Third Eye: Beyond Immediate Gains in Recursive Self-Improvement

Sujato Dutta **Padala Indira Sneha** **Kypu Mahitha**
Shashank Vanga **Neeraj Choudhary**
Mahindra University

arXiv:2611.12988v1 [cs.LG] 8 Nov 2026

Abstract

Does a useful self-update remain useful after the model updates again? Third Eye separates immediate performance from the outcome of a further native update. We evaluate five 3–8B backbones on math and code: 64 fresh runs, three candidates per run, and three continuations per candidate. Selection and final scoring use disjoint items. In the balanced Qwen3-4B/Llama-3.2-3B comparison, immediate-selected utility is near zero, but continuation reduces it by $1.78\,\mathrm{pp}$ (simultaneous 95% interval $[0.68, 2.88]$). Candidate-averaged immediate utility is already negative: selection partly accounts for the apparently neutral first step. Continuation also lowers utility for SmolLM3-3B and Llama-3.1-8B; Qwen3-8B remains inconclusive. Target and out-of-distribution losses dominate retention changes. Matched recipe and recipient–donor controls identify neither a reliable mitigation nor a single mechanism. The results support evaluating persistent self-updates as adaptation transitions, with independent scoring and a no-update baseline, rather than interpreting immediate performance as their full value.

1 Introduction

Language models increasingly generate the data used to update their own parameters (Zelikman et al., 2022; Huang et al., 2022; Yuan et al., 2024; Zweiger et al., 2025). In a recursive loop, an update changes both the model's current answers and the model that will supply and learn from its next training batch. An immediate evaluation measures only the current checkpoint.

Verification addresses a different question. Answer checks and execution tests can reject incorrect corrections (Cobbe et al., 2021; Lightman et al., 2023; Liu et al., 2023), but passing examples do not establish the broader value of training on them. That value may also change after another update.

SEAL learns self-edits using post-adaptation feedback (Zweiger et al., 2025), while modification screening checks protected behavior (Tayebati et al., 2026). We examine the next native update, where persistent parameter changes also affect the generation of new training data.

Third Eye measures both horizons (Figure 1). It constructs three equal-budget LoRA candidates from a common parent, then follows each through three native continuations. Development items determine candidate selection; independent items score all checkpoints against the unchanged parent. The five-model, two-task study contains 64 fresh runs, 192 candidates, and 576 continuations.

In the balanced Qwen3-4B/Llama-3.2-3B comparison, immediate-selected utility is $+0.02\,\mathrm{pp}$ and continued utility is $-1.76\,\mathrm{pp}$. Their paired difference has an adjusted interval excluding zero. However, candidate-averaged immediate utility is already negative: selection partly explains the neutral first-step mean. SmolLM3-3B and Llama-3.1-8B instead show immediate damage followed by further decline; Qwen3-8B's continuation effect remains unresolved. The evidence concerns different assessments across horizons, with substantial model dependence.

We contribute:

1. A paired protocol with independent scoring, repeated continuations, and an explicit no-update baseline.

2. Five-model evidence separating selection effects and target/OOD/retention changes across horizons.

3. Matched recipe and recipient–donor controls that bound explanations and proposed mitigations.

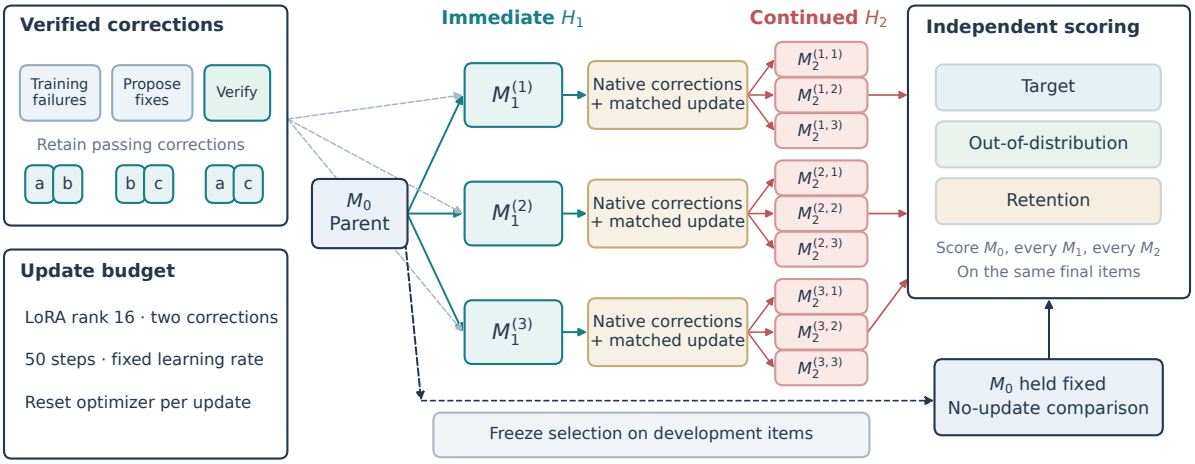

Figure 1: Third Eye follows every candidate through another self-update. Development items select candidates; independent items score both horizons against the unchanged parent.

2 Measuring a Self-Update Beyond Its First Step

2.1 Verified data and matched branches

A parent model M_0 first attempts training problems. Failed problems receive up to eight correction attempts each. Only solutions that pass the task verifier enter the correction pool. Math corrections are checked against answers; code corrections must pass execution tests in an isolated sandbox. We construct $K = 3$ compositionally distinct batches B_k, each containing two different problems. Batches may overlap with one another. This preserves candidate diversity without requiring six different verified corrections.

Let A denote the fixed update procedure. Candidate k is

$$M_1^{(k)} = A(M_0, B_k; \xi_k). \tag{1}$$

The parent is restored before each candidate is trained. For each draw $d \in \{1, 2, 3\}$, candidate $M_1^{(k)}$ generates a native correction batch $C_{k,d}$ and undergoes one further update:

$$M_2^{(k,d)} = A(M_1^{(k)}, C_{k,d}; \zeta_{k,d}). \tag{2}$$

All candidates continue, including those not chosen by the immediate selector. Continuation draws vary the correction/optimization procedure jointly; they do not separately randomize each source of variation.

