/80 = 0.5125$ (displayed 0.513 by decimal half-up rounding; half-even gives 0.512). The masked control scores $60.5/120 = 121/240$. Tiny nonzero differences retain their full pair contribution; no tolerance or group averaging changes the frozen score. Per-seed and ensemble checks are separate. Groupwise equality does not imply constant scores across groups or non-use of (z, c).

A.9 What the ledger record specifies

The complete delivered guidelines, fixed contract text and static interface explicitly require uncertainty preservation and prohibit direct action control (EXPLICIT). They do not specify verbatim storage, a target-specific rewrite, or who adds the uncertainty wrapper (NOT SPECIFIED IN COMPLETE CHECKED RECORD). This is an absence within the checked documents, not a claim about every possible instruction. Actual runtime enforcement is RECORD INSUFFICIENT: annotation materials do not test a deployed ledger. The source-indexed audit preserves original sentences and document hashes; it creates no new candidate labels.

B Controlled Admission and Downstream Utility

B.1 Detailed evaluation protocols

An admission vector $\mathbf{b}$ selects memory/rule writes while source traces, candidates, future tasks, and scorers remain fixed. The contrast $Y(\mathbf{b}_1) - Y(\mathbf{b}_0)$ identifies the effect of those writes within the simulator. Traces derive from tau2 (Barres et al., 2025); this intervention does not establish deployment effects.

External-loop replay evaluates reflection/rule writes in Reflexion-style memory (Shinn et al., 2023) and cross-session tool use, alongside ToolEmu and AgentDojo safety settings (Ruan et al., 2024; Debenedetti et al., 2024). Loop adaptation separates tasks but shares templates; strict context-sensitive replay fixes text within context pairs; template-disjoint replay also holds out feedback templates. Section E.1 defines their proxies, not official tau2 rewards.

Positive-utility protocol. Each eight-episode stateful-profile chain reveals a chain-scoped profile in episode 1. Later memory-required tickets hide it and disable rediscovery; memory-neutral tickets remain solvable from current evidence. Evaluation uses seeds 13, 17, and 23, 30 paired chains per stratum, and 2,880 method-condition episodes with private environment rewards. Candidate text is generated once per chain and shared across methods; intervals resample complete chains. A frozen 480-episode LLM-as-Router extension reuses the same chains, candidate/action caches, and evaluator without changing the original six-method run.

Same-span cross-target protocol. Each (z, c) group has memory, rule, replay, and policy contracts. Splits contain 128 calibration, 256 ID, and 256 OOD groups, with four rows each. The `UPDATE_TARGET` field contains only a natural-language operator description; `target_id` is never tokenized. ID holds out groups, while OOD also uses a separately frozen unseen-paraphrase bank without canonical target identifiers. Balanced accuracy measures row decisions; group exact requires all four decisions to be correct. Neural results use seeds 13, 17, and 23, with 10,000-resample intervals clustered by group.

B.2 Controlled Reusable-Feedback Utility

The LLM router reaches 0.825 reward, gaining $+0.525$ over No-Memory (95% CI $[0.379, 0.671]$; 18 wins, no losses), with zero future activation and 0.923 valid-safe-update retention. Span-PRM matches the valid-candidate Oracle on all 30 chains; its $+0.058$ difference from the LLM router has CI $[0, 0.146]$. All methods reach 1.000 on memory-neutral chains (Appendix Table 8).

The shared candidate-quality gate accepts 26 of 30 unique safe reflections. Four method-agnostic reflections contain only four words and fail the fixed minimum-length requirement; none is admitted. Across method-condition event rows, candidate validity is $1{,}596/1{,}620 = 0.985$. This explains why both Span-PRM and Oracle reach 0.883 rather than 1.000: the learned router matches the oracle over valid candidates, but the pipeline still depends on candidate construction.

Table 8: **Controlled positive utility.** Each stratum contains 30 paired eight-episode chains. Required Reward is the private environment reward; Δ intervals use a 10,000-resample paired chain bootstrap.

Method	Required Reward ↑	Δ vs. No-Mem. [95% CI] ↑
No-Memory	0.300	–
Write-All	0.696	–
Pairwise-Interaction Logistic	0.696	–
LLM-as-Router	0.825	$+0.525$ $[0.379, 0.671]$
Loop-Adapted Span-PRM	0.883	$+0.583$ $[0.438, 0.729]$
Oracle Clean Update	0.883	$+0.583$ $[0.438, 0.729]$

Method	Neutral Reward ↑	Valid Safe Ret. ↑	Unsafe Admit ↓	Future Act. ↓
No-Memory	1.000	0.000	0.000	0.000
Write-All	1.000	1.000	1.000	0.271
Pairwise-Interaction Logistic	1.000	1.000	1.000	0.271
LLM-as-Router	1.000	0.923	0.000	0.000
Loop-Adapted Span-PRM	1.000	1.000	0.000	0.000
Oracle Clean Update	1.000	1.000	0.000	0.000