Updates use LoRA (Hu et al., 2021), rank 16 and scaling 32, on attention projections. Each full update has 50 AdamW steps, learning rate 10^{-4}, microbatch size one, and gradient accumulation eight. Thus two corrections receive 400 example presentations per update. Optimizer state is reset between updates, while the candidate's parameter changes persist into its continuation. If a native correction pool is exhausted without supplying a full batch, the branch is terminal: $M_2 = M_1$, with no continuation training. None of the 576 fresh continuation draws was terminal.

2.2 Immediate and continued utility

For role $r \in \{\text{target}, \text{OOD}, \text{retention}\}$, let $a_r(M)$ be accuracy on a fixed item set. Utility is the equal-role mean change from the parent:

$$U(M; M_0) = \frac{1}{3} \sum_r \big[a_r(M) - a_r(M_0)\big]. \tag{3}$$

We report percentage points, multiplying this quantity by 100. Equal role weights prevent the larger retention set from dominating. They remain a specified objective, rather than a universal measure of quality.

The immediate selector maximizes development utility. For run p, let w_{pk} assign equal mass to tied best candidates and zero elsewhere. On independent scoring items, its two endpoints are

$$G_{1p} = \sum_k w_{pk} U(M_1^{(k)}; M_0), \tag{4}$$

$$G_{2p} = \sum_k w_{pk} \frac{1}{3} \sum_{d=1}^{3} U(M_2^{(k,d)}; M_0). \tag{5}$$

Their paired difference, $\Delta_p = G_{2p} - G_{1p}$, measures the increment from continuation. Both endpoints use the same parent reference. Here H_2

Role	Math	Code
Training	GSM8K: 1,024	MBPP: 306
Dev. target	GSM8K: 64	MBPP: 64
Dev. OOD	MATH: 64	MBPP val.: 64
Final target	GSM8K: 128	MBPP test: 128
Final OOD	MATH-500: 128	HumanEval: 128
Retention	MMLU: 256 in each suite	

Table 1: Fixed evaluation suites. Development and final items are disjoint. Retention is measured on separate development and final subsets.

means an observed two-update outcome, not a prediction.

We also report uniform-candidate endpoints, replacing w_{pk} by $1/3$. This separates the behavior of the update procedure from the effect of selecting a candidate on development performance. HOLD retains M_0 and has zero utility by definition. REPLAY admits the immediate-selected candidate only if its development target gain is positive and both OOD and retention gains are nonnegative. These actions are recorded before final scoring. Holding does not consume an equal training budget; it is the decision alternative to updating.

2.3 Models, tasks, and experimental units

We evaluate Qwen3-4B and Qwen3-8B (Yang et al., 2025), Llama-3.2-3B-Instruct (Meta AI, 2024), Llama-3.1-8B-Instruct (Grattafiori et al., 2024), and SmolLM3-3B (Bakouch et al., 2025). The three smaller backbones have eight fresh runs per task; each 8B backbone has four. This gives 64 runs, balanced between math and code within every model. A run repeats the adaptation experiment from the same released checkpoint with a different random seed; it is not an independently pretrained model.

Table 1 specifies the data. GSM8K (Cobbe et al., 2021) and MBPP (Austin et al., 2021) provide training failures and target evaluation. MATH (Hendrycks et al., 2021) and HumanEval (Chen et al., 2021) provide final OOD problems; MMLU (Hendrycks et al., 2020) measures retention. Code's development OOD set is MBPP validation, whereas final OOD scoring uses HumanEval. This domain change is relevant when interpreting admission. Eight HumanEval problems used for timing calibration are excluded from final scoring.

Every run scores 13 checkpoints: the parent, three candidates, and nine continuations. The study therefore contains 832 checkpoint evaluations. Checkpoints and draws are repeated measurements within a run, not additional independent runs. Hash-locked manifests, templates, actions, and checkpoint identities link these measurements. We freeze the date in model templates, preventing a changing system date from changing prompts across checkpoints. Appendix A documents settings and exact revisions.

2.4 Uncertainty and declared comparisons

We use 5,000 bootstrap replicates, resampling parent runs within model/task strata and items within evaluation roles. Checkpoints sharing an exact item pool share its item resample. This retains pairing and avoids treating reuse of the same test questions as independent evidence (Bouthillier et al., 2021; Dror et al., 2018; Berg-Kirkpatrick et al., 2012). The three native continuation draws are averaged within a candidate before resampling; they are not separately resampled. Intervals are consequently conditional on these finite draws and the frozen selector.

Declared tests and descriptive comparisons are kept separate. For Qwen3-4B/Llama-3.2-3B, the primary is the model difference in G_2, with a joint 95% interval. The pooled continuation contrast belongs to a six-contrast secondary family, with simultaneous 95% intervals and Holm-adjusted centered-bootstrap p-values (Holm, 1979). SmolLM3's declared primary is its paired continuation contrast, with a 95% interval. Each 8B model has the same endpoint with a 97.5% interval, giving joint 95% coverage for the two-model test family. Each declaration was frozen before its corresponding outcomes. All five-model plots and role/task decompositions are descriptive; they do not create a new pooled confirmatory primary.

3 What Changes After Another Update?

3.1 A neutral immediate mean can precede a decline

In the 32-run Qwen3-4B/Llama-3.2-3B comparison, immediate-selected utility averages $+0.016$ pp (95% interval $[-1.319, 1.326]$). After continuation, utility is -1.760 pp ($[-3.105, -0.420]$). The paired decrement is -1.776 pp, with simultaneous interval $[-2.876, -0.679]$ and Holm-adjusted $p = 0.0012$. The same selected candidates therefore receive materially different assessments at the two horizons.

The near-zero immediate mean is not an equiv-

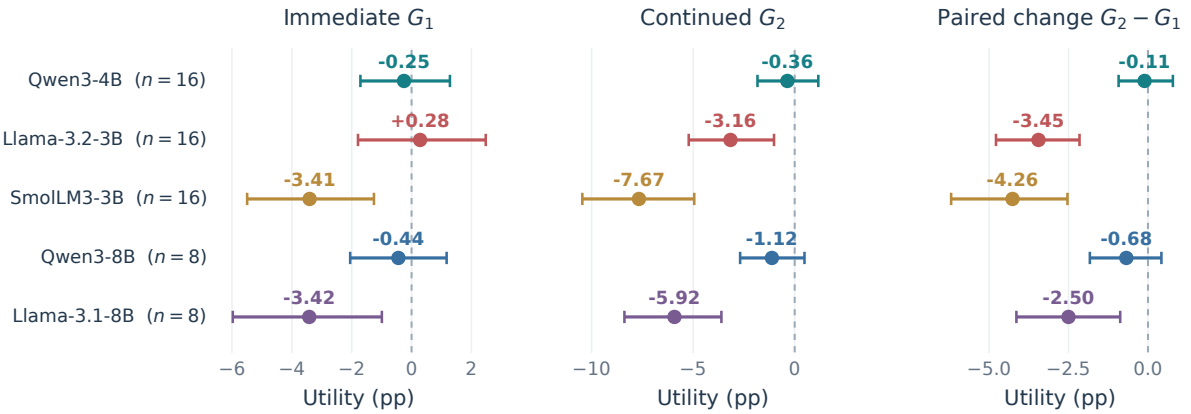

Figure 2: Immediate and continued utility across five models. Bars are descriptive joint run/item 95% intervals; adjusted intervals for declared contrasts appear in Table 3. The paired change uses within-run differences, rather than differences between interval endpoints.

alence test: its interval permits both benefit and harm. It also combines different model responses. Llama-3.2-3B has immediate utility $+0.285$ pp and continued utility -3.160 pp; Qwen3-4B has -0.252 pp and -0.359 pp. Most of the pooled continuation decrement comes from Llama. The predeclared Qwen-minus-Llama G_2 contrast is $+2.801$ pp ($[0.266, 5.329]$). This is evidence of different outcomes under a common recipe, not evidence that Qwen adaptation is beneficial: its mean remains slightly negative.

3.2 Selection accounts for part of the first-step assessment

Uniform candidate weights change the interpretation of the first step (Table 2). Across the same 32 runs, candidate-averaged immediate utility is -0.686 pp, rather than the selector's $+0.016$ pp. Candidate-averaged continued utility is -2.059 pp. Immediate selection thus raises both endpoints, especially the first, but does not eliminate the continuation decrement.

This control limits the claim that immediate evaluation simply misses all update damage. Some damage is already visible before selection, and SmolLM3 and Llama-3.1-8B have negative immediate outcomes even after selection. What the paired design establishes is that the immediate endpoint alone does not describe the full update procedure.

3.3 Continuation effects differ across backbones

Figure 2 presents the full model matrix. SmolLM3-3B declines by 4.262 pp after continuation, with its

	Selected		Uniform	
Model	G_1	G_2	G_1	G_2
Qwen3-4B	-0.25	-0.36	-0.50	-0.41
Llama-3.2-3B	$+0.28$	-3.16	-0.87	-3.71
SmolLM3-3B	-3.41	-7.67	-3.65	-8.02
Qwen3-8B	-0.44	-1.12	-0.75	-1.55
Llama-3.1-8B	-3.42	-5.92	-4.20	-7.56

Table 2: Mean utility in percentage points. Uniform weights average all three candidates, removing development-based candidate selection.

declared 95% interval $[-6.163, -2.581]$. Llama-3.1-8B declines by 2.501 pp, with adjusted 97.5% interval $[-4.411, -0.583]$. Neither replicates the near-neutral first step: their immediate means are -3.406 pp and -3.418 pp, respectively. They support an additional continuation cost under this recipe, while showing that delayed loss need not be hidden initially.

Qwen3-8B's mean continuation change is -0.681 pp, with adjusted interval $[-1.914, 0.629]$. This result is inconclusive. Its interval admits both a meaningful decline and a modest improvement. The five-model pattern is therefore neither uniform degradation nor a monotonic size effect. Only Qwen3-8B was executed on A100 hardware; the remaining fresh runs used GH200 hardware. Model size, family, training history, and runtime are not independently randomized.

3.4 Target and OOD loss dominate retention loss

Figure 3 decomposes the continuation increment by role. For Llama-3.2-3B, target accuracy falls by 6.45 pp and OOD accuracy by 3.88 pp, while re-

Declared contrast	Mean	Interval
Qwen4B − Llama3B, G_2	$+2.80$	$[+0.27, +5.33]$
Pooled Δ, these models	-1.78	$[-2.88, -0.68]$
SmolLM3-3B, Δ	-4.26	$[-6.16, -2.58]$
Qwen3-8B, Δ	-0.68	$[-1.91, +0.63]$
Llama-3.1-8B, Δ	-2.50	$[-4.41, -0.58]$

Table 3: Declared contrasts, in percentage points. Coverage is 95% for the first and third rows, simultaneous 95% for the second, and 97.5% per model for the two 8B rows. These are separate declared test families.

tention changes by only -0.01 pp. SmolLM3-3B shows target/OOD changes of $-6.46/-6.00$ pp, versus -0.33 pp in retention. Llama-3.1-8B shows the same ordering, with $-4.30/-2.38/-0.83$ pp for target/OOD/retention. Thus the observed loss is not adequately described as forgetting unrelated MMLU knowledge. It also concerns the capabilities adaptation is meant to serve.

Qwen3-4B differs: target declines by 0.74 pp, OOD improves by 0.79 pp, and retention declines by 0.37 pp. These are descriptive role means, not three independently established effects. An exploratory model-by-task interaction in the 32-run comparison is $+2.07$ pp for the continuation contrast, with 95% interval $[-0.98, 5.11]$; it does not resolve a task-specific explanation. Appendix B gives both horizons for every role.

Formatting also changes. On GSM8K, strict numeric-extraction failures rise from 4.49% at H_1 to 10.19% at H_2 for Llama-3.2-3B, and from 4.69% to 15.12% for SmolLM3. These are observable failure modes within the measured loss. They do not establish how much loss reflects formatting rather than problem-solving ability. We preserve the original scoring and do not repair answers after seeing outcomes. Appendix E reports syntax, extraction, and length diagnostics.

Run-level signs also differ (Figure 4). Seven of nine Llama-3.2-3B runs with positive measured G_1 have negative measured G_2; none of Qwen3-4B's eight positive-G_1 runs does. SmolLM3 and Llama-3.1-8B have no positive-G_1 runs in this sample. These observed signs describe finite-item outcomes, not noise-corrected individual effects.