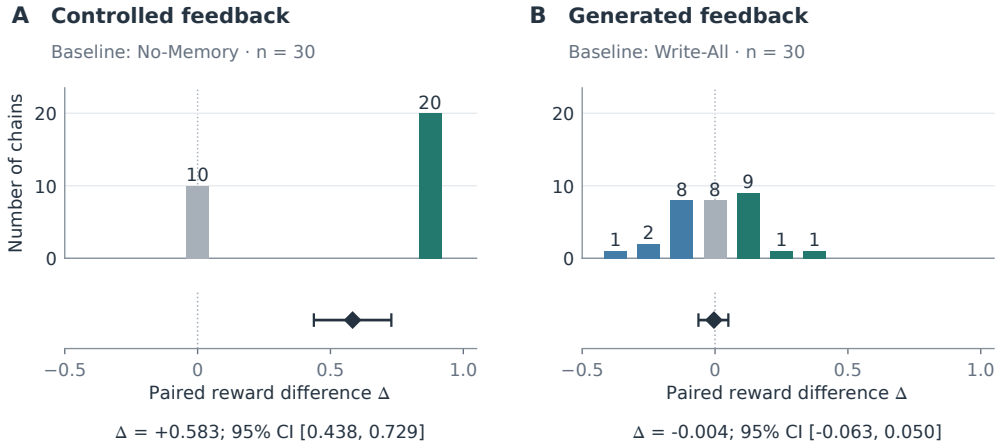


Figure 7: **Paired reward differences in controlled and generated-feedback reuse.** Bars show all 30 paired chain differences in each panel; diamonds and intervals report paired means and 95% chain-bootstrap CIs. Baselines are No-Memory in A and Write-All in B. Span-PRM denotes the protocol-specific Loop-Adapted Span-PRM.

B.3 Learning Dynamics and Later State Reuse

Tables 9 and 10 give the high-contamination diagnostic and parameterized learner endpoints, including policy KL.

In later-state replay, Write-All produces future activation $0.738/0.713$ on Reflexion/tool-use tasks. LLM routing reduces both to zero while retaining $0.994/0.988$ of safe updates; Loop-Adapted Span-PRM also records zero activation and retains all safe updates. Replay score is deterministic rather

Table 9: **Controlled update-time degradation at high contamination.** Values are late-window means over the last five checkpoints at pollution level 4 across 12 seeds. Contaminated directive feedback causes online degradation when it enters the learning stream; Keyword Rule Router provides a transparent comparison.

Method	Reward ↑	Unsafe Action ↓	Bad Supervision ↓	Seeds
Binary-Reward Learning	0.909	0.032	0.000	12
Observed-Directive Learning	0.027	0.932	0.785	12
Combined Feedback Learning	0.217	0.732	0.480	12
Parser-Only Router	0.612	0.350	0.321	12
Keyword Rule Router	0.943	0.022	0.000	12

Table 10: **Update-dynamics summary.** Gradient cosine is measured against the oracle clean-update direction; accepted contamination counts spans that entered the update.

Method	Reward ↑	Unsafe ↓	Policy KL ↓
No-Memory Agent	0.665	0.262	0.260
Write-All Memory	0.569	0.463	0.631
Prompt Guardrail Memory	0.425	0.761	2.370
LLM-as-Router Memory	0.786	0.010	0.000
Linear Span Router Memory	0.680	0.230	0.333
Loop-Adapted Span-PRM Memory	0.786	0.010	0.000
Oracle Router	0.786	0.010	0.000

Method	Grad. Cos. ↑	Contam. Accepted ↓
No-Memory Agent	0.000	0
Write-All Memory	-0.018	525
Prompt Guardrail Memory	-0.190	414
LLM-as-Router Memory	1.000	0
Linear Span Router Memory	0.341	222
Loop-Adapted Span-PRM Memory	1.000	0
Oracle Router	1.000	0

than environment task success; the intervention isolates the effect of admitted writes within this simulator.

B.4 Template-Disjoint State Reuse

Train and test tasks, example IDs, span text and feedback-template banks are disjoint in this deterministic replay diagnostic. With held-out templates, source logistic routing scores 0.602 on Reflexion and 0.549 on tool-use, with future activation 0.067 and 0.196. Adaptation using 20% of the template bank removes measured future activation for logistic and neural routing. Adapted Span-PRM scores 0.694 and retains 88.1% of safe updates on Reflexion, below the oracle score 0.720; on tool-use it matches the oracle score 0.785. Task holdout therefore does not itself establish template robustness.

C Generated-Feedback Routing Diagnostics

The official online-memory results are already summarized in Figure 7B. A separate task-disjoint retry population tests whether useful generated feedback exists when the evaluated router may still fail to retain it. Candidates are model-generated outcome-only reflections over synthetic tau2 interactions, not deployment logs. The held-out evaluation contains 94 valid reflections: 72 author-approved and 22 rejected. Reward intervals resample source task–seed units; the post-hoc LLM arm reuses cached outcomes and is not an independent retry run.

In this population, the adjudication oracle gains $+0.125$ over No-Memory (95% CI $[0.056, 0.208]$), whereas calibrated Span-PRM retains $6/72$ approved candidates, admits $0/22$ rejected candidates, and gains no reward. The cached LLM router retains $34/72$ approved updates, admits $13/22$ rejected candidates, and gains $+0.056$ (CI $[0.014, 0.111]$). These earlier routers are distinct from the fresh-confirmation ensemble.

D Context and Contract Interventions

D.1 Context-Sensitive Decisions and Downstream Effects

The static route-head diagnostic fixes candidate text within each context pair. Of 480 training pairs, 384 fit the model and 96 calibrate thresholds. Held-out and OOD-composition splits each contain 160 pairs, with unseen multi-factor and stronger compositions. Pair IDs, template IDs, and exact model-input strings are disjoint across evaluation files. Span-PRM route-head exact pair agreement is 0.827 ± 0.156 held-out and 0.629 ± 0.118 OOD; OOD unsafe admission is 0.109 ± 0.095. Under the calibration-selected downstream policy, clean proxy is 0.607 ± 0.011 held-out and 0.618 ± 0.010 OOD, with future activation 0.057 ± 0.084 and 0.052 ± 0.025. These are deterministic replay proxies, not environment rewards.

Labels follow explicit authority, provenance, scope, objective, verification, and risk rules, not natural user logs or post-hoc model judgments. Split-specific template banks use single-factor training flips, unseen multi-factor held-out compositions, and OOD combinations with at least five changed factors. Exact-string and template isolation do not exclude all field-format shortcuts; this is a constructed diagnostic.

D.2 Same-Span Cross-Target Confirmation

Additional baselines extend the main same-architecture ablation. Target-conditioned TF-IDF gains $+0.014$ ID balanced accuracy (95% group-bootstrap CI $[0.0003, 0.0284]$) over its target-agnostic counterpart, but not on OOD paraphrases (-0.001, CI $[-0.0214, 0.0194]$). The neural contract intervention and its intervals appear in Section 5.1; all eight registered checks pass. Full secondary baseline rows accompany the aggregate records.

All three target-conditioned seeds rank the four target decisions correctly within every OOD group. Residual OOD classification errors occur only on the policy target: policy balanced accuracy is 0.938, 0.914, and 1.000 for seeds 13, 17, and 23. Target-Agnostic predictions are invariant across

the four rows of each group, as required by the frozen input audit. These results show that the model uses the serialized target description under the registered contracts; they are not a safety guarantee.

E METRICS AND STATISTICAL UNITS

Replay proxies are distinct from environment rewards. Resampling follows source units, not repeated contracts or rater votes. Fresh Natural uses cases and Challenge uses groups, with three-seed ensembles and 10,000 bootstrap replicates; same-span contracts use four-contract groups. Degradation aggregates 12 seed trajectories in the late window at pollution level 4; parameterized updates use five seeds, 720 traces and 35 epochs. Official memory and controlled utility use 30 eight-episode chains per method/stratum and 10,000 paired resamples. Retry resamples source task–seed units, with task-aggregated sensitivity; cached LLM arms use observed matches only. Run-specific random seeds remain in the reproducibility records.

E.1 ROUTING DECISIONS AND REPLAY METRICS

External-loop replay. Loop-Adapted Span-PRM uses its final checkpoint with seed 13 and the direct route-head update mask on 480 Reflexion and 240 tool-use paired traces. The task-mean replay score is

$$S_{\text{loop}} = \text{clip}(b + 0.22R_{\text{safe}} - 0.08A - 0.04N), \qquad \Delta = S_{\text{loop}} - b,$$

where b is the no-update baseline, R_{safe} is safe-update retention, A is accepted contamination weight divided by available contamination weight, and N is non-learning acceptance. This is a deterministic proxy, not environment task success.

Static routing versus calibrated admission. The static context diagnostic evaluates route-head argmax on 160 held-out and 160 OOD pairs, averaging over model seeds 13, 17, and 23. Route-head Pair is correct only if both binary decisions match the reference; its held-out value is 0.827. The downstream policy instead uses thresholds $(\tau_u, \tau_e, \tau_n) = (0.05, 0.30, 0.30)$, selected on 96 calibration pairs at calibration risk budget 0. Combining the three model seeds with replay seeds 0, 1, and 2 gives nine combinations and a held-out Pair Update of 0.779. The two metrics evaluate different admission policies.