4 Controls on the Update Procedure

The main comparison couples the receiving parameters, the corrections they generate, and subsequent optimization. We examine these factors in 16 paired follow-up parents, four per Qwen3-4B/Llama-3.2-3B task stream. They reuse inspected parents and are not additional fresh replications. One fixed candidate branch and two draws per parent yield 272 scored cells. Native baselines and controls are re-scored on the same A100 runtime. The nine declared control contrasts use joint run/item resampling, simultaneous 95% intervals, and Holm correction.

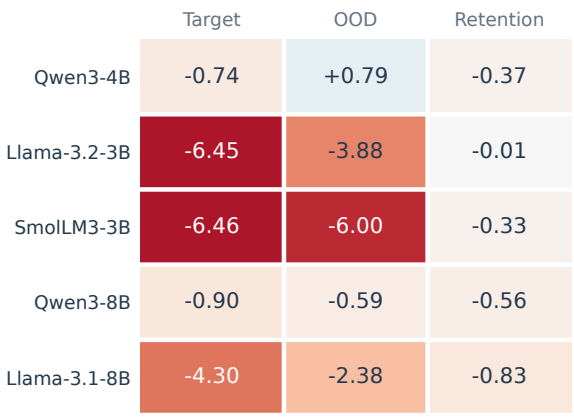

Figure 3: Role-wise continuation changes, in percentage points. Target and OOD losses dominate in the backbones with the largest mean declines. Colors show descriptive means, not significance.

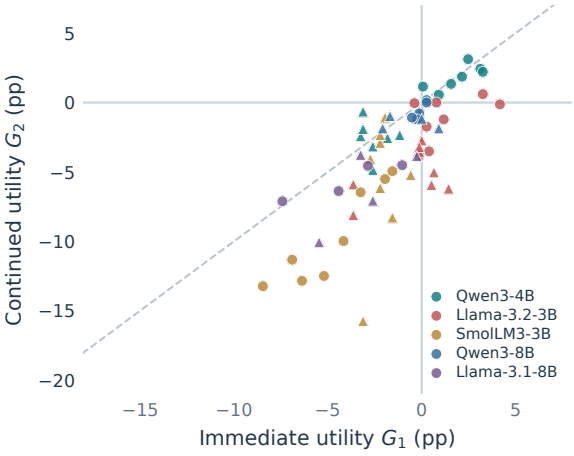

Figure 4: Each point is one fresh run; circles denote code and triangles math. The diagonal marks unchanged utility after continuation. Observed signs are finite-item estimates, not noise-corrected labels.

4.1 Fewer steps and pooled corrections do not give a resolved rescue

To test the intensity of the two-example recipe, we compare native 50-step continuations with 10 steps, 25 steps, and the union of verified corrections from that candidate's archived continuation batches. The alternatives change the receiving model's continua-

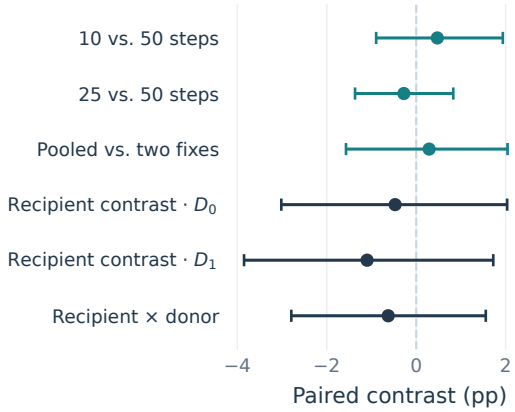

Figure 5: Matched recipe and recipient–donor contrasts. Bars are simultaneous 95% intervals for the nine-control family. No displayed intervention resolves an improvement or a single mechanism.

tion recipe; they do not redefine the main study's training budget.

Relative to 50 steps, the 10-step contrast is $+0.47$ pp ($[-0.90, 1.94]$), and the 25-step contrast is -0.28 pp ($[-1.37, 0.83]$). Pooling yields $+0.29$ pp ($[-1.57, 2.04]$). All intervals cross zero (Figure 5). These controls neither establish a mitigation nor prove recipe invariance: they remain compatible with useful improvements. The native continuation increment in this smaller subset is -1.20 pp ($[-3.06, 0.55]$), itself inconclusive.

4.2 Crossed recipients and donors leave the mechanism open

We apply correction batches from the unchanged parent (D_0) and the first-step model (D_1) to both receiving parameter states. Comparing receivers on the same donor data tests whether receiving state changes the update increment. The recipient contrasts are -0.47 pp for D_0 ($[-3.02, 2.03]$) and -1.10 pp for D_1 ($[-3.85, 1.72]$). Their interaction is -0.63 pp ($[-2.79, 1.56]$). The intervals do not isolate parameter interference, endogenous data change, or their interaction. Applying the same native recipe directly to M_0 gives -0.44 pp ($[-2.40, 1.36]$). Thus these controls do not establish that the additional loss requires the first update, rather than another application of this small-batch recipe.

The donor comparison also has a specific limit: D_0 was harvested on the control runtime, while D_1 was archived. Its source contrast includes generation/runtime differences alongside correction composition. The crossed design is informative about these matched updates, but it does not justify a causal decomposition of all native continuation loss.

4.3 Reference solutions expose another recipe dependence

A matched reference-solution arm replaces verified self-generated solutions with gold solutions for the same training problems. Under the same 50-step recipe, it is 11.56 pp worse than the self-generated arm (simultaneous interval $[-14.45, -8.64]$). This large difference is not evidence that correct reference solutions are intrinsically inferior.

A saved-output audit confirms identical receiving states, training seeds, problem identities, and step budgets. Yet gold solutions contain only 41–61% as many assistant-label tokens per example. Their resulting completions are shorter, and code syntax-failure rates increase. For Llama code, target syntax failures rise from 1.66% to 18.55%; for Qwen, from 0.98% to 4.88%. Equal example presentations therefore do not match text style or label-token exposure. These observations motivate caution about data substitution, rather than a claim that syntax or exposure mediates the entire utility loss.