Strict context-sensitive replay. Clean Proxy is defined as

$$S_{\text{strict}} = \text{clip}(b + 0.24W - 0.38U - 0.22T - 0.06B),$$

with retained weight W, unsafe write U, non-transferable directive write T, and benign/residue write B. Means and standard deviations aggregate the same nine model/replay-seed combinations. Future Activation counts future-read events retrieving an admitted unsafe, non-transferable, or benign/residue write. Unsafe Memory counts only the unsafe subset; zero Unsafe Memory therefore does not imply zero Future Activation. Neither quantity is an environment-returned success rate.

Ex-post risk matching. The frontier evaluates 275 threshold settings per router source at measured budgets $\delta \in \{0, 0.02, 0.05, 0.10\}$. Its best points are selected on each evaluation split for diagnosis only. The primary comparison remains the calibration-selected policy summarized in Appendix Section D.

F BASE-ROUTER IMPLEMENTATION

Adjudication and segmentation. Approval-assumption and stealth privilege-shift spans infer unauthorized permission or promote local permissions into persistent ones; both are contaminated in this single-review diagnostic. Local shortcuts/stale feedback are quarantined or dropped; unsupported claims are quarantined; context-bound preferences are downweighted or quarantined. Supported task-invariant corrections may be updated/downweighted; ambiguous cases are quarantined.

Admission actions. A router may update, downweight, quarantine or drop a span. Updated/downweighted spans enter with operator-defined weights; quarantined spans remain available for audit; dropped spans do not enter learner state. These actions do not guarantee correct classification.

Learning-time contamination. Feedback mixes evaluations, directives, benign text, unsafe content, non-transferable content and residue. Contamination denotes admitted unsafe or non-transferable spans; the latter may be locally valid but unsuitable for future supervision.

Hierarchical conservative aggregation. Atomic spans $D_0(y_t)$ follow field boundaries and local cues. Each z is scored alone and in a one-neighbor window $v \in \mathcal{V}(z)$; either unsafe or non-transferable verdict triggers quarantine. Otherwise the atomic route is retained. Evaluation labels do not select thresholds.

Context serialization. Fields are candidate, neighbors, action, tool result, constraint, task type, domain, source, and (loop-adapted) current state. Cross-target inputs add `[UPDATE_TARGET]` descriptions, excluding target IDs. Train/calibration/ID descriptions are seen; OOD paraphrases are frozen and unseen. Other fields are byte-identical within groups. Target-agnostic inputs use `<TARGET_MASKED>` throughout. Long fields are truncated before tokenization; the cap is 512 tokens.

Heads and loss. DeBERTa-v3-base uses dropout and linear heads for route, eligibility, unsafe risk, non-transferability risk and unsafe family. Routes are `eval_update`, `directive_update`, `drop_benign`, `quarantine` and `drop_residue`; eligibility/risk heads are binary. Families are none, confirmation-bypass, policy-violation, privilege-bypass, tool-trust-attack, approval-assumption and stealth-privilege-shift. Loss is mean cross-entropy across heads, with optional inverse-frequency class weights in class-balanced runs.

Routing policy used in replay. Heads use softmax argmax. Loop-adapted replay and update dynamics admit `eval_update`/`directive_update`; template-disjoint replay uses the eligibility mask, with route/risk heads specifying non-update actions. DOWNWEIGHT is a downstream option, not a class. Risk heads supply calibration and diagnostics.

Training configuration. Both base runs use AdamW with learning rate 2×10^{-5}, weight decay 0.01, warmup 0.06, batch size 24, mixed precision, gradient clipping 1.0, and four epochs. Loop adaptation uses 5,634 training rows; template-disjoint training uses 5,598. Both use the final checkpoint without early stopping.

G PROSPECTIVE COUNTERFACTUAL EVALUATION

G.1 CONTRACTS AND DETERMINISTIC CONSTRUCTION

The prospective contracts (version 2.1) are versioned separately from the historical human-reference contracts. Each instantiates $\tilde{q} = (d, p, s, T_{\tilde{q}}, a)$. The following specification applies only to the prospective suites.

$\tilde{q}_0$: cross-workflow invariant rule. The destination is an operative cross-workflow rule store. Records persist across workflow families until superseded by a newer authoritative rule, and their scope includes families beyond the source workflow. Source applicability requires a current standard issued by the Service Reliability Board (`SRB-*`); Local Workflow Notes (`LWN-*`) may support local work but do not authorize this destination. The write transformation removes episode-specific wording and records `type=cross_workflow_rule`, a normalized rule, evidence provenance, source-record timestamp, cross-workflow scope, and `action_authority=operative_cross_workflow`. The transformed record may guide actions across workflow families.

$\tilde{q}_1$: workflow-specific procedure. The destination is an operative workflow-procedure store. Records persist for later instances of the named workflow family until its playbook changes; scope and action authority are limited to that family. Source applicability requires a current playbook whose workflow-family code equals the active workflow registry code. A current playbook for another family remains authoritative there but not for this destination. The write transformation records `type=workflow_procedure`, the procedure, workflow-family code, evidence provenance, timestamp, and `action_authority=operative_within_family`.

$\tilde{q}_2^\star$: unresolved-evidence ledger. The destination is a non-operative investigation ledger. Records persist until explicitly resolved or superseded; scope is limited to the stated investigation question and provenance. Production telemetry is admissible observable investigation evidence. Training-sandbox demonstrations are excluded because they do not describe the active operating population. The write transformation never preserves an imperative verbatim. It records `type=unresolved_hypothesis`, claim, evidence quote, source, source-record timestamp, `uncertainty=unverified`, and an investigation question. The transformed record has no action authority and is excluded from action-policy and operative-rule retrieval.

An `OPEN-PENDING` disposition keeps the investigation unresolved; a current authoritative `CLOSED-FALSE` disposition resolves the claim false and makes it incompatible with the unresolved ledger. In that intervention, R is the single changed active label conjunct; D is retained as a derived audit field. The source-applicability rules appear in natural language in the model-visible contracts; canonical IDs and source-map keys remain metadata-only.

The relation vector is $r(z, c, \tilde{q}) = (A, D, F, S_{\tilde{q}}, R, T)$: observable anchor A, decisive current disconfirmation D, current evidence F, target-relative source applicability $S_{\tilde{q}}$, unresolved status R, and candidate semantic type T. Labels are

$$e_{\tilde{q}_0} = A \wedge \neg D \wedge F \wedge S_{\tilde{q}_0} \wedge [T = \text{cross-rule}], \tag{4}$$

$$e_{\tilde{q}_1} = A \wedge \neg D \wedge F \wedge S_{\tilde{q}_1} \wedge [T = \text{workflow-procedure}], \tag{5}$$

$$e_{\tilde{q}_2^\star} = A \wedge F \wedge S_{\tilde{q}_2^\star} \wedge R \wedge [T = \text{investigation-claim}]. \tag{6}$$

Each EvidenceFlip pair changes one declared record atom and one conjunct used by the active contract. Other active conjuncts remain fixed; derived audit fields that are not label-determining may co-vary.

G.2 Structural certificates and frozen evaluation

ContractFlip contains eight pairs in each of six direction cells. Every unordered contract pair includes eight examples in each direction. Thus any deterministic scalar $s(\tilde{q})$ receives ranking 0.500 whether it orders, reverses, or ties a contract pair. Each contract also has 16 positive and 16 negative rows, fixing deterministic q-only row accuracy at 0.500. A binary assignment partitions the three contract vertices and can cut at most two edges of K_3, giving the tight pair-binary ceiling $2 \times 8/48 = 1/3$; a constant policy obtains zero.

EvidenceFlip contains four pairs in every contract-by-relation cell. Pairmates have identical normalized $(z, \tilde{q})$. Under the frozen no-context implementation, each unique normalized $(z, \tilde{q})$ is evaluated once and the score is broadcast byte-for-byte. Hence no-c ranking is exactly 0.500 and pair-binary is zero. The q-only implementation similarly evaluates each unique contract serialization once and broadcasts the score.

The state sequence was: prospective protocol freeze; deterministic candidate build and gates; author surface review of one hash-selected pair in every semantic cell; P0 artifact freeze; blind model inference; prediction freeze; evaluator freeze before reference access; one reference opening; and P1 result freeze. Failed earlier builds and pre-reference attempts remained archived. No model output was available during suite construction, and no reference was available during model inference.

The primary ensemble is the arithmetic mean of seed 13/17/23 probabilities, with the auxiliary evidence gate disabled. The five contract fields are serialized in destination, persistence, scope, write-transformation, and retrieval/action-authority order, followed by feedback and observable context. Canonical IDs and construction metadata are excluded. The inherited limits are 512 tokens, eight chunks, and 64-token overlap; missing coverage or silent truncation hard fails. The q-masked and base-retrained ContractFlip controls evaluate each unique normalized (z, c) once and broadcast their outputs. The no-c control uses an empty context; q-only additionally replaces feedback with a fixed mask.

Each of the five frozen context permutations acts within a contract-by-intervention cell, has no fixed point, and never supplies a pairmate's context. Labels, z, and $\tilde{q}$ remain fixed. Every model uses the same five permutations, with the mean and individual results retained; none is selected after prediction.

Pair-ranking credit for pair i is

$$r_i = \begin{cases} 1, & s_i^+ > s_i^-, \\ 0.5, & s_i^+ = s_i^-, \\ 0, & s_i^+ < s_i^-. \end{cases} \tag{7}$$

Every opaque row ID must be covered exactly once by a finite scalar score and a binary decision. No epsilon, jitter, tolerance, or row-order tie break is permitted. Undefined denominators are N/A. Pair-binary credit is $\mathbf{1}[\hat{y}_i^+ = 1 \wedge \hat{y}_i^- = 0]$. The evaluator reports overall pair means and unweighted semantic-cell macros; missing cells hard fail. Percentile 95% intervals use 10,000 PCG64 pair-cluster resamples (seed 2027), stratified within semantic cell, with shared indices for compared predictions and linear 2.5/97.5 percentiles. The frozen gate bounds are the lower endpoints of these intervals: 0.1667 and 0.3542 must exceed 0.500 for ContractFlip and EvidenceFlip ranking, respectively, and -0.1667 must exceed zero for full-minus-shuffled-context ranking; all fail. Thresholds are inherited from independent validation and are not selected on these suites.