5 Implications for Evaluating Self-Improvement

Measure transitions as well as checkpoints. The result is not that immediate evaluation is useless. It detects substantial first-step damage in several backbones, and candidate selection improves the immediate mean. Its limit is that a current score omits the subsequent update transition. A useful evaluation should report G_1, G_2, and their paired difference, with an unchanged parent as a decision baseline. The role breakdown matters: a retention-only gate would overlook much of the target/OOD loss observed here.

Abstention is a baseline, not a demonstrated solution. In the 32-run comparison, REPLAY commits 21.9% of the selected candidate mass and yields -0.27 pp ($[-0.80, 0.21]$) relative to holding. It does not establish improvement over HOLD. A fixed development-derived Qwen-commit/Llama-hold rule likewise fails to outperform holding or replay. Better performance than forced updating is not sufficient evidence of a useful admission policy.

Selection and uncertainty require separate controls. Selecting and scoring the same outcome

creates optimism (Cawley and Talbot, 2010). An eight-parent diagnostic shows same-outcome selection optimism increasing from zero at $K = 1$ to $2.09\,\mathrm{pp}$ at $K = 3$ (Appendix D). Our main study instead freezes selection on separate items. Uncertainty still depends on shared final questions: for the Qwen–Llama primary, the run-only interval is $[1.85, 3.78]\,\mathrm{pp}$, compared with the joint interval $[0.27, 5.33]\,\mathrm{pp}$. Repeated checkpoints cannot substitute for broader evaluation support.

Prediction remains a distinct empirical claim. Training-response forecasting (Luo et al., 2026; Liu et al., 2026) and modification screening (Ji and Ding, 2026; Xie et al., 2026) address useful but different decisions. The present results establish a measurement contrast; they do not establish that a forecaster can exploit it. Earlier held-out forecasting and admission diagnostics did not yield a validated advantage and are not the basis of the five-model claims. No learned predictor is fitted to, or selected using, these fresh outcomes.

6 Related Work

Self-generated adaptation and agent improvement. STaR (Zelikman et al., 2022) and self-improvement with generated reasoning (Huang et al., 2022) show how model outputs can become training data. Self-Rewarding Language Models (Yuan et al., 2024) extend this idea to preference supervision. SEAL (Zweiger et al., 2025) learns to generate self-adaptation data and directives using downstream performance after training. Our study does not introduce self-generated parameter adaptation; it measures how a candidate's immediate value differs from its value after another native update. Self-Refine (Madaan et al., 2023), Reflexion (Shinn et al., 2023), and analyses of self-correction (Huang et al., 2023) concern refinement and feedback that need not persist as weight changes. The Darwin Gödel Machine (Zhang et al., 2025) evolves agent programs. These settings motivate recursive evaluation but differ from the fixed LoRA branches studied here.

Predicting and screening modifications. TuneAhead (Luo et al., 2026) predicts full fine-tuning outcomes from data descriptors and short probes. Target-independent micro-interventions (Liu et al., 2026) characterize training response across models. Rehearse (Ji and Ding, 2026) and Predict Before You Modify (Xie et al., 2026) use prior modification outcomes to guide agent changes. The Self-Healing Harness (Tayebati et al., 2026) screens persistent instruction changes against triggering failures and protected behavior. Thus beneficial local edits with collateral costs are already documented. Our distinction is the paired immediate/continued evaluation of persistent weight updates, with native correction regeneration and independent scoring.

Forgetting and adaptation geometry. Continual-learning methods constrain important parameters or replay past experience (Kirkpatrick et al., 2016; Zenke et al., 2017; Lopez-Paz and Ranzato, 2017). Forgetting during language-model fine-tuning (Luo et al., 2023), distorted pretrained features (Kumar et al., 2022), and LoRA-specific intruder dimensions (Shuttleworth et al., 2024) offer plausible explanations for adaptation failures. LoRA (Hu et al., 2021) and QLoRA (Dettmers et al., 2023) make parameter-efficient experiments accessible, but efficiency does not establish update quality. We do not measure spectral mediation or test a continual-learning remedy; our recipient–donor and recipe controls leave the mechanism unresolved.

Verification and credible evaluation. Outcome verification (Cobbe et al., 2021), process supervision (Lightman et al., 2023), and stronger code testing (Liu et al., 2023) improve supervision. They answer different questions from the broader effect of a parameter update. Evaluation reproducibility (Melis et al., 2017), model-selection bias (Cawley and Talbot, 2010), benchmark variation (Bouthillier et al., 2021), and paired statistical testing (Dror et al., 2018; Berg-Kirkpatrick et al., 2012) motivate our independent scoring and shared-item resampling. Third Eye applies these safeguards to the adaptation transition, where candidates and descendants reuse both ancestors and evaluation items.

7 Conclusion

Third Eye evaluates a verified self-update immediately and after a further native update. The five-model study shows why these endpoints should be reported together: continuation can worsen utility substantially, while the size and visibility of that change depend on the backbone. Candidate selection explains part of an apparently neutral first step, and target/OOD losses can dominate retention changes. The matched controls provide neither a

universal mechanism nor a reliable repair. The practical contribution is a controlled way to measure update consequences before interpreting immediate gains as recursive improvement.

Limitations

The evidence covers five 3–8B backbones, two tasks, two-example LoRA updates, and H_2. Transfer to larger batches, full fine-tuning, or longer loops is untested. Fifty steps on two corrections may over-specialize; the controls cannot rule out better recipes. Eight runs per 8B model provide less precision. Hardware, checkpoint training history, and parameter count are not isolated causal factors.

Fixed item sets offer finite coverage and may overlap pretraining data. Intervals condition on three averaged continuations, the development selector, and the equal-role objective. Code's development/final OOD domains differ. Observed sign changes are noisy estimates. The 16-parent controls reuse inspected parents; their reference-source audit does not establish a mediator.

Ethical Considerations

We use public math, code, and general-knowledge benchmarks, with no human participants or private personal data. Generated code runs in resource-limited isolated sandboxes. Models and datasets retain their respective license requirements. Benchmark verification cannot certify unattended self-modification. We report adverse and unresolved outcomes, and do not interpret nonsignificance as evidence of safety.

References

Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, et al. 2021. Program Synthesis with Large Language Models. *arXiv preprint arXiv:2108.07732*.