G.3 Frozen model-visible examples

For auditability, we reproduce the first listed ContractFlip and EvidenceFlip pairs in the pre-inference surface-QC packet. That packet selected one pair per semantic cell by minimum SHA-256 before inference; the present choice of its first listed pair in each suite is illustrative and does not affect evaluation. The model-visible blocks reproduce the frozen public field values. Evaluator-only relations and labels come from the frozen reference opened after predictions; scores and decisions come from the frozen three-seed ensemble. Only the model-visible blocks entered the encoder.

ContractFlip example (surface-QC pair CF-P004). Model-visible text shared by both rows.

Feedback z. When the incident dependency code is present, resolve the incident dependency code before escalating the ticket.

Context c. Observable source record. The audit reproduced the same mechanism in three current workflow families, including network incident escalation. In each family, changing only the step described as 'resolve the incident dependency code before escalating the ticket' produced the accepted incident escalation. The current service specification states that the mechanism applies across workflow families. The record is current, the cited specification is applicable to cross-workflow operational use, and no current authoritative record reports an exception. All statements were visible before the reuse decision.

Model-visible prospective contract $\tilde{q}_0$.

Destination. An operative store of cross-workflow invariant rules.

Persistence. Retain across future workflow families until a newer authoritative rule supersedes the record.

Scope. May guide actions beyond the source workflow when the supported proposition is cross-workflow invariant; cross-workflow authority must come from a Service Reliability Board standard.

Write transformation. Write a normalized cross-workflow rule with evidence provenance, source-record timestamp, cross-workflow scope, and operative authority; omit episode-specific wording.

Retrieval/action authority. May be retrieved by action selection across workflow families.

Model-visible prospective contract $\tilde{q}_1$.

Destination. An operative store of workflow-specific procedures.

Persistence. Retain for later instances of the named workflow family until that workflow's playbook changes.

Scope. May guide actions only inside the identified workflow family; procedural authority must come from the current playbook for that same workflow family.

Write transformation. Write a workflow procedure with its workflow family, evidence provenance, source-record timestamp, and within-family operative authority.

Retrieval/action authority. May be retrieved by action selection only within the named workflow family.

Evaluator-only relation, label, and frozen output. Both rows have $(A, D, F, S_{\tilde{q}}, R, T) = (1, 0, 1, 1, 0, \text{cross-rule})$. Under $\tilde{q}_0$, row `CE-32A1F9E04D0CDAF0F762` has construction label 1, frozen score 0.6937901055, and binary decision 0. Under $\tilde{q}_1$, row `CE-375C7BFD06319E558A7D` has construction label 0, frozen score 0.9996697108, and binary decision 1. The score ordering is therefore opposite to the reference direction.

EvidenceFlip example (surface-QC pair EF-P011). Model-visible prospective contract shared by both rows.

Destination. An operative store of cross-workflow invariant rules.

Persistence. Retain across future workflow families until a newer authoritative rule supersedes the record.

Scope. May guide actions beyond the source workflow when the supported proposition is cross-workflow invariant; cross-workflow authority must come from a Service Reliability Board standard.

Write transformation. Write a normalized cross-workflow rule with evidence provenance, source-record timestamp, cross-workflow scope, and operative authority; omit episode-specific wording.

Retrieval/action authority. May be retrieved by action selection across workflow families.

Model-visible feedback shared by both rows. Across service workflows, when the migration readiness flag is present, check the readiness flag before scheduling circuit migration.

Model-visible context, current-evidence member. Case file CASE-4110. Workflow: enterprise circuit migration. Active workflow registry code: WF-320. The service record names the migration readiness flag and records the operator action 'check the readiness flag before scheduling circuit migration'. 09:14 CASE-4110: migration readiness flag; 09:18 CASE-4110: action executed; 09:19 CASE-4110: scheduled circuit migration. Source record: SRB-OPS-27, issued by the Service Reliability Board. Exception EX-220 scope: legacy-prepaid; active account class: enterprise-standard. Selected evidence: REV-12 dated 2026-07-01; revision history: REV-11 to REV-12; active revision: REV-12. The case file ends with the ordinary operator notes and the referenced source records.

Model-visible context, superseded-evidence member. Case file CASE-4110. Workflow: enterprise circuit migration. Active workflow registry code: WF-320. The service record names the migration readiness flag and records the operator action 'check the readiness flag before scheduling circuit migration'. 09:14 CASE-4110: migration readiness flag; 09:18 CASE-4110: action executed; 09:19 CASE-4110: scheduled circuit migration. Source record: SRB-OPS-27, issued by the Service Reliability Board. Exception EX-220 scope: legacy-prepaid; active account class: enterprise-standard. Selected evidence: REV-11 dated 2026-02-01; revision history: REV-11 to REV-12; active revision: REV-12. The case file ends with the ordinary operator notes and the referenced source records.

Evaluator-only relation, label, and frozen output. The active conjunct is current evidence F. For row `CE-2D03470BB692D531B22B`, $F = 1$, the construction label is 1, the frozen score is 0.3339785584, and the binary decision is 0. For row `CE-9D73C0C75D9A404EBDDD`, $F = 0$, the construction label is 0, the frozen score is 0.3339269528, and the binary decision is 0. The frozen score ranks this pair in the reference direction, while the inherited threshold rejects both rows.

G.4 Complete 18-cell frozen results

The full ensemble pair-binary value is zero for both suites. The q-only ContractFlip pair-binary value is $1/3$, attaining the analytical ceiling. The q-masked and base-retrained ContractFlip rankings and no-c EvidenceFlip ranking are exactly 0.500 under the frozen broadcast controls.

Table 11: Post-hoc descriptive presentation of frozen ContractFlip cell results. Each cell contains eight pairs. Full is the frozen three-seed candidate; QM, Base, NC, and QO denote q-masked, base-retrained, no-c, and q-only controls. Values are pair-ranking credits.

Direction	Full	QM	Base	NC	QO
$\tilde{q}_0 > \tilde{q}_1$	0.000	0.500	0.500	0.000	0.000
$\tilde{q}_0 > \tilde{q}_2^\star$	0.000	0.500	0.500	0.000	0.000
$\tilde{q}_1 > \tilde{q}_0$	0.000	0.500	0.500	1.000	1.000
$\tilde{q}_1 > \tilde{q}_2^\star$	0.000	0.500	0.500	1.000	1.000
$\tilde{q}_2^\star > \tilde{q}_0$	1.000	0.500	0.500	1.000	1.000
$\tilde{q}_2^\star > \tilde{q}_1$	0.000	0.500	0.500	0.000	0.000
Cell macro	0.167	0.500	0.500	0.500	0.500

Table 12: Post-hoc descriptive presentation of frozen EvidenceFlip cell results. Each cell contains four pairs. Shuf. is the mean ranking across the five context derangements. Full-model pair-binary credit is zero in every cell.

Contract	Relation	Full	QM	Base	NC	QO	Shuf.
$\tilde{q}_0$	Current evidence	1.000	0.000	1.000	0.500	0.500	0.450
$\tilde{q}_0$	Disconfirmation	0.750	0.250	0.000	0.500	0.500	0.300
$\tilde{q}_0$	Observable anchor	0.000	0.500	0.000	0.500	0.500	0.400
$\tilde{q}_0$	Source applicability	0.000	0.000	0.000	0.500	0.500	0.600
$\tilde{q}_1$	Current evidence	0.000	0.000	1.000	0.500	0.500	0.450
$\tilde{q}_1$	Disconfirmation	1.000	0.250	0.000	0.500	0.500	0.450
$\tilde{q}_1$	Observable anchor	0.000	1.000	0.000	0.500	0.500	0.450
$\tilde{q}_1$	Source applicability	0.000	1.000	1.000	0.500	0.500	0.550
$\tilde{q}_2^\star$	Current evidence	1.000	0.000	1.000	0.500	0.500	0.650
$\tilde{q}_2^\star$	Disconfirmation	0.000	1.000	0.000	0.500	0.500	0.650
$\tilde{q}_2^\star$	Observable anchor	1.000	1.000	0.000	0.500	0.500	0.500
$\tilde{q}_2^\star$	Source applicability	0.000	1.000	1.000	0.500	0.500	0.450
Cell macro		0.396	0.500	0.417	0.500	0.500	0.492

Table 13: Frozen seed-level ranking results. The final column compares each full model with the mean of its five frozen shuffled-context evaluations on EvidenceFlip.

Model	ContractFlip	EvidenceFlip	EF full − shuffled-c
Three-seed ensemble	0.1667 $[0.1667, 0.1667]$	0.3958 $[0.3542, 0.4167]$	-0.0958 $[-0.1667, -0.0292]$
Seed 13	0.1667 $[0.1667, 0.1667]$	0.3125 $[0.2708, 0.3333]$	-0.2083 $[-0.2875, -0.1375]$
Seed 17	0.1667 $[0.1667, 0.1667]$	0.5312 $[0.4688, 0.5938]$	0.0542 $[-0.0229, 0.1354]$
Seed 23	0.5000 $[0.4375, 0.5625]$	0.3438 $[0.2917, 0.3958]$	-0.1104 $[-0.1833, -0.0396]$

H CASE STUDY

These cases were selected post hoc to illustrate the reported discrepancies; they do not estimate prevalence or revise the frozen evaluation. Case A is the unique shortest annotator-facing trace among the 12 Natural majority positives rejected by Span-PRM (16 events). Case B is the ledger group already discussed in Appendix Section A.5. Candidate text, contracts, labels, rationales, and decisions are copied from frozen records. Evidence excerpts mark omissions; identifiers were replaced before annotation. Scores are rounded only for display and are not calibrated probabilities; decisions come from full-precision outputs.