Elie Bakouch, Carlos Miguel Patiño, Anton Lozhkov, Edward Beeching, Aymeric Roucher, Nouamane Tazi, et al. 2025. SmolLM3: Smol, Multilingual, Long-context Reasoner. Hugging Face technical blog.

Taylor Berg-Kirkpatrick, David Burkett, and Dan Klein. 2012. An empirical investigation of statistical significance in NLP. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 995–1005.

Xavier Bouthillier, Pierre Delaunay, Mirko Bronzi, Assya Trofimov, Brennan Nichyporuk, Justin Szeto, et al. 2021. Accounting for Variance in Machine Learning Benchmarks. *arXiv preprint arXiv:2103.03098*.

Gavin C. Cawley and Nicola L. C. Talbot. 2010. On over-fitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research*, 11:2079–2107.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, et al. 2021. Evaluating Large Language Models Trained on Code. *arXiv preprint arXiv:2107.03374*.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, et al. 2021. Training Verifiers to Solve Math Word Problems. *arXiv preprint arXiv:2110.14168*.

Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. QLoRA: Efficient Finetuning of Quantized LLMs. *arXiv preprint arXiv:2305.14314*.

Rotem Dror, Gili Baumer, Segev Shlomov, and Roi Reichart. 2018. The hitchhiker's guide to testing statistical significance in natural language processing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1383–1392.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, et al. 2024. The Llama 3 Herd of Models. *arXiv preprint arXiv:2407.21783*.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring Massive Multitask Language Understanding. *arXiv preprint arXiv:2009.03300*.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring Mathematical Problem Solving With the MATH Dataset. *arXiv preprint arXiv:2103.03874*.

Sture Holm. 1979. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2):65–70.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. LoRA: Low-Rank Adaptation of Large Language Models. *arXiv preprint arXiv:2106.09685*.

Jiaxin Huang, Shixiang Shane Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. 2022. Large Language Models Can Self-Improve. *arXiv preprint arXiv:2210.11610*.

Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. 2023. Large Language Models Cannot Self-Correct Reasoning Yet. *arXiv preprint arXiv:2310.01798*.

Jiazhen Ji and Shouhong Ding. 2026. Rehearse: Stepping Back from the Confidence Cliff in Self-Improving Autoresearch. *arXiv preprint arXiv:2607.27687*.

James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, et al. 2016. Overcoming catastrophic forgetting in neural networks. *arXiv preprint arXiv:1612.00796*.

Ananya Kumar, Aditi Raghunathan, Robbie Jones, Tengyu Ma, and Percy Liang. 2022. Fine-Tuning can Distort Pretrained Features and Underperform Out-of-Distribution. *arXiv preprint arXiv:2202.10054*.

Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, et al. 2023. Let's Verify Step by Step. *arXiv preprint arXiv:2305.20050*.

Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. 2023. Is Your Code Generated by ChatGPT Really Correct? Rigorous Evaluation of Large Language Models for Code Generation. *arXiv preprint arXiv:2305.01210*.

Zhongxuan Liu, Sicheng Zhou, and Hongzhi Wang. 2026. Target-Independent Micro-Interventions for Predicting Training Response Across Language-Model Families. *arXiv preprint arXiv:2609.08618*.

David Lopez-Paz and Marc'Aurelio Ranzato. 2017. Gradient Episodic Memory for Continual Learning. *arXiv preprint arXiv:1706.08840*.

Yun Luo, Zhen Yang, Fandong Meng, Yafu Li, Jie Zhou, and Yue Zhang. 2023. An Empirical Study of Catastrophic Forgetting in Large Language Models During Continual Fine-tuning. *arXiv preprint arXiv:2308.08747*.

Yuxiang Luo, Haonan Long, Chen Wang, Qiqi Duan, Xiaotian Lin, Yanwei Xu, et al. 2026. TuneAhead: Predicting Fine-tuning Performance Before Full Training Begins. *arXiv preprint arXiv:2606.17660*.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, et al. 2023. Self-Refine: Iterative Refinement with Self-Feedback. *arXiv preprint arXiv:2303.17651*.

Gábor Melis, Chris Dyer, and Phil Blunsom. 2017. On the State of the Art of Evaluation in Neural Language Models. *arXiv preprint arXiv:1707.05589*.

Meta AI. 2024. Llama 3.2 Model Card. Model documentation.

Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language Agents with Verbal Reinforcement Learning. *arXiv preprint arXiv:2303.11366*.

Reece Shuttleworth, Jacob Andreas, Antonio Torralba, and Pratyusha Sharma. 2024. LoRA vs Full Fine-tuning: An Illusion of Equivalence. *arXiv preprint arXiv:2410.21228*.

Sina Tayebati, Divake Kumar, Nastaran Darabi, Ranganath Krishnan, and Amit Ranjan Trivedi. 2026. Self-Healing Harness for Runtime Oversight of Agent Self-Modification. *arXiv preprint arXiv:2609.24130*.

Zixuan Xie, Chang Xu, Yuhua Jiang, and Li Zhao. 2026. Predict Before You Modify: Learning Improvement Dynamics for Self-Evolving Agents. *arXiv preprint arXiv:2610.13991*.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, et al. 2025. Qwen3 Technical Report. *arXiv preprint arXiv:2505.09388*.

Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Xian Li, Sainbayar Sukhbaatar, Jing Xu, and Jason Weston. 2024. Self-Rewarding Language Models. *arXiv preprint arXiv:2401.10020*.

Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. 2022. STaR: Bootstrapping Reasoning With Reasoning. *arXiv preprint arXiv:2203.14465*.

Friedemann Zenke, Ben Poole, and Surya Ganguli. 2017. Continual Learning Through Synaptic Intelligence. *arXiv preprint arXiv:1703.04200*.

Jenny Zhang, Shengran Hu, Cong Lu, Robert Lange, and Jeff Clune. 2025. Darwin Godel Machine: Open-Ended Evolution of Self-Improving Agents. *arXiv preprint arXiv:2505.22954*.

Adam Zweiger, Jyothish Pari, Han Guo, Ekin Akyürek, Yoon Kim, and Pulkit Agrawal. 2025. Self-Adapting Language Models. *arXiv preprint arXiv:2506.10943*.

A Implementation and Frozen Settings

Study coverage. Table 4 summarizes the model/task matrix. All models use unquantized bfloat16 arithmetic.