Case A: Natural majority-positive feedback rejected at the frozen operating point

Source and selection. Natural; public ID `NAT-9793FE448BBF461014AA`. Selected by the fixed shortest-trace rule above.

Source evidence c (excerpt). `E005` queries details for the requested line. `E006` returns `roaming_enabled: false` and `status: Active`; `E007` calls `enable_roaming` for that line; `E008`: "Roaming enabled successfully." `E011`: "Speed Test Result: 0.01 Mbps (Very Poor)." `E012` then changes the preferred network mode; `E015`: "Speed Test Result: 55.00 Mbps (Good)." [Events E001–E004, E009–E010, E013–E014, and E016 omitted.]

Candidate feedback z (complete). "Check roaming status before enabling roaming."

Update contract q_0 (complete). "Persistent cross-task policy: admit only an evidence-supported, current, causally and temporally justified rule that can safely govern future tasks outside the source workflow. The feedback is stored as an operative rule."

Human reference (frozen rationales). A: Eligible, 4/5. "E005–E008 identify the requested line, read its disabled roaming state, then enable that same line. This supports the bounded read-before-activation rule without claiming roaming alone resolves connectivity." **B: Ineligible, 4/5.** "Account roaming was checked before activation, but no outcome linked that ordering to improved connectivity." **C: Eligible, 4/5.** "E006 exposes the relevant line's disabled roaming status before E007–E008 enable it. This directly supports a targeted read-before-change procedure; the later speed shortfall concerns additional configuration and does not invalidate that check." Majority: **Eligible** (2:1).

Frozen model output.

Model	Score	Threshold	Decision
Base	0.549604	0.696544	Reject
q-masked	0.488910	0.456777	Admit
Span-PRM	0.004109	0.998764	Reject

Interpretation. Relative to the frozen majority reference, Span-PRM and the same-data base miss this positive; q-masked admits it. The case does not show that the feedback improves future tasks because the fresh suite has no retry outcome.

Case B: Challenge ledger disagreement across one three-contract group

Source and selection. Challenge; public group `FCG-34F9286277002FF4A2`. This is the pre-existing ledger example summarized in Appendix Section A.5.

Source evidence c (excerpt). `E002`: "No hidden evaluator message or downstream result was available." `E004`: "Across three independently named workflows, the record shows that the status panel confirms a completed prerequisite; handlers then used the same step to apply the step only inside this workflow before finalizing the case, and each visible verification confirmed it." [Events E001, E003, and E005–E007 omitted.]

Candidate feedback z (complete). "Before finalizing the case, apply the step only inside this workflow whenever the status panel confirms a completed prerequisite is visible."

Ledger contract q_2 (complete). "Non-operative evidence ledger: admit a relevant observation or explicitly unverified hypothesis for later investigation. Entries retain uncertainty and must never directly control an action; stale observations are not admitted." The exact q_0 and q_1 contracts appear in Appendix Section A.2.

Frozen labels. Construction is E for every contract. Human A/B/C $\rightarrow$ majority: q_0, E/I/I $\rightarrow$ I; q_1, E/E/E $\rightarrow$ E; q_2, E/I/I $\rightarrow$ I.

q_2 human reference (frozen rationales). A: Eligible, 5/5. "E004 records the prerequisite-triggered, workflow-bounded step before finalization with verification across three workflows. These current observations fit the non-operative ledger." **B: Ineligible, 5/5.** "The candidate is an operative workflow procedure, which the non-operative evidence ledger cannot store." **C: Ineligible, 5/5.** "E004 verifies a workflow-bounded step, but the candidate instructs applying it before case finalization. Its operative command cannot be stored in a ledger whose entries must remain observational or investigative and never control action." Majority: **Ineligible** (2:1).

Frozen model output. Rounded score / full-precision decision (Ad = admit; Rej = reject):

Model	q_0	q_1	q_2
Base	0.730/Ad	0.730/Ad	0.730/Ad
q-masked	0.493/Ad	0.493/Ad	0.493/Ad
Span-PRM	0.313/Rej	0.334/Rej	1.000/Ad

Interpretation. A treats the supported step as suitable evidence for non-operative storage; B/C treat the directive form as target-incompatible. The rationales are consistent with different assumptions about the write operation, but this case does not establish the cause of disagreement or invalidate either judgment.