Model	Code / math runs	GPU
Qwen3-4B	8 / 8	GH200
Llama-3.2-3B	8 / 8	GH200
SmolLM3-3B	8 / 8	GH200
Qwen3-8B	4 / 4	A100
Llama-3.1-8B	4 / 4	GH200

Table 4: Study coverage. Counts denote fresh adaptation runs from a fixed released checkpoint, not independent pretraining runs.

Optimization. We use unquantized bfloat16 models with attention query, key, value, and output LoRA projections; rank 16, alpha 32, and zero adapter dropout. The learning rate is 10^{-4}, with 50 AdamW steps, microbatch one, gradient accumulation eight, and maximum training sequence length 1,024. The two-example batch is repeatedly presented, giving 400 presentations. The optimizer is reinitialized between updates. Parent restoration and adapter hashes prevent candidate branches from inheriting one another. The exact optimizer and tokenizer configurations accompany the experiment manifests.

Generation and scoring. Generation is capped at 512 new tokens. Evaluation uses greedy pass@1, with conditional-likelihood scoring for multiple-choice retention. Generation batches are 32 for greedy decoding and 16 for sampling. Reasoning-capable checkpoints use non-thinking mode. Training failure harvesting scans at most 1,024 GSM8K or 306 MBPP problems, with eight correction attempts per failed problem. Only verified examples are retained. Scoring uses separate immutable development and final manifests. Prompt and item hashes are checked for reuse; this does not audit external pretraining corpora.

Scarce corrections. Candidate batches must differ in composition and contain no repeated problem within a batch. Sharing a correction across candidate batches is permitted. Candidates have equal batch size and optimizer budget. A continuation with fewer than two verified corrections after its fixed harvest budget preserves the incoming adapter exactly and records `continuation_available=0`. No synthetic repetition is used to force an update. All fresh study draws had available continuations.

Frozen prompts. A template's automatically inserted date can otherwise change between model evaluations. We fix the date string to `11 Oct 2026` and bind template and prompt hashes to the scoring manifest. Historical runs with changing templates are excluded from the fresh five-model estimates.

Checkpoint revisions. Qwen3-8B is pinned to `b968826d9c46`. SmolLM3-3B is pinned to `a07cc9a04f16`. The Llama-3.1-8B-Instruct checkpoint uses the NousResearch mirror at revision `0e9e39f249a1`, with recorded upstream revision `d10aef7999a2`. These are abbreviated revision identifiers. Full revisions and file hashes are included in the machine-readable configurations.

B Role and Task Decomposition

The tables use compact model labels: Qwen4B and Qwen8B denote Qwen3; Llama3B denotes Llama-3.2-3B, Llama8B denotes Llama-3.1-8B, and Smol3B denotes SmolLM3-3B.

The decompositions are exploratory. Their ordinary 95% intervals are not adjusted for every displayed role, task, and model comparison. The model/task interaction reported in the main text is a difference of paired continuation changes, not an experimentally identified causal effect of belonging to a model family.

Target accuracy. The target head evaluates the task distribution used to harvest corrections.

Model	G_1	G_2	Δ	95% CI
Qwen4B	$+1.64$	$+0.90$	-0.74	$[-2.73, +1.24]$
Llama3B	$+0.68$	-5.76	-6.45	$[-9.21, -3.88]$
Smol3B	-5.00	-11.47	-6.46	$[-11.30, -2.48]$
Qwen8B	-0.34	-1.24	-0.90	$[-3.45, +1.50]$
Llama8B	-4.69	-8.98	-4.30	$[-7.98, -0.75]$

Table 5: Target accuracy endpoints and paired changes (pp). Intervals are descriptive joint run/item intervals.

Out-of-distribution accuracy. The OOD head tests transfer to MATH-500 or HumanEval.

Model	G_1	G_2	Δ	95% CI
Qwen4B	-1.73	-0.94	$+0.79$	$[-0.93, +2.58]$
Llama3B	-0.51	-4.39	-3.88	$[-6.56, -1.27]$
Smol3B	-5.30	-11.30	-6.00	$[-8.71, -3.42]$
Qwen8B	$+0.05$	-0.54	-0.59	$[-2.93, +2.02]$
Llama8B	-6.74	-9.11	-2.38	$[-5.27, +0.52]$

Table 6: Out-of-distribution accuracy endpoints and paired changes (pp). Intervals are descriptive joint run/item intervals.

Retention accuracy. The retention head evaluates the fixed MMLU anchor.

Model	G_1	G_2	Δ	95% CI
Qwen4B	-0.66	-1.03	-0.37	$[-0.98, +0.16]$
Llama3B	$+0.68$	$+0.68$	-0.01	$[-1.14, +1.14]$
Smol3B	$+0.09$	-0.24	-0.33	$[-1.35, +0.64]$
Qwen8B	-1.03	-1.59	-0.56	$[-1.75, +0.50]$
Llama8B	$+1.17$	$+0.34$	-0.83	$[-2.20, +0.47]$

Table 7: Retention accuracy endpoints and paired changes (pp). Intervals are descriptive joint run/item intervals.

Code task. Each model/task stratum uses its own fixed seed grid.

Model	G_1	G_2	Δ	95% CI
Qwen4B	$+1.73$	$+1.61$	-0.12	$[-1.13, +0.91]$
Llama3B	$+1.19$	-1.23	-2.42	$[-4.07, -0.86]$
Smol3B	-4.74	-9.60	-4.87	$[-6.80, -3.12]$
Qwen8B	-0.16	-0.77	-0.61	$[-1.81, +0.58]$
Llama8B	-3.94	-5.64	-1.70	$[-3.82, +0.24]$

Table 8: Code endpoints and paired continuation changes (pp). There are eight runs per smaller-model stratum and four per 8B stratum.

Math task. Each model/task stratum uses its own fixed seed grid.

Model	G_1	G_2	Δ	95% CI
Qwen4B	-2.23	-2.33	-0.10	$[-1.42, +1.34]$
Llama3B	-0.62	-5.09	-4.47	$[-6.58, -2.46]$
Smol3B	-2.08	-5.73	-3.66	$[-6.98, -0.76]$
Qwen8B	-0.72	-1.47	-0.75	$[-2.69, +1.15]$
Llama8B	-2.90	-6.20	-3.30	$[-5.81, -0.67]$

Table 9: Math endpoints and paired continuation changes (pp). There are eight runs per smaller-model stratum and four per 8B stratum.

C Matched Control Results

The recipient contrasts compare update increments from M_1 and M_0 on fixed donor corrections. D_0 denotes corrections generated by the unchanged parent on the control runtime; D_1 denotes archived candidate-generated corrections. Their interaction is the difference of these receiver contrasts. All paired receivers and native baselines are scored on the control runtime. The reference-source comparison matches training problems and receiving states, but not solution style or assistant-label token counts. Recorded completion lengths are descriptive; no missing finish reason is reconstructed as a known truncation event.

Contrast	Mean	Simultaneous 95% CI
Native first step	-0.44	$[-2.40, +1.36]$
Native continuation	-1.20	$[-3.06, +0.55]$
10 vs. 50 steps	$+0.47$	$[-0.90, +1.94]$
25 vs. 50 steps	-0.28	$[-1.37, +0.83]$
Pooled vs. two fixes	$+0.29$	$[-1.57, +2.04]$
Recipient, D_0	-0.47	$[-3.02, +2.03]$
Recipient, D_1	-1.10	$[-3.85, +1.72]$
Recipient $\times$ donor	-0.63	$[-2.79, +1.56]$
Gold vs. self-generated	-11.56	$[-14.45, -8.64]$

Table 10: The nine declared controls (pp). Positive recipe contrasts favor the alternative over its matched baseline.

Holm-adjusted $p = 0.0018$ for gold versus self-generated corrections, $p = 0.5407$ for the native continuation increment, and $p = 1.0$ for the other seven contrasts.

D Selection Optimism Diagnostic

This corrected-template diagnostic evaluates candidate selection under repeated continuation outcomes. Choosing and scoring a candidate on the same draw yields optimism of 0, 1.25, and 2.09 pp for $K = 1, 2, 3$. The corresponding cross-draw selection gaps are 0, 1.20, and 2.25 pp. Synthetic null and recoverable-signal controls check that the analysis can distinguish noise-only selection from a usable signal. These results motivate independent selection and scoring; they are not additional independent evidence for the fresh continuation contrast.

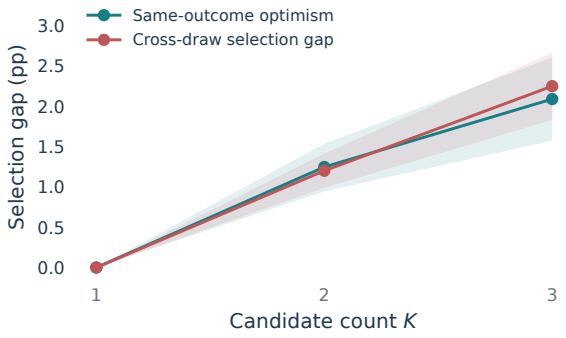

Figure 6: Selection gaps in an eight-parent diagnostic. Intervals resample parents conditional on its scored items. This diagnostic is separate from the fresh five-model estimates.

E Output Format and Length Audit

The audit reads saved completions and reuses the scoring parser. GSM8K failures mean that no unambiguous numeric answer could be extracted. Code failures mean that the extracted program, including HumanEval's required prefix where appli-

cable, fails Python syntax parsing. No code is executed again and no answer is repaired. MATH's symbolic extraction is not assigned a comparable diagnostic rate. Lengths count whitespace-separated words, rather than model tokens. Generation finish reasons were not recorded, so these observations cannot establish which outputs were truncated at the generation limit.

Model	Head	Parent	H_1	H_2
Qwen4B	Code T	0.00	0.83	1.01
Qwen4B	Code O	2.34	1.56	1.56
Qwen4B	Math T	4.69	3.22	4.82
Llama3B	Code T	1.56	1.95	2.10
Llama3B	Code O	0.78	1.22	1.79
Llama3B	Math T	2.34	4.49	10.19
Smol3B	Code T	0.00	1.95	5.50
Smol3B	Code O	0.78	2.54	6.22
Smol3B	Math T	3.12	4.69	15.12
Qwen8B	Code T	0.00	0.00	0.00
Qwen8B	Code O	1.56	1.56	1.43
Qwen8B	Math T	8.59	5.27	6.02
Llama8B	Code T	2.34	0.20	0.85
Llama8B	Code O	0.78	1.37	1.24
Llama8B	Math T	2.34	2.73	7.55

Table 11: Syntax or numeric-extraction failure rates (%). T denotes target and O denotes OOD.

The same development-selected candidate weights are used for these format rates and for utility.

Completion length. Word counts characterize output changes without inferring token-level truncation.

Model	Head	H_1 words	H_2 words
Qwen4B	Code T	44.3	47.0
Qwen4B	Code O	124.8	123.7
Qwen4B	Math T	147.8	144.9
Llama3B	Code T	73.4	85.5
Llama3B	Code O	149.9	141.9
Llama3B	Math T	149.2	164.1
Smol3B	Code T	130.9	159.8
Smol3B	Code O	185.5	190.1
Smol3B	Math T	147.4	187.6
Qwen8B	Code T	25.8	28.0
Qwen8B	Code O	146.2	139.4
Qwen8B	Math T	156.4	169.3
Llama8B	Code T	61.0	59.1
Llama8B	Code O	99.9	91.9
Llama8B	Math T	144.5	161.0

Table 12: Mean whitespace-separated completion lengths. These are descriptive output lengths, not generated-token counts.

F Statistical and Artifact Accounting

The main five-model analysis contains 64 fresh runs, 192 immediate candidates, 576 continuations, and 832 scored checkpoints. Each checkpoint has 128 target, 128 OOD, and 256 retention outcomes. Role weights are equal irrespective of item count. The 16 control parents overlap previously inspected runs and are reported separately.

Bootstrap resampling preserves parent strata and shared item identities. The three native draws are averaged before bootstrap resampling. Consequently, neither checkpoints nor native draws are used as independent parent observations. Primary reports retain their declared multiplicity families; the five-model plotting analysis is a separate descriptive re-expression of saved outcomes.

The accompanying source includes vector figures, aggregate plotting data, analysis code, and manifest/report checksums. The saved summaries reproduce the figures without further model execution.